

WORKING DRAFT

J3/06-007r1

25th September 2006 18:59

This is an internal working document of J3.

Contents

1	Overview	1
1.1	Scope	1
1.2	Processor	1
1.3	Inclusions	1
1.4	Exclusions	1
1.5	Conformance	2
1.6	Compatibility	3
1.6.1	New intrinsic procedures	3
1.6.2	New intrinsic data type and operator	3
1.6.3	Fortran 2003 compatibility	3
1.6.4	Fortran 95 compatibility	3
1.6.5	Fortran 90 compatibility	4
1.6.6	FORTTRAN 77 compatibility	4
1.7	Notation used in this part of ISO/IEC 1539	5
1.7.1	Applicability of requirements	5
1.7.2	Informative notes	5
1.7.3	Syntax rules	5
1.7.4	Constraints	6
1.7.5	Assumed syntax rules	6
1.7.6	Syntax conventions and characteristics	6
1.7.7	Text conventions	7
1.8	Deleted and obsolescent features	7
1.8.1	General	7
1.8.2	Nature of deleted features	7
1.8.3	Nature of obsolescent features	7
1.9	Normative references	7
2	Fortran terms and concepts	9
2.1	High level syntax	9
2.2	Program unit concepts	12
2.2.1	Program units and scoping units	12
2.2.2	Program	12
2.2.3	Main program	12
2.2.4	Procedure	12
2.2.5	Module	13
2.2.6	Submodule	13
2.3	Execution concepts	14
2.3.1	Statement classification	14
2.3.2	Program execution	14
2.3.3	Executable/nonexecutable statements	14
2.3.4	Statement order	14
2.3.5	The END statement	16
2.3.6	Execution sequence	16
2.4	Data concepts	17
2.4.1	Type	17
2.4.2	Data value	17
2.4.3	Data entity	17

2.4.4	Scalar	19
2.4.5	Array	19
2.4.6	Co-array	20
2.4.7	Pointer	20
2.4.8	Storage	20
2.5	Fundamental terms	21
2.5.1	Name and designator	21
2.5.2	Keyword	21
2.5.3	Association	21
2.5.4	Declaration	21
2.5.5	Definition	22
2.5.6	Reference	22
2.5.7	Intrinsic	22
2.5.8	Operator	22
2.5.9	Sequence	22
2.5.10	Companion processors	23
3	Lexical tokens, source form, and macro processing	25
3.1	Processor character set	25
3.1.1	Letters	25
3.1.2	Digits	25
3.1.3	Underscore	25
3.1.4	Special characters	26
3.1.5	Other characters	26
3.2	Low-level syntax	26
3.2.1	Names	26
3.2.2	Constants	27
3.2.3	Operators	27
3.2.4	Statement labels	28
3.2.5	Delimiters	29
3.3	Source form	29
3.3.1	Free source form	29
3.3.2	Fixed source form	31
3.4	Including source text	32
3.5	Macro processing	32
3.5.1	Macro definition	32
3.5.2	Macro expansion	35
4	Types	41
4.1	The concept of type	41
4.1.1	Set of values	41
4.1.2	Constants	41
4.1.3	Operations	41
4.2	Type parameters	41
4.3	Relationship of types and values to objects	43
4.3.1	Type specifiers and type compatibility	43
4.4	Intrinsic types	45
4.4.1	Classification and specification	45
4.4.2	Integer type	45
4.4.3	Real type	46
4.4.4	Complex type	48
4.4.5	Character type	49
4.4.6	Logical type	53
4.4.7	Bits type	53

4.5	Derived types	54
4.5.1	Derived type concepts	54
4.5.2	Derived-type definition	55
4.5.3	Derived-type parameters	59
4.5.4	Components	60
4.5.5	Type-bound procedures	68
4.5.6	Final subroutines	70
4.5.7	Type extension	72
4.5.8	Derived-type values	74
4.5.9	Derived-type specifier	74
4.5.10	Construction of derived-type values	75
4.5.11	Derived-type operations and assignment	77
4.6	Enumerations and enumerators	77
4.7	Construction of array values	79
5	Attribute declarations and specifications	83
5.1	General	83
5.2	Type declaration statements	83
5.2.1	Syntax	83
5.2.2	Automatic data objects	84
5.2.3	Initialization	85
5.2.4	Examples of type declaration statements	85
5.3	Attributes	85
5.3.1	Constraints	85
5.3.2	Accessibility attribute	86
5.3.3	ALLOCATABLE attribute	86
5.3.4	ASYNCHRONOUS attribute	86
5.3.5	BIND attribute for data entities	87
5.3.6	CONTIGUOUS attribute	87
5.3.7	DIMENSION attribute	88
5.3.8	EXTERNAL attribute	92
5.3.9	INTENT attribute	93
5.3.10	INTRINSIC attribute	95
5.3.11	OPTIONAL attribute	95
5.3.12	PARAMETER attribute	96
5.3.13	POINTER attribute	96
5.3.14	PROTECTED attribute	96
5.3.15	SAVE attribute	97
5.3.16	TARGET attribute	98
5.3.17	VALUE attribute	98
5.3.18	VOLATILE attribute	98
5.4	Attribute specification statements	99
5.4.1	Accessibility statement	99
5.4.2	ALLOCATABLE statement	100
5.4.3	ASYNCHRONOUS statement	100
5.4.4	BIND statement	100
5.4.5	CONTIGUOUS statement	100
5.4.6	DATA statement	100
5.4.7	DIMENSION statement	103
5.4.8	INTENT statement	103
5.4.9	OPTIONAL statement	103
5.4.10	PARAMETER statement	104
5.4.11	POINTER statement	104
5.4.12	PROTECTED statement	104

5.4.13	SAVE statement	104
5.4.14	TARGET statement	105
5.4.15	VALUE statement	105
5.4.16	VOLATILE statement	105
5.5	IMPLICIT statement	105
5.6	NAMelist statement	107
5.7	Storage association of data objects	108
5.7.1	EQUIVALENCE statement	108
5.7.2	COMMON statement	111
5.7.3	Restrictions on common and equivalence	113
6	Use of data objects	115
6.1	Scalars	116
6.1.1	Substrings	116
6.1.2	Structure components	117
6.1.3	Complex parts	119
6.1.4	Type parameter inquiry	119
6.2	Arrays	120
6.2.1	Whole arrays	120
6.2.2	Array elements and array sections	120
6.2.3	Image selectors	124
6.3	Dynamic association	124
6.3.1	ALLOCATE statement	124
6.3.2	NULLIFY statement	129
6.3.3	DEALLOCATE statement	129
7	Expressions and assignment	133
7.1	Expressions	133
7.1.1	Form of an expression	133
7.1.2	Intrinsic operations	137
7.1.3	Defined operations	138
7.1.4	Type, type parameters, and shape of an expression	139
7.1.5	Conformability rules for elemental operations	141
7.1.6	Specification expression	142
7.1.7	Initialization expression	143
7.1.8	Evaluation	145
7.2	Interpretation of operations	150
7.2.1	General	150
7.2.2	Numeric intrinsic operations	151
7.2.3	Character intrinsic operation	151
7.2.4	Relational intrinsic operations	152
7.2.5	Logical intrinsic operations	153
7.2.6	Bits intrinsic operations	154
7.3	Precedence of operators	154
7.4	Assignment	156
7.4.1	Assignment statement	156
7.4.2	Pointer assignment	162
7.4.3	Masked array assignment – WHERE	166
7.4.4	FORALL	168
8	Execution control	175
8.1	Executable constructs containing blocks	175
8.1.1	General	175
8.1.2	Rules governing blocks	175

8.1.3	ASSOCIATE construct	176
8.1.4	BLOCK construct	177
8.1.5	CASE construct	178
8.1.6	CRITICAL construct	180
8.1.7	DO construct	182
8.1.8	IF construct and statement	188
8.1.9	SELECT TYPE construct	189
8.1.10	EXIT statement	192
8.2	Branching	192
8.2.1	GO TO statement	192
8.2.2	Computed GO TO statement	193
8.2.3	Arithmetic IF statement	193
8.3	CONTINUE statement	193
8.4	STOP statement	193
8.5	Image execution control	195
8.5.1	Image control statements	195
8.5.2	SYNC ALL statement	197
8.5.3	SYNC TEAM statement	197
8.5.4	SYNC IMAGES statement	199
8.5.5	NOTIFY and QUERY statements	200
8.5.6	SYNC MEMORY statement	202
8.5.7	STAT= and ERRMSG= specifiers in image execution control statements	203
9	Input/output statements	205
9.1	Records	205
9.1.1	Formatted record	205
9.1.2	Unformatted record	206
9.1.3	Endfile record	206
9.2	External files	206
9.2.1	File existence	207
9.2.2	File access	207
9.2.3	File position	209
9.2.4	File storage units	211
9.3	Internal files	212
9.4	File connection	212
9.4.1	Connection modes	213
9.4.2	Unit existence	213
9.4.3	Connection of a file to a unit	214
9.4.4	Preconnection	215
9.4.5	OPEN statement	215
9.4.6	CLOSE statement	220
9.5	Data transfer statements	221
9.5.1	General	221
9.5.2	Control information list	222
9.5.3	Data transfer input/output list	227
9.5.4	Execution of a data transfer input/output statement	229
9.5.5	Termination of data transfer statements	241
9.6	Waiting on pending data transfer	241
9.6.1	Wait operation	241
9.6.2	WAIT statement	242
9.7	File positioning statements	243
9.7.1	Syntax	243
9.7.2	BACKSPACE statement	243
9.7.3	ENDFILE statement	244

9.7.4	REWIND statement	244
9.8	FLUSH statement	244
9.9	File inquiry statement	245
9.9.1	Forms of the INQUIRE statement	245
9.9.2	Inquiry specifiers	246
9.9.3	Inquire by output list	252
9.10	Error, end-of-record, and end-of-file conditions	252
9.10.1	Error conditions and the ERR= specifier	253
9.10.2	End-of-file condition and the END= specifier	253
9.10.3	End-of-record condition and the EOR= specifier	254
9.10.4	IOSTAT= specifier	254
9.10.5	IOMSG= specifier	255
9.11	Restrictions on input/output statements	255
10	Input/output editing	257
10.1	Format specifications	257
10.2	Explicit format specification methods	257
10.2.1	FORMAT statement	257
10.2.2	Character format specification	257
10.3	Form of a format item list	258
10.3.1	Syntax	258
10.3.2	Edit descriptors	258
10.3.3	Fields	260
10.4	Interaction between input/output list and format	260
10.5	Positioning by format control	262
10.6	Decimal symbol	262
10.7	Data edit descriptors	262
10.7.1	General	262
10.7.2	Numeric and bits editing	263
10.7.3	Logical editing	269
10.7.4	Character editing	270
10.7.5	Generalized editing	270
10.7.6	User-defined derived-type editing	272
10.8	Control edit descriptors	272
10.8.1	Position editing	273
10.8.2	Slash editing	274
10.8.3	Colon editing	274
10.8.4	SS, SP, and S editing	274
10.8.5	P editing	274
10.8.6	BN and BZ editing	275
10.8.7	RU, RD, RZ, RN, RC, and RP editing	275
10.8.8	DC and DP editing	275
10.9	Character string edit descriptors	276
10.10	List-directed formatting	276
10.10.1	General	276
10.10.2	Values and value separators	276
10.10.3	List-directed input	277
10.10.4	List-directed output	279
10.11	Namelist formatting	280
10.11.1	General	280
10.11.2	Name-value subsequences	280
10.11.3	Namelist input	280
10.11.4	Namelist output	284

11	Program units	287
11.1	Main program	287
11.2	Modules	288
11.2.1	General	288
11.2.2	The USE statement and use association	289
11.2.3	Submodules	291
11.3	Block data program units	292
12	Procedures	295
12.1	Concepts	295
12.2	Procedure classifications	295
12.2.1	Procedure classification by reference	295
12.2.2	Procedure classification by means of definition	295
12.3	Characteristics	296
12.3.1	Characteristics of procedures	296
12.3.2	Characteristics of dummy arguments	296
12.3.3	Characteristics of function results	297
12.4	Procedure interface	297
12.4.1	General	297
12.4.2	Implicit and explicit interfaces	297
12.4.3	Specification of the procedure interface	298
12.5	Procedure reference	308
12.5.1	Syntax	308
12.5.2	Actual arguments, dummy arguments, and argument association	310
12.5.3	Function reference	322
12.5.4	Subroutine reference	322
12.5.5	Resolving named procedure references	323
12.5.6	Resolving type-bound procedure references	325
12.6	Procedure definition	325
12.6.1	Intrinsic procedure definition	325
12.6.2	Procedures defined by subprograms	326
12.6.3	Definition and invocation of procedures by means other than Fortran	333
12.6.4	Statement function	333
12.7	Pure procedures	334
12.8	Elemental procedures	336
12.8.1	Elemental procedure declaration and interface	336
12.8.2	Elemental function actual arguments and results	336
12.8.3	Elemental subroutine actual arguments	337
13	Intrinsic procedures and modules	339
13.1	Classes of intrinsic procedures	339
13.2	Arguments to intrinsic procedures	339
13.2.1	General rules	339
13.2.2	The shape of array arguments	340
13.2.3	Mask arguments	340
13.2.4	Arguments to collective subroutines	340
13.3	Bit model	340
13.4	Numeric models	341
13.5	Standard generic intrinsic procedures	342
13.5.1	Numeric functions	342
13.5.2	Mathematical functions	343
13.5.3	Character functions	343
13.5.4	Kind functions	344
13.5.5	Miscellaneous type conversion functions	344

13.5.6	Numeric inquiry functions	344
13.5.7	Array inquiry functions	344
13.5.8	Other inquiry functions	345
13.5.9	Bit manipulation procedures	345
13.5.10	Floating-point manipulation functions	345
13.5.11	Vector and matrix multiply functions	346
13.5.12	Array reduction functions	346
13.5.13	Array construction functions	346
13.5.14	Array location functions	346
13.5.15	Collective subroutines	347
13.5.16	Null function	347
13.5.17	Allocation transfer procedure	347
13.5.18	Random number subroutines	347
13.5.19	System environment procedures	347
13.6	Specific names for standard intrinsic functions	348
13.7	Specifications of the standard intrinsic procedures	349
13.8	Standard intrinsic modules	434
13.8.1	General	434
13.8.2	The IEEE intrinsic modules	435
13.8.3	The ISO_FORTRAN_ENV intrinsic module	435
13.8.4	The ISO_C_BINDING intrinsic module	438
14	Exceptions and IEEE arithmetic	439
14.1	Derived types and constants defined in the modules	440
14.2	The exceptions	441
14.3	The rounding modes	442
14.4	Underflow mode	443
14.5	Halting	443
14.6	The floating point status	444
14.7	Exceptional values	444
14.8	IEEE arithmetic	444
14.9	Tables of the procedures	445
14.9.1	Inquiry functions	445
14.9.2	Elemental functions	446
14.9.3	Kind function	446
14.9.4	Elemental subroutines	446
14.9.5	Nonelemental subroutines	446
14.10	Specifications of the procedures	447
14.11	Examples	462
15	Interoperability with C	467
15.1	General	467
15.2	The ISO_C_BINDING intrinsic module	467
15.2.1	Summary of contents	467
15.2.2	Named constants and derived types in the module	467
15.2.3	Procedures in the module	468
15.3	Interoperability between Fortran and C entities	472
15.3.1	General	472
15.3.2	Interoperability of intrinsic types	472
15.3.3	Interoperability with C pointer types	474
15.3.4	Interoperability of derived types and C struct types	475
15.3.5	Interoperability of scalar variables	476
15.3.6	Interoperability of array variables	476
15.3.7	Interoperability of procedures and procedure interfaces	477

15.4	Interoperation with C global variables	479
15.4.1	General	479
15.4.2	Binding labels for common blocks and variables	480
15.5	Interoperation with C functions	480
15.5.1	Definition and reference of interoperable procedures	480
15.5.2	Binding labels for procedures	481
15.5.3	Exceptions and IEEE arithmetic procedures	481
16	Scope, association, and definition	483
16.1	Identifiers and entities	483
16.2	Scope of global identifiers	483
16.3	Scope of local identifiers	484
16.3.1	Classes of local identifiers	484
16.3.2	Local identifiers that are the same as common block names	485
16.3.3	Function results	485
16.3.4	Components, type parameters, and bindings	485
16.3.5	Argument keywords	486
16.4	Statement and construct entities	486
16.5	Association	487
16.5.1	Name association	487
16.5.2	Pointer association	491
16.5.3	Storage association	494
16.5.4	Inheritance association	498
16.5.5	Establishing associations	498
16.6	Definition and undefinition of variables	498
16.6.1	Definition of objects and subobjects	498
16.6.2	Variables that are always defined	499
16.6.3	Variables that are initially defined	499
16.6.4	Variables that are initially undefined	499
16.6.5	Events that cause variables to become defined	499
16.6.6	Events that cause variables to become undefined	501
16.6.7	Variable definition context	503
16.6.8	Pointer association context	504
Annex A	(informative)Glossary of technical terms	505
Annex B	(informative)Decremental features	519
B.1	Deleted features	519
B.2	Obsolescent features	520
B.2.1	Alternate return	520
B.2.2	Computed GO TO statement	520
B.2.3	Statement functions	520
B.2.4	DATA statements among executables	521
B.2.5	Assumed character length functions	521
B.2.6	Fixed form source	521
B.2.7	CHARACTER* form of CHARACTER declaration	521
Annex C	(informative)Extended notes	523
C.1	Clause 4 notes	523
C.1.1	Selection of the approximation methods (4.4.3)	523
C.1.2	Type extension and component accessibility (4.5.2.2, 4.5.4)	523
C.1.3	Abstract types	524
C.1.4	Pointers (4.5.2)	525
C.1.5	Structure constructors and generic names	526
C.1.6	Generic type-bound procedures	528

	C.1.7	Final subroutines (4.5.6, 4.5.6.2, 4.5.6.3, 4.5.6.4)	529
C.2		Clause 5 notes	531
	C.2.1	The POINTER attribute (5.3.13)	531
	C.2.2	The TARGET attribute (5.3.16)	532
	C.2.3	The VOLATILE attribute (5.3.18)	532
C.3		Clause 6 notes	533
	C.3.1	Structure components (6.1.2)	533
	C.3.2	Allocation with dynamic type (6.3.1)	535
	C.3.3	Pointer allocation and association	535
C.4		Clause 7 notes	536
	C.4.1	Character assignment	536
	C.4.2	Evaluation of function references	536
	C.4.3	Pointers in expressions	536
	C.4.4	Pointers on the left side of an assignment	537
	C.4.5	An example of a FORALL construct containing a WHERE construct	537
	C.4.6	Examples of FORALL statements	538
C.5		Clause 8 notes	539
	C.5.1	Loop control	539
	C.5.2	The CASE construct	539
	C.5.3	Examples of DO constructs	539
	C.5.4	Examples of invalid DO constructs	542
C.6		Clause 9 notes	543
	C.6.1	External files (9.2)	543
	C.6.2	Nonadvancing input/output (9.2.3.1)	544
	C.6.3	Asynchronous input/output	545
	C.6.4	OPEN statement (9.4.5)	546
	C.6.5	Connection properties (9.4.3)	548
	C.6.6	CLOSE statement (9.4.6)	548
C.7		Clause 10 notes	548
	C.7.1	Number of records (10.4, 10.5, 10.8.2)	548
	C.7.2	List-directed input (10.10.3)	549
C.8		Clause 11 notes	550
	C.8.1	Main program and block data program unit (11.1, 11.3)	550
	C.8.2	Dependent compilation (11.2)	550
	C.8.3	Examples of the use of modules	552
	C.8.4	Modules with submodules	558
C.9		Clause 12 notes	563
	C.9.1	Portability problems with external procedures (12.4.3.4)	563
	C.9.2	Procedures defined by means other than Fortran (12.6.3)	563
	C.9.3	Procedure interfaces (12.4)	564
	C.9.4	Abstract interfaces (12.4) and procedure pointer components (4.5)	564
	C.9.5	Argument association and evaluation (12.5.2)	566
	C.9.6	Pointers and targets as arguments (12.5.2.5, 12.5.2.7, 12.5.2.8)	567
	C.9.7	Polymorphic Argument Association (12.5.2.10)	568
C.10		Clause 13 notes	570
	C.10.1	Module for THIS_IMAGE and IMAGE_INDEX	570
	C.10.2	Collective co-array subroutine variations	570
C.11		Clause 15 notes	571
	C.11.1	Runtime environments	571
	C.11.2	Examples of Interoperation between Fortran and C Functions	571
C.12		Clause 16 notes	577
	C.12.1	Examples of host association (16.5.1.4)	577
	C.12.2	Rules ensuring unambiguous generics (12.4.3.3.4)	578
C.13		Array feature notes	582

- C.13.1 Summary of features 582
- C.13.2 Examples 583
- C.13.3 FORmula TRANslation and array processing 588
- C.13.4 Sum of squared residuals 589
- C.13.5 Vector norms: infinity-norm and one-norm 589
- C.13.6 Matrix norms: infinity-norm and one-norm 589
- C.13.7 Logical queries 589
- C.13.8 Parallel computations 590
- C.13.9 Example of element-by-element computation 590
- C.13.10 Bit manipulation and inquiry procedures 591

- Annex D (informative)Syntax rules 593
 - D.1 Extract of all syntax rules 593
 - D.2 Syntax rule cross-reference 638

- Annex E (informative)Index 651

List of Tables

2.1	Requirements on statement ordering	15
2.2	Statements allowed in scoping units	15
3.1	Special characters	26
6.1	Subscript order value	121
7.1	Type of operands and results for intrinsic operators	137
7.3	Interpretation of the numeric intrinsic operators	151
7.4	Interpretation of the character intrinsic operator //	151
7.5	Interpretation of the relational intrinsic operators	152
7.6	Interpretation of the logical intrinsic operators	153
7.7	The values of operations involving logical intrinsic operators	153
7.8	Interpretation of the bits intrinsic operators	154
7.9	The values of bits intrinsic operations other than //	154
7.10	Categories of operations and relative precedence	154
7.11	Type conformance for the intrinsic assignment statement	157
7.12	Numeric conversion and the assignment statement	160
7.13	Bits conversion and the assignment statement	160
10.1	E and D exponent forms	266
10.2	EN exponent forms	267
10.3	ES exponent forms	268
13.1	Characteristics of the result of NULL ()	410
15.1	Names of C characters with special semantics	468
15.2	Interoperability between Fortran and C types	472

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and nongovernmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of the joint technical committee is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 1539-1 was prepared by Joint Technical Committee ISO/IEC/JTC1, *Information technology*, Subcommittee SC22, *Programming languages, their environments and system software interfaces*.

This fifth edition cancels and replaces the fourth edition (ISO/IEC 1539-1:2004), which has been technically revised. It also incorporates the Technical Corrigenda ISO/IEC 1539-1:2004/Cor. 1:2005 and ISO/IEC 15391:2004/Cor. 2:2006.

ISO/IEC 1539 consists of the following parts, under the general title *Information technology — Programming languages — Fortran*:

- *Part 1: Base language*
- *Part 2: Varying length character strings*
- *Part 3: Conditional Compilation*

Introduction

International Standard programming language Fortran

This part of the International Standard comprises the specification of the base Fortran language, informally known as Fortran 2008. With the limitations noted in 1.6.3, the syntax and semantics of Fortran 2003 are contained entirely within Fortran 2008. Therefore, any standard-conforming Fortran 2003 program not affected by such limitations is a standard-conforming Fortran 2008 program. New features of Fortran 2008 can be compatibly incorporated into such Fortran 2003 programs, with any exceptions indicated in the text of this part of the standard.

Fortran 2008 contains several extensions to Fortran 2003; some of these are listed below.

- (1) The maximum rank of an array has been increased from seven to fifteen.
- (2) Performance enhancements: The DO CONCURRENT construct, which allows loop iterations to be executed in any order or potentially concurrently.
- (3) Pointers can be initialized to point to a target.
- (4) Performance enhancements: CONTIGUOUS attribute.
- (5) The ATAN intrinsic is extended so that ATAN (Y, X) is ATAN2 (Y,X).
- (6) Allocatable components of recursive type.
- (7) The MOLD= specifier has been added to the ALLOCATE statement.
- (8) OPEN statement enhancements that allow the processor to select a unit number when opening an external unit. Such a unit number is guaranteed not to interfere with any program-managed unit numbers.
- (9) The BLOCK construct (allows declarations within executable statements).
- (10) A disassociated or deallocated actual argument can correspond to an optional nonpointer nonallocatable dummy argument.
- (11) The concept of *variable* now includes references to pointer functions which return associated pointers.
- (12) The COMPILER_VERSION and COMPILER_OPTIONS functions provide information about the translation phase of the execution of a program.
- (13) The real and imaginary parts of a COMPLEX variable can be selected using a component-like syntax.
- (14) Scoped macros which can generate whole Fortran statements and subprograms.
- (15) A FINDLOC intrinsic was added and a BACK= argument was also added to MAXLOC and MINLOC.
- (16) Parallel programming support: SPMD parallel programming, co-arrays for data exchange between images, image control statements, and collective procedures.
- (17) A BITS data type for non-numeric programming and enhanced handling of BOZ constants.
- (18) The G0 edit descriptor.
- (19) Additional mathematical intrinsic functions for computing Bessel functions, the error function, the Gamma function, and generalized L_2 norms.

J3 internal note

Unresolved Technical Issue 080

The laundry list needs to be redone at a later time. RAH suggests going down the list of things in spread sheet and having a big feature and a little feature list.

Organization of this part of ISO/IEC 1539

This part of ISO/IEC 1539 is organized in 16 clauses, dealing with 8 conceptual areas. These 8 areas, and the clauses in which they are treated, are:

High/low level concepts	Clauses 1, 2, 3
Data concepts	Clauses 4, 5, 6
Computations	Clauses 7, 13, 14
Execution control	Clause 8
Input/output	Clauses 9, 10
Program units	Clauses 11, 12
Interoperability with C	Clause 15
Scoping and association rules	Clause 16

It also contains the following nonnormative material:

Glossary	A
Decremental features	B
Extended notes	C
Syntax rules	D
Index	E

1 Information technology — Programming languages — 2 Fortran —

3 Part 1: 4 Base Language

5 1 Overview

6 1.1 Scope

7 ISO/IEC 1539 is a multipart International Standard; the parts are published separately. This publi-
8 cation, ISO/IEC 1539-1, which is the first part, specifies the form and establishes the interpretation
9 of programs expressed in the base Fortran language. The purpose of this part of ISO/IEC 1539 is to
10 promote portability, reliability, maintainability, and efficient execution of Fortran programs for use on
11 a variety of computing systems. The second part, ISO/IEC 1539-2, defines additional facilities for the
12 manipulation of character strings of variable length; this has been largely subsumed by allocatable char-
13 acters with deferred length parameters. The third part, ISO/IEC 1539-3, defines a standard conditional
14 compilation facility for Fortran. A processor conforming to part 1 need not conform to ISO/IEC 1539-2
15 or ISO/IEC 1539-3; however, conformance to either assumes conformance to this part. Throughout this
16 publication, the term “this standard” refers to ISO/IEC 1539-1.

17 1.2 Processor

18 The combination of a computing system and the mechanism by which programs are transformed for use
19 on that computing system is called a **processor** in this part of ISO/IEC 1539.

20 1.3 Inclusions

21 This part of ISO/IEC 1539 specifies

- 22 (1) the forms that a program written in the Fortran language may take,
- 23 (2) the rules for interpreting the meaning of a program and its data,
- 24 (3) the form of the input data to be processed by such a program, and
- 25 (4) the form of the output data resulting from the use of such a program.

26 1.4 Exclusions

27 This part of ISO/IEC 1539 does not specify

- 28 (1) the mechanism by which programs are transformed for use on computing systems,
- 29 (2) the operations required for setup and control of the use of programs on computing systems,
- 30 (3) the method of transcription of programs or their input or output data to or from a storage
31 medium,

- 1 (4) the program and processor behavior when this part of ISO/IEC 1539 fails to establish an
2 interpretation except for the processor detection and reporting requirements in items (2) to
3 (8) of 1.5,
- 4 (5) the size or complexity of a program and its data that will exceed the capacity of any
5 particular computing system or the capability of a particular processor,
- 6 (6) the physical properties of the representation of quantities and the method of rounding,
7 approximating, or computing numeric values on a particular processor,
- 8 (7) the physical properties of input/output records, files, and units, or
- 9 (8) the physical properties and implementation of storage.

10 1.5 Conformance

11 A program (2.2.2) is a **standard-conforming program** if it uses only those forms and relationships
12 described herein and if the program has an interpretation according to this part of ISO/IEC 1539. A
13 program unit (2.2.1) conforms to this part of ISO/IEC 1539 if it can be included in a program in a
14 manner that allows the program to be standard conforming.

15 A processor conforms to this part of ISO/IEC 1539 if:

- 16 (1) it executes any standard-conforming program in a manner that fulfills the interpretations
17 herein, subject to any limits that the processor may impose on the size and complexity of
18 the program;
- 19 (2) it contains the capability to detect and report the use within a submitted program unit of
20 a form designated herein as obsolescent, insofar as such use can be detected by reference to
21 the numbered syntax rules and constraints;
- 22 (3) it contains the capability to detect and report the use within a submitted program unit of
23 an additional form or relationship that is not permitted by the numbered syntax rules or
24 constraints, including the deleted features described in Annex B;
- 25 (4) it contains the capability to detect and report the use within a submitted program unit of
26 an intrinsic type with a kind type parameter value not supported by the processor (4.4);
- 27 (5) it contains the capability to detect and report the use within a submitted program unit of
28 source form or characters not permitted by Clause 3;
- 29 (6) it contains the capability to detect and report the use within a submitted program of name
30 usage not consistent with the scope rules for names, labels, operators, and assignment
31 symbols in Clause 16;
- 32 (7) it contains the capability to detect and report the use within a submitted program unit of
33 intrinsic procedures whose names are not defined in Clause 13; and
- 34 (8) it contains the capability to detect and report the reason for rejecting a submitted program.

35 However, in a format specification that is not part of a FORMAT statement (10.2.1), a processor need not
36 detect or report the use of deleted or obsolescent features, or the use of additional forms or relationships.

37 A standard-conforming processor may allow additional forms and relationships provided that such ad-
38 ditions do not conflict with the standard forms and relationships. However, a standard-conforming
39 processor may allow additional intrinsic procedures even though this could cause a conflict with the
40 name of a procedure in a standard-conforming program. If such a conflict occurs and involves the name
41 of an external procedure, the processor is permitted to use the intrinsic procedure unless the name is
42 given the EXTERNAL attribute (5.3.8) in the scoping unit (2.2.1). A standard-conforming program
43 shall not use nonstandard intrinsic procedures or modules that have been added by the processor.

44 Because a standard-conforming program may place demands on a processor that are not within the
45 scope of this part of ISO/IEC 1539 or may include standard items that are not portable, such as
46 external procedures defined by means other than Fortran, conformance to this part of ISO/IEC 1539

- 1 does not ensure that a program will execute consistently on all or any standard-conforming processors.
- 2 In some cases, this part of ISO/IEC 1539 allows the provision of facilities that are not completely specified
3 in the standard. These facilities are identified as **processor dependent**. They shall be provided, with
4 methods or semantics determined by the processor.

NOTE 1.1

The processor should be accompanied by documentation that specifies the limits it imposes on the size and complexity of a program and the means of reporting when these limits are exceeded, that defines the additional forms and relationships it allows, and that defines the means of reporting the use of additional forms and relationships and the use of deleted or obsolescent forms. In this context, the use of a deleted form is the use of an additional form.

The processor should be accompanied by documentation that specifies the methods or semantics of processor-dependent facilities.

5 **1.6 Compatibility**

6 **1.6.1 New intrinsic procedures**

- 7 Each Fortran International Standard since ISO 1539:1980 (informally referred to as FORTRAN 77), defines
8 more intrinsic procedures than the previous one. Therefore, a Fortran program conforming to an older
9 standard may have a different interpretation under a newer standard if it invokes an external procedure
10 having the same name as one of the new standard intrinsic procedures, unless that procedure is specified
11 to have the EXTERNAL attribute.

12 **1.6.2 New intrinsic data type and operator**

- 13 This part of ISO/IEC 1539 specifies a new intrinsic type, BITS, which will conflict with a derived type
14 of the same name.

- 15 This part of ISO/IEC 1539 specifies a new intrinsic operator, .XOR., which might conflict with a user-
16 defined operator of the same name, has a different precedence from that of a user-defined operator, and
17 a different syntax from that of a user-defined unary operator.

18 **1.6.3 Fortran 2003 compatibility**

- 19 Except as identified in this subclause, this part of ISO/IEC 1539 is an upward compatible extension
20 to the preceding Fortran International Standard, ISO/IEC 1539-1:2004 (Fortran 2003). Any standard-
21 conforming Fortran 2003 program that does not use a derived type called BITS, and does not use a
22 user-defined operator called .XOR., remains standard-conforming under this part of ISO/IEC 1539.

23 **1.6.4 Fortran 95 compatibility**

- 24 Except as identified in this subclause, this part of ISO/IEC 1539 is an upward compatible extension to
25 ISO/IEC 1539-1:1997 (Fortran 95). Any standard-conforming Fortran 95 program that does not use a
26 derived type called BITS or a user-defined operator called .XOR. remains standard-conforming under
27 this part of ISO/IEC 1539. The following Fortran 95 features may have different interpretations in this
28 part of ISO/IEC 1539.

- 29 (1) Earlier Fortran standards had the concept of printing, meaning that column one of formatted
30 output had special meaning for a processor-dependent (possibly empty) set of external
31 files. This could be neither detected nor specified by a standard-specified means. The
32 interpretation of the first column is not specified by this part of ISO/IEC 1539.

- 1 (2) This part of ISO/IEC 1539 specifies a different output format for real zero values in list-
2 directed and namelist output.
- 3 (3) If the processor can distinguish between positive and negative real zero, this part of ISO/IEC
4 1539 requires different returned values for ATAN2(Y,X) when $X < 0$ and Y is negative real
5 zero and for LOG(X) and SQRT(X) when X is complex with $\text{REAL}(X) < 0$ and negative
6 zero imaginary part.

7 **1.6.5 Fortran 90 compatibility**

8 Except for the deleted features noted in Annex B.1, and except as identified in this subclause, this part
9 of ISO/IEC 1539 is an upward compatible extension to ISO/IEC 1539:1991 (Fortran 90). Any standard-
10 conforming Fortran 90 program that does not use a derived type called BITS, a user-defined operator
11 called .XOR., or one of the deleted features remains standard-conforming under this part of ISO/IEC
12 1539.

13 The PAD= specifier in the INQUIRE statement in this part of ISO/IEC 1539 returns the value UNDE-
14 FINED if there is no connection or the connection is for unformatted input/output. Fortran 90 specified
15 YES.

16 Fortran 90 specified that if the second argument to MOD or MODULO was zero, the result was processor
17 dependent. this part of ISO/IEC 1539 specifies that the second argument shall not be zero.

18 **1.6.6 FORTRAN 77 compatibility**

19 Except for the deleted features noted in Annex B.1, and except as identified in this subclause, this part
20 of ISO/IEC 1539 is an upward compatible extension to ISO 1539:1980 (FORTRAN 77). Any standard-
21 conforming FORTRAN 77 program that does not use one of the deleted features noted in Annex B.1 and
22 that does not depend on the differences specified here remains standard-conforming under this part of
23 ISO/IEC 1539. This part of ISO/IEC 1539 restricts the behavior for some features that were processor
24 dependent in FORTRAN 77. Therefore, a standard-conforming FORTRAN 77 program that uses one of
25 these processor-dependent features may have a different interpretation under this part of ISO/IEC 1539,
26 yet remain a standard-conforming program. The following FORTRAN 77 features may have different
27 interpretations in this part of ISO/IEC 1539.

- 28 (1) FORTRAN 77 permitted a processor to supply more precision derived from a real constant
29 than can be represented in a real datum when the constant is used to initialize a data object
30 of type double precision real in a DATA statement. This part of ISO/IEC 1539 does not
31 permit a processor this option.
- 32 (2) If a named variable that was not in a common block was initialized in a DATA statement and
33 did not have the SAVE attribute specified, FORTRAN 77 left its SAVE attribute processor
34 dependent. This part of ISO/IEC 1539 specifies (5.4.6) that this named variable has the
35 SAVE attribute.
- 36 (3) FORTRAN 77 specified that the number of characters required by the input list was to be
37 less than or equal to the number of characters in the record during formatted input. This
38 part of ISO/IEC 1539 specifies (9.5.4.4.2) that the input record is logically padded with
39 blanks if there are not enough characters in the record, unless the PAD= specifier with the
40 value 'NO' is specified in an appropriate OPEN or READ statement.
- 41 (4) A value of 0 for a list item in a formatted output statement will be formatted in a different
42 form for some G edit descriptors. In addition, this part of ISO/IEC 1539 specifies how
43 rounding of values will affect the output field form, but FORTRAN 77 did not address this
44 issue. Therefore, some FORTRAN 77 processors may produce an output form different from
45 the output form produced by Fortran 2003 processors for certain combinations of values and
46 G edit descriptors.
- 47 (5) If the processor can distinguish between positive and negative real zero, the behavior of the

1 SIGN intrinsic function when the second argument is negative real zero is changed by this
2 part of ISO/IEC 1539.

3 **1.7 Notation used in this part of ISO/IEC 1539**

4 **1.7.1 Applicability of requirements**

5 In this part of ISO/IEC 1539, “shall” is to be interpreted as a requirement; conversely, “shall not” is
6 to be interpreted as a prohibition. Except where stated otherwise, such requirements and prohibitions
7 apply to programs rather than processors.

8 **1.7.2 Informative notes**

9 Informative notes of explanation, rationale, examples, and other material are interspersed with the
10 normative body of this part of ISO/IEC 1539. The informative material is nonnormative; it is identified
11 by being in shaded, framed boxes that have numbered headings beginning with “NOTE.”

12 **1.7.3 Syntax rules**

13 Syntax rules describe the forms that Fortran lexical tokens, statements, and constructs may take. These
14 syntax rules are expressed in a variation of Backus-Naur form (BNF) with the following conventions.

- 15 (1) Characters from the Fortran character set (3.1) are interpreted literally as shown, except
16 where otherwise noted.
- 17 (2) Lower-case italicized letters and words (often hyphenated and abbreviated) represent gen-
18 eral syntactic classes for which particular syntactic entities shall be substituted in actual
19 statements.

20 Common abbreviations used in syntactic terms are:

<i>arg</i>	for	argument	<i>attr</i>	for	attribute
<i>decl</i>	for	declaration	<i>def</i>	for	definition
<i>desc</i>	for	descriptor	<i>expr</i>	for	expression
<i>int</i>	for	integer	<i>op</i>	for	operator
<i>spec</i>	for	specifier	<i>stmt</i>	for	statement

- 21 (3) The syntactic metasymbols used are:

is	introduces a syntactic class definition
or	introduces a syntactic class alternative
[]	encloses an optional item
[] ...	encloses an optionally repeated item that may occur zero or more times
■	continues a syntax rule

- 22 (4) Each syntax rule is given a unique identifying number of the form Rsnn, where s is a
23 one- or two-digit clause number and nn is a two-digit sequence number within that clause.
24 The syntax rules are distributed as appropriate throughout the text, and are referenced by
25 number as needed. Some rules in Clauses 2 and 3 are more fully described in later clauses; in
26 such cases, the clause number s is the number of the later clause where the rule is repeated.
- 27 (5) The syntax rules are not a complete and accurate syntax description of Fortran, and cannot
28 be used to generate a Fortran parser automatically; where a syntax rule is incomplete, it is
29 restricted by corresponding constraints and text.

NOTE 1.2

An example of the use of the syntax rules is:

digit-string **is** *digit* [*digit*] ...

The following are examples of forms for a digit string allowed by the above rule:

digit
digit digit
digit digit digit digit
digit digit digit digit digit digit digit digit

If particular entities are substituted for *digit*, actual digit strings might be:

4
67
1999
10243852

1 1.7.4 Constraints

2 Each constraint is given a unique identifying number of the form Csn, where s is a one or two digit clause
3 number and nn is a two or three digit sequence number within that clause.

4 Often a constraint is associated with a particular syntax rule. Where that is the case, the constraint is
5 annotated with the syntax rule number in parentheses. A constraint that is associated with a syntax
6 rule constitutes part of the definition of the syntax term defined by the rule. It thus applies in all places
7 where the syntax term appears.

8 Some constraints are not associated with particular syntax rules. The effect of such a constraint is similar
9 to that of a restriction stated in the text, except that a processor is required to have the capability to
10 detect and report violations of constraints (1.5). In some cases, a broad requirement is stated in text
11 and a subset of the same requirement is also stated as a constraint. This indicates that a standard-
12 conforming program is required to adhere to the broad requirement, but that a standard-conforming
13 processor is required only to have the capability of diagnosing violations of the constraint.

14 1.7.5 Assumed syntax rules

15 In order to minimize the number of additional syntax rules and convey appropriate constraint informa-
16 tion, the following rules are assumed.

17 R101 *xyz-list* **is** *xyz* [, *xyz*] ...
18 R102 *xyz-name* **is** *name*
19 R103 *scalar-xyz* **is** *xyz*

20 C101 (R103) *scalar-xyz* shall be scalar.

21 The letters “*xyz*” stand for any syntactic class phrase. An explicit syntax rule for a term overrides an
22 assumed rule.

23 1.7.6 Syntax conventions and characteristics

24 (1) Any syntactic class name ending in “-*stmt*” follows the source form statement rules: it shall
25 be delimited by end-of-line or semicolon, and may be labeled unless it forms part of another
26 statement (such as an IF or WHERE statement). Conversely, everything considered to be
a source form statement is given a “-*stmt*” ending in the syntax rules.

- 1 (2) The rules on statement ordering are described rigorously in the definition of *program-unit*
2 (R202). Expression hierarchy is described rigorously in the definition of *expr* (R722).
- 3 (3) The suffix “-*spec*” is used consistently for specifiers, such as input/output statement speci-
4 fiers. It also is used for type declaration attribute specifications (for example, “*array-spec*”
5 in R510), and in a few other cases.
- 6 (4) Where reference is made to a type parameter, including the surrounding parentheses, the
7 suffix “-*selector*” is used. See, for example, “*kind-selector*” (R405) and “*length-selector*”
8 (R421).
- 9 (5) The term “*subscript*” (for example, R619, R620, and R621) is used consistently in array
10 definitions.

11 1.7.7 Text conventions

12 In the descriptive text, an equivalent English word is frequently used in place of a syntactic term.
13 Particular statements and attributes are identified in the text by an upper-case keyword, e.g., “END
14 statement”. Boldface words are used in the text where they are first defined with a specialized meaning.
15 The descriptions of obsolescent features appear in a smaller type size.

NOTE 1.3

This sentence is an example of the type size used for obsolescent features.

16 1.8 Deleted and obsolescent features

17 1.8.1 General

18 This part of ISO/IEC 1539 protects the users’ investment in existing software by including all but five
19 of the language elements of Fortran 90 that are not processor dependent. This part of ISO/IEC 1539
20 identifies two categories of outmoded features. There are five in the first category, **deleted features**,
21 which consists of features considered to have been redundant in FORTRAN 77 and largely unused in
22 Fortran 90. Those in the second category, **obsolescent features**, are considered to have been redundant
23 in Fortran 90 and Fortran 95, but are still frequently used.

24 1.8.2 Nature of deleted features

25 Better methods existed in FORTRAN 77 for each deleted feature. These features were not included in
26 Fortran 95 or Fortran 2003, and are not included in this revision of Fortran.

27 1.8.3 Nature of obsolescent features

28 Better methods existed in Fortran 90 and Fortran 95 for each obsolescent feature. It is recommended
29 that programmers use these better methods in new programs and convert existing code to these methods.

30 The obsolescent features are identified in the text of this part of ISO/IEC 1539 by a distinguishing type
31 font (1.7.7).

32 A future revision of this part of ISO/IEC 1539 might delete an obsolescent feature if its use has become
33 insignificant.

34 1.9 Normative references

35 The following referenced standards are indispensable for the application of this part of ISO/IEC 1539.
36 For dated references, only the edition cited applies. For undated references, the latest edition of the

- 1 referenced standard (including any amendments) applies.
- 2 ISO/IEC 646:1991, *Information technology—ISO 7-bit coded character set for information interchange*.
- 3 ISO 8601:1988, *Data elements and interchange formats—Information interchange—*
4 *Representation of dates and times*.
- 5 ISO/IEC 9899:1999, *Information technology—Programming languages—C*.
- 6 ISO/IEC 10646-1:2000, *Information technology—Universal multiple-octet coded character set (UCS)—*
7 *Part 1: Architecture and basic multilingual plane*.
- 8 IEC 60559 (1989-01), *Binary floating-point arithmetic for microprocessor systems*.
- 9 ISO/IEC 646:1991 (International Reference Version) is the international equivalent of ANSI X3.4-1986,
10 commonly known as ASCII.
- 11 This part of ISO/IEC 1539 refers to ISO/IEC 9899:1999 as the C International Standard.
- 12 Because IEC 60559 (1989-01) was originally IEEE 754-1985, *Standard for binary floating-point arith-*
13 *metic*, and is widely known by this name, this standard refers to it as the IEEE International Standard.

1 2 Fortran terms and concepts

2 2.1 High level syntax

3 This clause introduces the terms associated with program units and other Fortran concepts above the
 4 construct, statement, and expression levels and illustrates their relationships. The notation used in this
 5 part of ISO/IEC 1539 is described in 1.7.

NOTE 2.1

Constraints and other information related to the rules that do not begin with R2 appear in the appropriate clause.

6 R201 *program* **is** *program-unit*
 7 [*program-unit*] ...

8 A *program* shall contain exactly one *main-program program-unit* or a main program defined by means
 9 other than Fortran, but not both.

10 R202 *program-unit* **is** *main-program*
 11 **or** *external-subprogram*
 12 **or** *module*
 13 **or** *submodule*
 14 **or** *block-data*
 15 R1101 *main-program* **is** [*program-stmt*]
 16 [*specification-part*]
 17 [*execution-part*]
 18 [*internal-subprogram-part*]
 19 *end-program-stmt*
 20 R203 *external-subprogram* **is** *function-subprogram*
 21 **or** *subroutine-subprogram*
 22 R1225 *function-subprogram* **is** *function-stmt*
 23 [*specification-part*]
 24 [*execution-part*]
 25 [*internal-subprogram-part*]
 26 *end-function-stmt*
 27 R1233 *subroutine-subprogram* **is** *subroutine-stmt*
 28 [*specification-part*]
 29 [*execution-part*]
 30 [*internal-subprogram-part*]
 31 *end-subroutine-stmt*
 32 R1104 *module* **is** *module-stmt*
 33 [*specification-part*]
 34 [*module-subprogram-part*]
 35 *end-module-stmt*
 36 R1116 *submodule* **is** *submodule-stmt*
 37 [*specification-part*]
 38 [*module-subprogram-part*]
 39 *end-submodule-stmt*
 40 R1120 *block-data* **is** *block-data-stmt*
 41 [*specification-part*]

1			<i>end-block-data-stmt</i>
2	R204	<i>specification-part</i>	is [<i>use-stmt</i>] ...
3			[<i>import-stmt</i>] ...
4			[<i>implicit-part</i>]
5			[<i>declaration-construct</i>] ...
6	R205	<i>implicit-part</i>	is [<i>implicit-part-stmt</i>] ...
7			<i>implicit-stmt</i>
8	R206	<i>implicit-part-stmt</i>	is <i>implicit-stmt</i>
9			or <i>parameter-stmt</i>
10			or <i>format-stmt</i>
11			or <i>entry-stmt</i>
12	R207	<i>declaration-construct</i>	is <i>derived-type-def</i>
13			or <i>entry-stmt</i>
14			or <i>enum-def</i>
15			or <i>format-stmt</i>
16			or <i>interface-block</i>
17			or <i>macro-definition</i>
18			or <i>parameter-stmt</i>
19			or <i>procedure-declaration-stmt</i>
20			or <i>specification-stmt</i>
21			or <i>type-declaration-stmt</i>
22			or <i>stmt-function-stmt</i>
23	R208	<i>execution-part</i>	is <i>executable-construct</i>
24			[<i>execution-part-construct</i>] ...
25	R209	<i>execution-part-construct</i>	is <i>executable-construct</i>
26			or <i>format-stmt</i>
27			or <i>entry-stmt</i>
28			or <i>data-stmt</i>
29	R210	<i>internal-subprogram-part</i>	is <i>contains-stmt</i>
30			[<i>internal-subprogram</i>] ...
31	R211	<i>internal-subprogram</i>	is <i>function-subprogram</i>
32			or <i>subroutine-subprogram</i>
33	R1107	<i>module-subprogram-part</i>	is <i>contains-stmt</i>
34			[<i>module-subprogram</i>] ...
35	R1108	<i>module-subprogram</i>	is <i>function-subprogram</i>
36			or <i>subroutine-subprogram</i>
37			or <i>separate-module-subprogram</i>
38	R212	<i>specification-stmt</i>	is <i>access-stmt</i>
39			or <i>allocatable-stmt</i>
40			or <i>asynchronous-stmt</i>
41			or <i>bind-stmt</i>
42			or <i>common-stmt</i>
43			or <i>data-stmt</i>
44			or <i>dimension-stmt</i>
45			or <i>equivalence-stmt</i>
46			or <i>external-stmt</i>
47			or <i>intent-stmt</i>
48			or <i>intrinsic-stmt</i>
49			or <i>namelist-stmt</i>
50			or <i>optional-stmt</i>
51			or <i>pointer-stmt</i>
52			or <i>protected-stmt</i>
53			or <i>save-stmt</i>
54			or <i>target-stmt</i>

1		or <i>volatile-stmt</i>
2		or <i>value-stmt</i>
3	R213	<i>executable-construct</i>
4		is <i>action-stmt</i>
5		or <i>associate-construct</i>
6		or <i>block-construct</i>
7		or <i>case-construct</i>
8		or <i>critical-construct</i>
9		or <i>do-construct</i>
10		or <i>forall-construct</i>
11		or <i>if-construct</i>
12		or <i>select-type-construct</i>
13	R214	<i>action-stmt</i>
14		is <i>allocate-stmt</i>
15		or <i>assignment-stmt</i>
16		or <i>backspace-stmt</i>
17		or <i>call-stmt</i>
18		or <i>close-stmt</i>
19		or <i>continue-stmt</i>
20		or <i>cycle-stmt</i>
21		or <i>deallocate-stmt</i>
22		or <i>endfile-stmt</i>
23		or <i>end-function-stmt</i>
24		or <i>end-program-stmt</i>
25		or <i>end-subroutine-stmt</i>
26		or <i>exit-stmt</i>
27		or <i>flush-stmt</i>
28		or <i>forall-stmt</i>
29		or <i>goto-stmt</i>
30		or <i>if-stmt</i>
31		or <i>inquire-stmt</i>
32		or <i>notify-stmt</i>
33		or <i>nullify-stmt</i>
34		or <i>open-stmt</i>
35		or <i>pointer-assignment-stmt</i>
36		or <i>print-stmt</i>
37		or <i>query-stmt</i>
38		or <i>read-stmt</i>
39		or <i>return-stmt</i>
40		or <i>rewind-stmt</i>
41		or <i>stop-stmt</i>
42		or <i>sync-all-stmt</i>
43		or <i>sync-images-stmt</i>
44		or <i>sync-memory-stmt</i>
45		or <i>sync-team-stmt</i>
46		or <i>wait-stmt</i>
47		or <i>where-stmt</i>
48		or <i>write-stmt</i>
49		or <i>arithmetic-if-stmt</i>
		or <i>computed-goto-stmt</i>
50	C201	(R208) An <i>execution-part</i> shall not contain an <i>end-function-stmt</i> , <i>end-program-stmt</i> , or <i>end-</i>
51		<i>subroutine-stmt</i> .
52		Additionally, an EXPAND statement may occur anywhere that any statement may occur other than

1 as the first statement of a program unit. The syntax rules are applied to the program after macro
2 expansion, i.e. with each EXPAND statement replaced by the statements it produces.

3 2.2 Program unit concepts

4 2.2.1 Program units and scoping units

5 Program units are the fundamental components of a Fortran program. A **program unit** may be
6 a main program, an external subprogram, a module, a submodule, or a block data program unit. A
7 **subprogram** may be a function subprogram or a subroutine subprogram. A module contains definitions
8 that are to be made accessible to other program units. A submodule is an extension of a module; it may
9 contain the definitions of procedures declared in a module or another submodule. A block data program
10 unit is used to specify initial values for data objects in named common blocks. Each type of program
11 unit is described in Clause 11 or 12. An **external subprogram** is a subprogram that is not in a main
12 program, a module, a submodule, or another subprogram. An **internal subprogram** is a subprogram
13 that is in a main program or another subprogram. A **module subprogram** is a subprogram that is in
14 a module or submodule but is not an internal subprogram.

15 A program unit consists of a set of nonoverlapping scoping units. A **scoping unit** is

- 16 (1) a program unit or subprogram, excluding any scoping units in it,
- 17 (2) a derived-type definition (4.5.2), or
- 18 (3) an interface body, excluding any scoping units in it.

19 A scoping unit that immediately surrounds another scoping unit is called the **host scoping unit** (often
20 abbreviated to **host**). A module or submodule is also the host scoping unit of its child submodules.

21 2.2.2 Program

22 A **program** consists of exactly one main program, any number (including zero) of other kinds of program
23 units, and any number (including zero) of external procedures and other entities defined by means other
24 than Fortran.

NOTE 2.2

There is a restriction that there shall be no more than one unnamed block data program unit (11.3).

This part of ISO/IEC 1539 places no ordering requirement on the program units that constitute a program, but because the public portions of a module are required to be available by the time a module reference (11.2.2) is processed, a processor may require a particular order of processing of the program units.

25 2.2.3 Main program

26 The Fortran main program is described in 11.1.

27 2.2.4 Procedure

28 A **procedure** encapsulates an arbitrary sequence of actions that may be invoked directly during program
29 execution. Procedures are either functions or subroutines. A **function** is a procedure that is invoked
30 in an expression; its invocation causes a value to be computed which is then used in evaluating the
31 expression. The variable that returns the value of a function is called the **result variable**. A **subroutine**
32 is a procedure that is invoked in a CALL statement, by a defined assignment statement, or by some
33 operations on derived-type entities. Unless it is a pure procedure, a subroutine may be used to change

1 the program state by changing the values of any of the data objects accessible to the subroutine; unless
2 it is a pure procedure, a function may do this in addition to computing the function value.

3 Procedures are described further in Clause 12.

4 **2.2.4.1 External procedure**

5 An **external procedure** is a procedure that is defined by an external subprogram or by means other
6 than Fortran. An external procedure may be invoked by the main program or by any procedure of a
7 program.

8 **2.2.4.2 Module procedure**

9 A **module procedure** is a procedure that is defined by a module subprogram (R1108). The module or
10 submodule containing the subprogram is the **host** scoping unit of the module procedure.

11 **2.2.4.3 Internal procedure**

12 An **internal procedure** is a procedure that is defined by an internal subprogram (R211). The containing
13 main program or subprogram is the **host** scoping unit of the internal procedure. An internal procedure
14 is local to its host in the sense that the internal procedure is accessible within the host scoping unit and
15 all its other internal procedures but is not accessible elsewhere.

16 **2.2.4.4 Interface block**

17 An **interface body** describes an abstract interface or the interface of a dummy procedure, external
18 procedure, procedure pointer, or type-bound procedure.

19 An **interface block** is a specific interface block, an abstract interface block, or a generic interface block.
20 A specific interface block is a collection of interface bodies. A generic interface block can also be used
21 to specify that a procedure can be invoked

- 22 (1) by using a generic name,
- 23 (2) by using a defined operator,
- 24 (3) by using a defined assignment, or
- 25 (4) for derived-type input/output.

26 **2.2.5 Module**

27 A **module** contains (or accesses from other modules) definitions that are to be made accessible to other
28 program units. These definitions include data object declarations, type definitions, procedure definitions,
29 and interface blocks. A scoping unit in another program unit may access the definitions in a module.
30 Modules are further described in Clause 11.

31 **2.2.6 Submodule**

32 A **submodule** is a program unit that extends a module or another submodule. It may provide definitions
33 (12.6) for procedures whose interfaces are declared (12.4.3.2) in an ancestor module or submodule. It
34 may also contain declarations and definitions of other entities, which are accessible in its descendants.
35 An entity declared in a submodule is not accessible by use association unless it is a module procedure
36 whose interface is declared in the ancestor module. Submodules are further described in Clause 11.

NOTE 2.3

The scoping unit of a submodule accesses the scoping unit of its parent module or submodule by host association.

1 2.3 Execution concepts

2 2.3.1 Statement classification

3 Each Fortran statement is classified as either an executable statement or a nonexecutable statement.
4 There are restrictions on the order in which statements may appear in a program unit, and not all
5 executable statements may appear in all contexts.

6 2.3.2 Program execution

7 An instance of a Fortran program is an **image**. Execution of a program consists of the asynchronous
8 execution of a fixed number (which may be one) of its images. Each image has its own execution state,
9 floating point status (14.6), and set of data objects and procedure pointers. Whether a file is available
10 on any image or only on a specific image is processor dependent. Each image is identified by an **image**
11 **index**, which is an integer value in the range one to the number of images.

NOTE 2.4

The programmer controls the progress of execution in each image through explicit use of Fortran control constructs (8.1, 8.2). Image control statements (8.5.1) affect the relative progress of execution between images. Co-arrays (2.4.6) provide a mechanism for accessing data on one image from another image.

NOTE 2.5

A processor might allow the number of images to be chosen at compile time, link time, or run time. It might be the same as the number of CPUs but this is not required. Compiling for a single image might permit the optimizer to eliminate overhead associated with parallel execution. Portable programs should not make assumptions about the exact number of images. The maximum number of images may be limited due to architectural constraints.

12 A **team** is a set of images formed by invoking the intrinsic collective subroutine FORM_TEAM (13.7.69)
13 for the purposes of collaboration. A team is identified by a scalar variable of type IMAGE_TEAM
14 (13.8.3.7).

15 2.3.3 Executable/nonexecutable statements

16 Image execution is a sequence, in time, of actions. An **executable statement** is an instruction to
17 perform or control one or more of these actions. Thus, the executable statements of a program unit
18 determine the behavior of the program unit. The executable statements are all of those that make up
19 the syntactic class *executable-construct*.

J3 internal note

Unresolved Technical Issue 095

The above definition is incorrect after the addition of the BLOCK construct, since it now includes the specification statements. It needs to be rewritten more carefully.

20 **Nonexecutable statements** do not specify actions; they are used to configure the program environment
21 in which actions take place. The nonexecutable statements are all those not classified as executable.

22 2.3.4 Statement order

23 The syntax rules of clause 2.1 specify the statement order within program units and subprograms. These
24 rules are illustrated in Table 2.1 and Table 2.2. Table 2.1 shows the ordering rules for statements and
25 applies to all program units, subprograms, and interface bodies. Vertical lines delineate varieties of

1 statements that may be interspersed and horizontal lines delineate varieties of statements that shall not
 2 be interspersed. Internal or module subprograms shall follow a CONTAINS statement. Between USE
 3 and CONTAINS statements in a subprogram, nonexecutable statements generally precede executable
 4 statements, although the ENTRY statement, FORMAT statement, and DATA statement may appear
 5 among the executable statements. Table 2.2 shows which statements are allowed in a scoping unit.

Table 2.1: **Requirements on statement ordering**

PROGRAM, FUNCTION, SUBROUTINE, MODULE, SUBMODULE, or BLOCK DATA statement		
USE statements		
IMPORT statements		
IMPLICIT NONE		
FORMAT and ENTRY statements	PARAMETER statements	IMPLICIT statements
	PARAMETER and DATA statements	Derived-type definitions, interface blocks, type declaration statements, enumeration definitions, procedure declarations, specification statements, and statement function statements
	DATA statements	Executable constructs
CONTAINS statement		
Internal subprograms or module subprograms		
END statement		

Table 2.2: **Statements allowed in scoping units**

Statement type	Kind of scoping unit ¹						
	Main program	Module or submodule	Block data	External subprog	Module subprog	Internal subprog	Interface body
USE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IMPORT	No	No	No	No	No	No	Yes
ENTRY	No	No	No	Yes	Yes	No	No
FORMAT	Yes	No	No	Yes	Yes	Yes	No
Misc. decl.s ²	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DATA	Yes	Yes	Yes	Yes	Yes	Yes	No
Derived-type	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interface	Yes	Yes	No	Yes	Yes	Yes	Yes
Executable	Yes	No	No	Yes	Yes	Yes	No
CONTAINS	Yes	Yes	No	Yes	Yes	No	No
Statement function	Yes	No	No	Yes	Yes	Yes	No

(1) The scoping unit of a module or submodule does not include any module subprograms that it contains.

Statements allowed in scoping units

Statement type	Kind of scoping unit ¹						
	Main program	Module or submodule	Block data	External subprog	Module subprog	Internal subprog	Interface body
(2) Miscellaneous	declarations are PARAMETER statements, IMPLICIT statements, type declaration statements, enumeration definitions, procedure declaration statements, and specification statements.						

1 **2.3.5 The END statement**

2 An *end-program-stmt*, *end-function-stmt*, *end-subroutine-stmt*, *end-mp-subprogram-stmt*, *end-module-stmt*, *end-submodule-stmt*, or *end-block-data-stmt* is an **END statement**. Each program unit, module
3 subprogram, and internal subprogram shall have exactly one END statement. The *end-program-stmt*,
4 *end-function-stmt*, *end-subroutine-stmt*, and *end-mp-subprogram-stmt* statements are executable, and
5 may be branch target statements (8.2). Executing an *end-program-stmt* causes normal termination of
6 execution of the program. Executing an *end-function-stmt*, *end-subroutine-stmt*, or *end-mp-subprogram-*
7 *stmt* is equivalent to executing a *return-stmt* with no *scalar-int-expr*.
8

9 The *end-module-stmt*, *end-submodule-stmt*, and *end-block-data-stmt* statements are nonexecutable.

10 **2.3.6 Execution sequence**

11 If a program contains a Fortran main program, execution of the program begins by creating a fixed
12 number of instances of the program; each image begins execution with the first executable construct of
13 the main program. The execution of a main program or subprogram involves execution of the executable
14 constructs within its scoping unit. When a Fortran procedure is invoked, the specification expressions
15 within the *specification-part* of the invoked procedure, if any, are evaluated in a processor dependent
16 order. Thereafter, execution proceeds to the first executable construct appearing after the invoked
17 entry point. With the following exceptions, the effect of execution is as if the executable constructs are
18 executed in the order in which they appear in the main program or subprogram until a STOP, RETURN,
19 or END statement is executed.

- 20 (1) Execution of a branching statement (8.2) changes the execution sequence. These statements
21 explicitly specify a new starting place for the execution sequence.
- 22 (2) CASE constructs, DO constructs, IF constructs, and SELECT TYPE constructs contain
23 an internal statement structure and execution of these constructs involves implicit internal
24 branching. See Clause 8 for the detailed semantics of each of these constructs.
- 25 (3) BLOCK constructs may contain specification expressions; see 8.1.4 for detailed semantics
26 of this construct.
- 27 (4) END=, ERR=, and EOR= specifiers may result in a branch.
- 28 (5) Alternate returns may result in a branch.

29 Internal subprograms may precede the END statement of a main program or a subprogram. The
30 execution sequence excludes all such definitions.

31 The relative ordering of the execution sequences of different images can be affected by image control
32 statements (8.5.1).

33 Normal termination of execution of an image occurs if a STOP statement is executed on that image. Ex-
34 ecution of an *end-program-stmt* results in a synchronization of all images followed by normal termination
35 of execution of all images. Normal termination of execution of an image also may occur during execution
36 of a procedure defined by a companion processor (C International Standard 5.1.2.2.3 and 7.20.4.3). If

1 normal termination of execution occurs within a Fortran program unit and the program incorporates
2 procedures defined by a companion processor, the process of execution termination shall include the
3 effect of executing the C exit() function (C International Standard 7.20.4.3).

4 **2.4 Data concepts**

5 Nonexecutable statements are used to specify the characteristics of the data environment. This includes
6 typing variables, declaring arrays, and defining new types.

7 **2.4.1 Type**

8 A **type** is a named category of data that is characterized by a set of values, a syntax for denoting
9 these values, and a set of operations that interpret and manipulate the values. This central concept is
10 described in 4.1.

11 A type may be parameterized, in which case the set of data values, the syntax for denoting them, and
12 the set of operations depend on the values of one or more parameters. Such a parameter is called a **type**
13 **parameter** (4.2).

14 There are two categories of types: intrinsic types and derived types.

15 **2.4.1.1 Intrinsic type**

16 An **intrinsic type** is a type that is defined by the language, along with operations, and is always
17 accessible. The intrinsic types are integer, real, complex, character, logical, and bits. The properties of
18 intrinsic types are described in 4.4. The intrinsic type parameters are KIND and LEN.

19 The **kind type parameter** indicates the representation method for the specified type.

20 **2.4.1.2 Derived type**

21 A **derived type** is a type that is defined by a type definition or by an intrinsic module. A scalar object of
22 derived type is called a **structure** (4.5). Derived types may be parameterized. Assignment of structures
23 is defined intrinsically (7.4.1.3), but there are no intrinsic operations for structures. For each derived
24 type, a structure constructor is available to provide values (4.5.10). In addition, data objects of derived
25 type may be used as procedure arguments and function results, and may appear in input/output lists.
26 If additional operations are needed for a derived type, they shall be supplied as procedure definitions.

27 Derived types are described further in 4.5.

28 **2.4.2 Data value**

29 Each intrinsic type has associated with it a set of values that a datum of that type may take, depending
30 on the values of the type parameters. The values for each intrinsic type are described in 4.4. The values
31 that objects of a derived type may assume are determined by the type definition, type parameter values,
32 and the sets of values of its components.

33 **2.4.3 Data entity**

34 A **data entity** is a data object, the result of the evaluation of an expression, or the result of the execution
35 of a function reference (called the function result). A data entity has a type and type parameters; it
36 may have a data value (an exception is an undefined variable). Every data entity has a rank and is thus
37 either a scalar or an array.

1 2.4.3.1 Data object

2 A **data object** (often abbreviated to **object**) is a constant (4.1.2), a variable (6), or a subobject of a
 3 constant. The type and type parameters of a named data object may be specified explicitly (5.2) or
 4 implicitly (5.5).

5 **Subobjects** are portions of certain objects that may be referenced and defined (variables only) inde-
 6 pendently of the other portions. These include portions of arrays (array elements and array sections),
 7 portions of character strings (substrings), portions of complex objects (real and imaginary parts), and
 8 portions of structures (components). Subobjects are themselves data objects, but subobjects are refer-
 9 enced only by object designators or intrinsic functions. A subobject of a variable is a variable. Subobjects
 10 are described in Clause 6.

11 Objects referenced by a name are:

12	a named scalar	(a scalar object)
	a named array	(an array object)

13 Subobjects referenced by an object designator are:

14	an array element	(a scalar subobject)
	an array section	(an array subobject)
	a structure component	(a scalar or an array subobject)
	a substring	(a scalar subobject)

15 2.4.3.1.1 Variable

16 A **variable** may have a value and may be defined and redefined during execution of a program.

17 A named **local variable** of the scoping unit of a module, submodule, main program, or subprogram,
 18 is a named variable that is a local entity of the scoping unit, is not a dummy argument, is not in
 19 COMMON, does not have the BIND attribute, and is not accessed by use or host association. A named
 20 local variable of a BLOCK construct is a named variable that is declared in that construct, is not in
 21 COMMON, does not have the BIND attribute, and is not accessed by use association. A subobject of a
 22 named local variable is also a local variable.

23 2.4.3.1.2 Constant

24 A **constant** has a value and cannot become defined, redefined, or undefined during execution of a
 25 program. A constant with a name is called a **named constant** and has the PARAMETER attribute
 26 (5.3.12). A constant without a name is called a **literal constant** (4.4).

27 2.4.3.1.3 Subobject of a constant

28 A **subobject of a constant** is a portion of a constant. The portion referenced may depend on the
 29 value of a variable.

NOTE 2.6

For example, given:

```
CHARACTER (LEN = 10), PARAMETER :: DIGITS = '0123456789'
CHARACTER (LEN = 1)           :: DIGIT
INTEGER :: I
...
DIGIT = DIGITS (I:I)
```

NOTE 2.6 (cont.)

DIGITS is a named constant and DIGITS (I:I) designates a subobject of the constant DIGITS.

1 2.4.3.2 Expression

2 An **expression** (7.1) produces a data entity when evaluated. An expression represents either a data
3 reference or a computation; it is formed from operands, operators, and parentheses. The type, type
4 parameters, value, and rank of an expression result are determined by the rules in Clause 7.

5 2.4.3.3 Function reference

6 A **function reference** (12.5.3) produces a data entity when the function is executed during expression
7 evaluation. The type, type parameters, and rank of a function result are determined by the interface of
8 the function (12.3.3). The value of a function result is determined by execution of the function.

9 2.4.4 Scalar

10 A **scalar** is a datum that is not an array. Scalars may be of any type.

NOTE 2.7

A structure is scalar even if it has arrays as components.

11 The **rank** of a scalar is zero. The shape of a scalar is represented by a rank-one array of size zero.

12 2.4.5 Array

13 An **array** is a set of scalar data, all of the same type and type parameters, whose individual elements
14 are arranged in a rectangular pattern. An **array element** is one of the individual elements in the array
15 and is a scalar. An **array section** is a subset of the elements of an array and is itself an array.

16 An array may have up to fifteen dimensions, and any **extent** (number of elements) in any dimension.
17 The **rank** of the array is the number of dimensions; its **size** is the total number of elements, which is
18 equal to the product of the extents. An array may have zero size. The **shape** of an array is determined
19 by its rank and its extent in each dimension, and may be represented as a rank-one array whose elements
20 are the extents. All named arrays shall be declared, and the rank of a named array is specified in its
21 declaration. The rank of a named array, once declared, is constant; the extents may be constant or may
22 vary during execution.

23 Two arrays are **conformable** if they have the same shape. A scalar is conformable with any array. Any
24 intrinsic operation defined for scalar objects may be applied to conformable objects. Such operations
25 are performed element-by-element to produce a resultant array conformable with the array operands.
26 Element-by-element operation means corresponding elements of the operand arrays are involved in a
27 scalar operation to produce the corresponding element in the result array. Such an operation is described
28 as **elemental**.

NOTE 2.8

If an elemental operation is intrinsically pure or is implemented by a pure elemental function (12.8),
the element operations may be performed simultaneously or in any order.

29 A rank-one array may be constructed from scalars and other arrays and may be reshaped into any
30 allowable array shape (4.7).

31 Arrays may be of any type and are described further in 6.2.

1 2.4.6 Co-array

2 A **co-array** is a data entity that has nonzero co-rank; it can be directly referenced or defined by any
3 image. It may be a scalar or an array.

4 For each co-array on an image, there is a corresponding co-array with the same type, type parameters,
5 and bounds on every other image. If a co-array is scalar, the set of corresponding co-arrays on all the
6 images is arranged in a rectangular pattern. If a co-array is an array, the set of corresponding co-array
7 elements on all the images is arranged in a rectangular pattern. In both cases, the dimensions of the
8 pattern are called **co-dimensions**.

9 A co-array on another image can be accessed directly by using *co-subscripts*. On its own image, a
10 co-array can be accessed without use of co-subscripts.

11 The **co-rank** of a co-array is the number of co-dimensions. The **co-size** of a co-array is always equal to
12 the number of images.

J3 internal note

Unresolved Technical Issue 007

The term "co-dimension" is not defined. It might be better to define co-array in terms of being a rectangular set of objects in co-rank "co-dimensions" etc., similarly to how we define arrays.

13 An object whose designator includes an *image-selector* is a **co-indexed object**. For a co-indexed object,
14 its co-subscript list determines the image index in the same way that a subscript list determines the
15 subscript order value for an array element (6.2.2.2).

16 Intrinsic procedures are provided for mapping between an image index and a list of co-subscripts.

NOTE 2.9

The mechanism for an image to reference and define a co-array on another image might vary according to the hardware. On a shared-memory machine, a co-array could be implemented as a section of an array of higher rank. On a distributed-memory machine with separate physical memory for each image, a processor might store a co-array at the same virtual address in each physical memory.

17 2.4.7 Pointer

18 A **data pointer** is a data entity that has the POINTER attribute. A **procedure pointer** is a procedure
19 entity that has the POINTER attribute. A **pointer** is either a data pointer or a procedure pointer.

20 A pointer is **associated** with a **target** by pointer assignment (7.4.2). A data pointer may also be
21 associated with a target by allocation (6.3.1). A pointer is **disassociated** following execution of a
22 NULLIFY statement, following pointer assignment with a disassociated pointer, by default initialization,
23 or by explicit initialization. A data pointer may also be disassociated by execution of a DEALLOCATE
24 statement. A disassociated pointer is not associated with a target (16.5.2).

25 A pointer that is not associated shall not be referenced or defined.

26 If a data pointer is an array, the rank is declared, but the extents are determined when the pointer is
27 associated with a target.

28 2.4.8 Storage

29 Many of the facilities of this part of ISO/IEC 1539 make no assumptions about the physical storage
30 characteristics of data objects. However, program units that include storage association dependent
31 features shall observe the storage restrictions described in 16.5.3.

1 2.5 Fundamental terms

2 For the purposes of this document, the terms and definitions in this subclause apply.

3 2.5.1 Name and designator

4 A **name** is used to identify a program constituent, such as a program unit, named variable, named
5 constant, dummy argument, or derived type. The rules governing the construction of names are given
6 in 3.2.1. A **designator** is a name followed by zero or more component selectors, complex part selectors,
7 array section selectors, array element selectors, image selectors, and substring selectors.

8 An **object designator** is a designator for a data object. A **procedure designator** is a designator for
9 a procedure.

NOTE 2.10

An object name is a special case of an object designator.

10 2.5.2 Keyword

11 The term **keyword** is used in two ways.

- 12 (1) It is used to describe a word that is part of the syntax of a statement. These keywords are
13 not reserved words; that is, names with the same spellings are allowed. In the syntax rules,
14 such keywords appear literally. In descriptive text, this meaning is denoted by the term
15 “keyword” without any modifier. Examples of statement keywords are: IF, READ, UNIT,
16 KIND, and INTEGER.
- 17 (2) It is used to denote names that identify items in a list. In actual argument lists, type
18 parameter lists, and structure constructors, items may be identified by a preceding *keyword*=
19 rather than their position within the list. An **argument keyword** is the name of a dummy
20 argument in the interface for the procedure being referenced, a **type parameter keyword**
21 is the name of a type parameter in the type being specified, and a **component keyword**
22 is the name of a component in a structure constructor.

23 R215 *keyword* **is** *name*

NOTE 2.11

Use of keywords rather than position to identify items in a list can make such lists more readable and allows them to be reordered. This facilitates specification of a list in cases where optional items are omitted.

24 2.5.3 Association

25 **Association** is name association (16.5.1), pointer association (16.5.2), storage association (16.5.3),
26 or inheritance association (16.5.4). Name association is argument association, host association, use
27 association, linkage association, or construct association.

28 Storage association causes different entities to use the same storage. Any association permits an entity
29 to be identified by different names in the same scoping unit or by the same name or different names in
30 different scoping units.

31 2.5.4 Declaration

32 The term **declaration** refers to the specification of attributes for various program entities. Often this
33 involves specifying the type of a named data object or specifying the shape of a named array object.

1 2.5.5 Definition

2 The term **definition** is used in two ways.

- 3 (1) It refers to the specification of derived types, enumerations, and procedures.
- 4 (2) When an object is given a valid value during program execution, it becomes **defined**. This
5 is often accomplished by execution of an assignment or input statement. When a variable
6 does not have a predictable value, it is **undefined**. Similarly, when a pointer is associated
7 with a target or nullified, its pointer association status is said to become **defined**. When
8 the association status of a pointer is not predictable, its pointer association status is said to
9 be **undefined**.

10 Clause 16 describes the ways in which variables may become defined and undefined.

11 2.5.6 Reference

12 A **data object reference** is the appearance of the data object designator in a context requiring its
13 value at that point during execution.

14 A **procedure reference** is the appearance of the procedure designator, operator symbol, or assignment
15 symbol in a context requiring execution of the procedure at that point. An occurrence of user-defined
16 derived-type input/output (10.7.6) or derived-type finalization (4.5.6.2) is also a procedure reference.

17 The appearance of a data object designator or procedure designator in an actual argument list does not
18 constitute a reference to that data object or procedure unless such a reference is necessary to complete
19 the specification of the actual argument.

20 A **module reference** is the appearance of a module name in a USE statement (11.2.2).

21 2.5.7 Intrinsic

22 The qualifier **intrinsic** has two meanings.

- 23 (1) The qualifier signifies that the term to which it is applied is defined in this part of ISO/IEC
24 1539. Intrinsic applies to types, procedures, modules, assignment statements, and operators.
25 All intrinsic types, procedures, assignments, and operators may be used in any scoping
26 unit without further definition or specification. Intrinsic modules may be accessed by use
27 association. Intrinsic procedures and modules defined in this part of ISO/IEC 1539 are called
28 standard intrinsic procedures and standard intrinsic modules, respectively.
- 29 (2) The qualifier applies to procedures or modules that are provided by a processor but are not
30 defined in this part of ISO/IEC 1539 (13, 14, 15.2). Such procedures and modules are called
31 nonstandard intrinsic procedures and nonstandard intrinsic modules, respectively.

32 2.5.8 Operator

33 An **operator** specifies a computation involving one (unary operator) or two (binary operator) data values
34 (**operands**). This part of ISO/IEC 1539 specifies a number of intrinsic operators (e.g., the arithmetic
35 operators +, -, *, /, and ** with numeric operands and the logical operators .AND., .OR., etc. with
36 logical operands). Additional operators may be defined within a program (4.5.5, 12.4.3.3).

37 2.5.9 Sequence

38 A **sequence** is a set ordered by a one-to-one correspondence with the numbers 1, 2, through n . The
39 number of elements in the sequence is n . A sequence may be empty, in which case it contains no elements.

1 The elements of a nonempty sequence are referred to as the first element, second element, etc. The
2 *n*th element, where *n* is the number of elements in the sequence, is called the last element. An empty
3 sequence has no first or last element.

4 **2.5.10 Companion processors**

5 A processor has one or more companion processors. A **companion processor** is a processor-dependent
6 mechanism by which global data and procedures may be referenced or defined. A companion processor
7 may be a mechanism that references and defines such entities by a means other than Fortran (12.6.3),
8 it may be the Fortran processor itself, or it may be another Fortran processor. If there is more than
9 one companion processor, the means by which the Fortran processor selects among them are processor
10 dependent.

11 If a procedure is defined by means of a companion processor that is not the Fortran processor itself, this
12 part of ISO/IEC 1539 refers to the C function that defines the procedure, although the procedure need
13 not be defined by means of the C programming language.

NOTE 2.12

A companion processor might or might not be a mechanism that conforms to the requirements of the C International Standard.

For example, a processor may allow a procedure defined by some language other than Fortran or C to be invoked if it can be described by a C prototype as defined in 6.5.5.3 of the C International Standard.

1 **3 Lexical tokens, source form, and macro processing**

2 **3.1 Processor character set**

3 The processor character set is processor dependent. Each character in a processor character set is either
 4 a **control character** or a **graphic character**. The set of graphics characters is further divided into
 5 letters (3.1.1), digits (3.1.2), underscore (3.1.3), special characters (3.1.4), and other characters (3.1.5).

6 The letters, digits, underscore, and special characters make up the **Fortran character set**.

7	R301	<i>character</i>	is	<i>alphanumeric-character</i>
8			or	<i>special-character</i>
9	R302	<i>alphanumeric-character</i>	is	<i>letter</i>
10			or	<i>digit</i>
11			or	<i>underscore</i>

12 Except for the currency symbol, the graphics used for the characters shall be as given in 3.1.1, 3.1.2,
 13 3.1.3, and 3.1.4. However, the style of any graphic is not specified.

14 **3.1.1 Letters**

15 The twenty-six **letters** are:

16 A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

17 The set of letters defines the syntactic class *letter*. The processor character set shall include lower-
 18 case and upper-case letters. A lower-case letter is equivalent to the corresponding upper-case letter in
 19 program units except in a character context (3.3).

NOTE 3.1

The following statements are equivalent:

```
CALL BIG_COMPLEX_OPERATION (NDATE)
call big_complex_operation (ndate)
Call Big_Complex_Operation (NDate)
```

20 **3.1.2 Digits**

21 The ten **digits** are:

22 0 1 2 3 4 5 6 7 8 9

23 The ten digits define the syntactic class *digit*.

24 **3.1.3 Underscore**

25 R303 *underscore* **is** `_`

26 The underscore may be used as a significant character in a name.

3.1.4 Special characters

1

2 The **special characters** are shown in Table 3.1.Table 3.1: **Special characters**

Character	Name of character	Character	Name of character
	Blank	;	Semicolon
=	Equals	!	Exclamation point
+	Plus	"	Quotation mark or quote
-	Minus	%	Percent
*	Asterisk	&	Ampersand
/	Slash	~	Tilde
\	Backslash	<	Less than
(Left parenthesis	>	Greater than
)	Right parenthesis	?	Question mark
[Left square bracket	'	Apostrophe
]	Right square bracket	`	Grave accent
{	Left curly bracket	^	Circumflex accent
}	Right curly bracket		Vertical line
,	Comma	\$	Currency symbol
.	Decimal point or period	#	Number sign
:	Colon	@	Commercial at

3 The special characters define the syntactic class *special-character*. Some of the special characters are
 4 used for operator symbols, bracketing, and various forms of separating and delimiting other lexical
 5 tokens.

6 3.1.5 Other characters

7 Additional characters may be representable in the processor, but may appear only in comments (3.3.1.1,
 8 3.3.2.1), character constants (4.4.5), input/output records (9.1.1), and character string edit descriptors
 9 (10.3.2).

10 3.2 Low-level syntax

11 The **low-level syntax** describes the fundamental lexical tokens of a program unit. **Lexical tokens** are
 12 sequences of characters that constitute the building blocks of a program. They are keywords, names,
 13 literal constants other than complex literal constants, operators, labels, delimiters, comma, =, =>, :, ::,
 14 ;, and %.

15 3.2.1 Names

16 **Names** are used for various entities such as variables, program units, dummy arguments, named con-
 17 stants, and derived types.

18 R304 *name* **is** *letter* [*alphanumeric-character*] ...

19 C301 (R304) The maximum length of a *name* is 63 characters.

NOTE 3.2

Examples of names:

NOTE 3.2 (cont.)

A1	
NAME_LENGTH	(single underscore)
S_P_R_E_A_D__O_U_T	(two consecutive underscores)
TRAILER_	(trailing underscore)

NOTE 3.3

The word “name” always denotes this particular syntactic form. The word “identifier” is used where entities may be identified by other syntactic forms or by values; its particular meaning depends on the context in which it is used.

1 3.2.2 Constants

- 2 R305 *constant* **is** *literal-constant*
3 **or** *named-constant*
4 R306 *literal-constant* **is** *int-literal-constant*
5 **or** *real-literal-constant*
6 **or** *complex-literal-constant*
7 **or** *logical-literal-constant*
8 **or** *char-literal-constant*
9 **or** *boz-literal-constant*
10 R307 *named-constant* **is** *name*
11 R308 *int-constant* **is** *constant*
- 12 C302 (R308) *int-constant* shall be of type integer.
- 13 R309 *char-constant* **is** *constant*
- 14 C303 (R309) *char-constant* shall be of type character.

15 3.2.3 Operators

- 16 R310 *intrinsic-operator* **is** *power-op*
17 **or** *mult-op*
18 **or** *add-op*
19 **or** *concat-op*
20 **or** *rel-op*
21 **or** *not-op*
22 **or** *and-op*
23 **or** *or-op*
24 **or** *equiv-op*
25 R707 *power-op* **is** ****
26 R708 *mult-op* **is** ***
27 **or** */*
28 R709 *add-op* **is** *+*
29 **or** *-*
30 R711 *concat-op* **is** *//*
31 R713 *rel-op* **is** *.EQ.*
32 **or** *.NE.*
33 **or** *.LT.*
34 **or** *.LE.*
35 **or** *.GT.*
36 **or** *.GE.*

1		or ==
2		or /=
3		or <
4		or <=
5		or >
6		or >=
7	R718	<i>not-op</i> is .NOT.
8	R719	<i>and-op</i> is .AND.
9	R720	<i>or-op</i> is .OR.
10	R721	<i>equiv-op</i> is .EQV.
11		or .NEQV.
12		or .XOR.
13	R311	<i>defined-operator</i> is <i>defined-unary-op</i>
14		or <i>defined-binary-op</i>
15		or <i>extended-intrinsic-op</i>
16	R703	<i>defined-unary-op</i> is . letter [letter]
17	R723	<i>defined-binary-op</i> is . letter [letter]
18	R312	<i>extended-intrinsic-op</i> is <i>intrinsic-operator</i>

19 3.2.4 Statement labels

20 A **statement label** provides a means of referring to an individual statement.

21 R313 *label* **is** *digit* [*digit* [*digit* [*digit* [*digit*]]]]

22 C304 (R313) At least one digit in a *label* shall be nonzero.

23 If a statement is labeled, the statement shall contain a nonblank character. The same statement label
24 shall not be given to more than one statement in a scoping unit. Leading zeros are not significant in
25 distinguishing between statement labels.

NOTE 3.4

For example:

```
99999
10
 010
```

are all statement labels. The last two are equivalent.

There are 99999 unique statement labels and a processor shall accept any of them as a statement label. However, a processor may have a limit on the total number of unique statement labels in one program unit.

26 Any statement may have a statement label, but the labels are used only in the following ways.

- 27 (1) The label on a branch target statement (8.2) is used to identify that statement as the
28 possible destination of a branch.
- 29 (2) The label on a FORMAT statement (10.2.1) is used to identify that statement as the format
30 specification for a data transfer statement (9.5).
- 31 (3) In some forms of the DO construct (8.1.7), the range of the DO construct is identified by
32 the label on the last statement in that range.

33 3.2.5 Delimiters

1 **Delimiters** are used to enclose syntactic lists. The following pairs are delimiters:

2 (...)

3 / ... /

4 [...]

5 (/ ... /)

6 **3.3 Source form**

7 A Fortran program unit is a sequence of one or more lines, organized as Fortran statements, comments,
8 and INCLUDE lines. A **line** is a sequence of zero or more characters. Lines following a program unit
9 END statement are not part of that program unit. A Fortran **statement** is a sequence of one or more
10 complete or partial lines.

11 A **character context** means characters within a character literal constant (4.4.5) or within a character
12 string edit descriptor (10.3.2).

13 A comment may contain any character that may occur in any character context.

14 There are two **source forms**: free and fixed. Free form and fixed form shall not be mixed in the same program unit.

15 The means for specifying the source form of a program unit are processor dependent.

16 **3.3.1 Free source form**

17 In **free source form** there are no restrictions on where a statement (or portion of a statement) may
18 appear within a line. A line may contain zero characters. If a line consists entirely of characters of
19 default kind (4.4.5), it may contain at most 132 characters. If a line contains any character that is not
20 of default kind, the maximum number of characters allowed on the line is processor dependent.

21 Blank characters shall not appear within lexical tokens other than in a character context or in a format
22 specification. Blanks may be inserted freely between tokens to improve readability; for example, blanks
23 may occur between the tokens that form a complex literal constant. A sequence of blank characters
24 outside of a character context is equivalent to a single blank character.

25 A blank shall be used to separate names, constants, or labels from adjacent keywords, names, constants,
26 or labels.

NOTE 3.5

For example, the blanks after REAL, READ, 30, and DO are required in the following:

```
REAL X
READ 10
30 DO K=1,3
```

27 One or more blanks shall be used to separate adjacent keywords except in the following cases, where
28 blanks are optional:

Adjacent keywords where separating blanks are optional

BLOCK DATA	END MODULE
DOUBLE PRECISION	END INTERFACE
ELSE IF	END PROCEDURE

Adjacent keywords where separating blanks are optional

ELSE WHERE	END PROGRAM
END ASSOCIATE	END SELECT
END BLOCK	END SUBMODULE
END BLOCK DATA	END SUBROUTINE
END CRITICAL	END TYPE
END DO	END WHERE
END ENUM	GO TO
END FILE	IN OUT
END FORALL	SELECT CASE
END FUNCTION	SELECT TYPE
END IF	

1 3.3.1.1 Free form commentary

2 The character “!” initiates a **comment** except where it appears within a character context. The
 3 comment extends to the end of the line. If the first nonblank character on a line is an “!”, the line
 4 is a comment line. Lines containing only blanks or containing no characters are also comment lines.
 5 Comments may appear anywhere in a program unit and may precede the first statement of a program
 6 unit or may follow the last statement of a program unit. Comments have no effect on the interpretation
 7 of the program unit.

NOTE 3.6

The standard does not restrict the number of consecutive comment lines.

8 3.3.1.2 Free form statement continuation

9 The character “&” is used to indicate that the current statement is continued on the next line that is not
 10 a comment line. Comment lines cannot be continued; an “&” in a comment has no effect. Comments may
 11 occur within a continued statement. When used for continuation, the “&” is not part of the statement.
 12 No line shall contain a single “&” as the only nonblank character or as the only nonblank character
 13 before an “!” that initiates a comment.

14 If a noncharacter context is to be continued, an “&” shall be the last nonblank character on the line,
 15 or the last nonblank character before an “!”. There shall be a later line that is not a comment; the
 16 statement is continued on the next such line. If the first nonblank character on that line is an “&”, the
 17 statement continues at the next character position following that “&”; otherwise, it continues with the
 18 first character position of that line.

19 If a lexical token is split across the end of a line, the first nonblank character on the first following
 20 noncomment line shall be an “&” immediately followed by the successive characters of the split token.

21 If a character context is to be continued, an “&” shall be the last nonblank character on the line and
 22 shall not be followed by commentary. There shall be a later line that is not a comment; an “&” shall be
 23 the first nonblank character on the next such line and the statement continues with the next character
 24 following that “&”.

25 3.3.1.3 Free form statement termination

26 If a statement is not continued, a comment or the end of the line terminates the statement.

27 A statement may alternatively be terminated by a “;” character that appears other than in a character
 28 context or in a comment. The “;” is not part of the statement. After a “;” terminator, another statement

1 may appear on the same line, or begin on that line and be continued. A sequence consisting only of zero
2 or more blanks and one or more “;” terminators, in any order, is equivalent to a single “;” terminator.

3 3.3.1.4 Free form statements

4 A label may precede any statement not forming part of another statement.

NOTE 3.7

No Fortran statement begins with a digit.

5 A statement shall not have more than 255 continuation lines.

6 3.3.2 Fixed source form

7 In **fixed source form**, there are restrictions on where a statement may appear within a line. If a source line contains only
8 default kind characters, it shall contain exactly 72 characters; otherwise, its maximum number of characters is processor
9 dependent.

10 Except in a character context, blanks are insignificant and may be used freely throughout the program.

11 3.3.2.1 Fixed form commentary

12 The character “!” initiates a **comment** except where it appears within a character context or in character position 6. The
13 comment extends to the end of the line. If the first nonblank character on a line is an “!” in any character position other
14 than character position 6, the line is a comment line. Lines beginning with a “C” or “*” in character position 1 and lines
15 containing only blanks are also comment lines. Comments may appear anywhere in a program unit and may precede the
16 first statement of the program unit or may follow the last statement of a program unit. Comments have no effect on the
17 interpretation of the program unit.

NOTE 3.8

The standard does not restrict the number of consecutive comment lines.

18 3.3.2.2 Fixed form statement continuation

19 Except within commentary, character position 6 is used to indicate continuation. If character position 6 contains a blank
20 or zero, the line is the initial line of a new statement, which begins in character position 7. If character position 6 contains
21 any character other than blank or zero, character positions 7–72 of the line constitute a continuation of the preceding
22 noncomment line.

NOTE 3.9

An “!” or “;” in character position 6 is interpreted as a continuation indicator unless it appears within commentary
indicated by a “C” or “*” in character position 1 or by an “!” in character positions 1–5.

23 Comment lines cannot be continued. Comment lines may occur within a continued statement.

24 3.3.2.3 Fixed form statement termination

25 If a statement is not continued, a comment or the end of the line terminates the statement.

26 A statement may alternatively be terminated by a “;” character that appears other than in a character context, in a
27 comment, or in character position 6. The “;” is not part of the statement. After a “;” terminator, another statement may
28 begin on the same line, or begin on that line and be continued. A “;” shall not appear as the first nonblank character
29 on an initial line. A sequence consisting only of zero or more blanks and one or more “;” terminators, in any order, is
30 equivalent to a single “;” terminator.

31 3.3.2.4 Fixed form statements

1 A label, if present, shall occur in character positions 1 through 5 of the first line of a statement; otherwise, positions 1
2 through 5 shall be blank. Blanks may appear anywhere within a label. A statement following a “;” on the same line shall
3 not be labeled. Character positions 1 through 5 of any continuation lines shall be blank. A statement shall not have more
4 than 255 continuation lines. The program unit END statement shall not be continued. A statement whose initial line
5 appears to be a program unit END statement shall not be continued.

6 **3.4 Including source text**

7 Additional text may be incorporated into the source text of a program unit during processing. This is
8 accomplished with the **INCLUDE line**, which has the form

9 `INCLUDE char-literal-constant`

10 The *char-literal-constant* shall not have a kind type parameter value that is a *named-constant*.

11 An INCLUDE line is not a Fortran statement.

12 An INCLUDE line shall appear on a single source line where a statement may appear; it shall be the
13 only nonblank text on this line other than an optional trailing comment. Thus, a statement label is not
14 allowed.

15 The effect of the INCLUDE line is as if the referenced source text physically replaced the INCLUDE line
16 prior to program processing. Included text may contain any source text, including additional INCLUDE
17 lines; such nested INCLUDE lines are similarly replaced with the specified source text. The maximum
18 depth of nesting of any nested INCLUDE lines is processor dependent. Inclusion of the source text
19 referenced by an INCLUDE line shall not, at any level of nesting, result in inclusion of the same source
20 text.

21 When an INCLUDE line is resolved, the first included statement line shall not be a continuation line
22 and the last included statement line shall not be continued.

23 The interpretation of *char-literal-constant* is processor dependent. An example of a possible valid inter-
24 pretation is that *char-literal-constant* is the name of a file that contains the source text to be included.

NOTE 3.10

In some circumstances, for example where source code is maintained in an INCLUDE file for use in programs whose source form might be either fixed or free, observing the following rules allows the code to be used with either source form:

- (1) Confine statement labels to character positions 1 to 5 and statements to character positions 7 to 72;
- (2) Treat blanks as being significant;
- (3) Use only the exclamation mark (!) to indicate a comment, but do not start the comment in character position 6;
- (4) For continued statements, place an ampersand (&) in both character position 73 of a continued line and character position 6 of a continuing line.

25 **3.5 Macro processing**

26 **3.5.1 Macro definition**

27 A macro definition defines a macro. A defined macro shall only be referenced by a USE statement,
28 IMPORT statement, or macro expansion statement. A defined macro shall not be redefined.

- 1 R314 *macro-definition* **is** *define-macro-stmt*
 2 [*macro-declaration-stmt*] ...
 3 *macro-body-block*
 4 *end-macro-stmt*
- 5 R315 *define-macro-stmt* **is** DEFINE MACRO [, *macro-attribute-list*] :: *macro-name* ■
 6 ■ [([*macro-dummy-arg-name-list*])]
- 7 C305 (R315) A *macro-dummy-arg-name* shall not appear more than once in a *macro-dummy-arg-*
 8 *name-list*.
- 9 R316 *macro-attribute* **is** *access-spec*
- 10 The DEFINE MACRO statement begins the definition of the macro *macro-name*. Appearance of an
 11 *access-spec* in the DEFINE MACRO statement explicitly gives the macro the specified attribute. Each
 12 *macro-dummy-arg-name* is a macro dummy argument. A macro dummy argument is a macro local
 13 variable.
- 14 R317 *macro-declaration-stmt* **is** *macro-type-declaration-stmt*
 15 **or** *macro-optional-decl-stmt*
- 16 R318 *macro-type-declaration-stmt* **is** MACRO *macro-type-spec* :: *macro-local-variable-name-list*
- 17 R319 *macro-optional-decl-stmt* **is** MACRO OPTIONAL :: *macro-dummy-arg-name-list*
- 18 R320 *macro-type-spec* **is** INTEGER [([KIND=] *macro-expr*)]
- 19 C306 (R318) A *macro-local-variable-name* shall not be the same as the name of a dummy argument
 20 of the macro being defined.
- 21 C307 (R319) A *macro-dummy-arg-name* shall be the name of a dummy argument of the macro being
 22 defined.
- 23 C308 (R320) If *macro-expr* appears, when the macro is expanded *macro-expr* shall be of type integer,
 24 and have a non-negative value that specifies a representation method that exists on the processor.
- 25 A macro type declaration statement specifies that the named entities are macro local variables of the
 26 specified type. If the kind is not specified, they are of default kind. A macro local variable that is not a
 27 macro dummy argument shall appear in a macro type declaration statement.
- 28 R321 *macro-body-block* **is** [*macro-body-construct*] ...
- 29 R322 *macro-body-construct* **is** *macro-definition*
 30 **or** *expand-stmt*
 31 **or** *macro-body-stmt*
 32 **or** *macro-do-construct*
 33 **or** *macro-if-construct*
- 34 C309 A statement in a macro definition that is not a *macro-body-construct* or *macro-definition* shall
 35 not appear on a line with any other statement.
- 36 R323 *macro-do-construct* **is** *macro-do-stmt*
 37 *macro-body-block*
 38 *macro-end-do-stmt*
- 39 R324 *macro-do-stmt* **is** MACRO DO *macro-do-variable-name* = *macro-do-limit* , ■
 40 ■ *macro-do-limit* [, *macro-do-limit*]

- 1 C310 (R324) A *macro-do-variable-name* shall be a local variable of the macro being defined, and shall
2 not be a macro dummy argument.
- 3 R325 *macro-do-limit* is *macro-expr*
- 4 C311 (R325) A *macro-do-limit* shall expand to an expression of type integer.
- 5 R326 *macro-end-do-stmt* is MACRO END DO
- 6 A macro DO construct iterates the expansion of its enclosed macro body block at macro expansion time.
7 The number of iterations is determined by the values of the expanded macro expressions in the MACRO
8 DO statement.
- 9 R327 *macro-if-construct* is *macro-if-then-stmt*
10 *macro-body-block*
11 [*macro-else-if-stmt*
12 *macro-body-block*] ...
13 [*macro-else-stmt*
14 *macro-body-block*]
15 *macro-end-if-stmt*
- 16 R328 *macro-if-then-stmt* is MACRO IF (*macro-condition*) THEN
- 17 R329 *macro-else-if-stmt* is MACRO ELSE IF (*macro-condition*) THEN
- 18 R330 *macro-else-stmt* is MACRO ELSE
- 19 R331 *macro-end-if-stmt* is MACRO END IF
- 20 R332 *macro-condition* is *macro-expr*
- 21 C312 (R332) A macro condition shall expand to an expression of type logical.
- 22 A macro IF construct provides conditional expansion of its enclosed macro body blocks at macro expansion
23 time. Whether the enclosed macro body blocks contribute to the macro expansion is determined by
24 the logical value of the expanded macro expressions in the MACRO IF and MACRO ELSE IF statements.
- 25 R333 *macro-body-stmt* is *result-token* [*result-token*] ... [&&]
- 26 C313 (R333) The first *result-token* shall not be MACRO unless the second *result-token* is not a keyword
27 or name.
- 28 R334 *result-token* is *token* [%% *token*] ...
- 29 C314 (R334) The concatenated textual *tokens* in a *result-token* shall have the form of a lexical token.
- 30 R335 *token* is any lexical token including labels, keywords, and semi-colon.
- 31 C315 && shall not appear in the last *macro-body-stmt* of a macro definition.
- 32 C316 When a macro is expanded, the last *macro-body-stmt* processed shall not end with &&.
- 33 R336 *end-macro-stmt* is END MACRO [*macro-name*]
- 34 C317 (R314) The *macro-name* in the END MACRO statement shall be the same as the *macro-name*
35 in the DEFINE MACRO statement.
- 36 R337 *macro-expr* is *basic-token-sequence*
- C318 (R337) A *macro-expr* shall expand to a scalar initialization expression.

1 Macro expressions are used to control the behavior of the MACRO DO and MACRO IF constructs when
 2 a macro is being expanded. The type, type parameters, and value of a macro expression are determined
 3 when that macro expression is expanded.

4 3.5.2 Macro expansion

5 3.5.2.1 General

6 Macro expansion is the conceptual replacement of the EXPAND statement with the Fortran statements
 7 that it produces. The semantics of an EXPAND statement are those of the Fortran statements that it
 8 produces. It is recommended that a processor be capable of displaying the results of macro expansion. It
 9 is processor-dependent whether comments in a macro definition appear in the expansion. It is processor-
 10 dependent whether continuations and consecutive blanks that are not part of a token are preserved.

11 The process of macro expansion produces Fortran statements consisting of tokens. The combined length
 12 of the tokens for a single statement, plus inter-token spacing, shall not be greater than 33280 characters.
 13 If a statement contains any character that is not of default kind, the maximum number of characters
 14 allowed is processor dependent.

NOTE 3.11

This length is so that the result of macro expansion can be formed into valid free form Fortran source, consisting of an initial line and 255 continuation lines, times 130 which allows for beginning and ending continuation characters (&) on each line.

Also, breaking tokens across continuation lines in macro definitions and in EXPAND statements does not affect macro expansion: it is as if they were joined together before replacement.

- 15 R338 *expand-stmt* **is** EXPAND *macro-name* [(*macro-actual-arg-list*)]
- 16 C319 (R338) *macro-name* shall be the name of a macro that was previously defined or accessed via
 17 use or host association.
- 18 C320 (R338) The macro shall expand to a sequence or zero or more complete Fortran statements.
- 19 C321 (R338) The statements produced by a macro expansion shall conform to the syntax rules and
 20 constraints as if they replaced the EXPAND statement prior to program processing.
- 21 C322 (R338) The statements produced by a macro expansion shall not include a statement which
 22 ends the scoping unit containing the EXPAND statement.
- 23 C323 (R338) If a macro expansion produces a statement which begins a new scoping unit, it shall also
 24 produce a statement which ends that scoping unit.
- 25 C324 (R338) If the EXPAND statement appears as the *action-stmt* of an *if-stmt*, it shall expand to
 26 exactly one *action-stmt* that is not an *if-stmt*, *end-program-stmt*, *end-function-stmt*, or *end-*
 27 *subroutine-stmt*.
- 28 C325 (R338) If the EXPAND statement appears as a *do-term-action-stmt*, it shall expand to exactly one *action-stmt*
 29 that is not a *continue-stmt*, a *goto-stmt*, a *return-stmt*, a *stop-stmt*, an *exit-stmt*, a *cycle-stmt*, an *end-function-*
 30 *stmt*, an *end-subroutine-stmt*, an *end-program-stmt*, or an *arithmetic-if-stmt*.
- 31 C326 (R338) If the EXPAND statement has a label, the expansion of the macro shall produce at least
 32 one statement, and the first statement produced shall not have a label.
- 33 C327 (R338) A *macro-actual-arg* shall appear corresponding to each nonoptional macro dummy ar-
 34 gument.

1 C328 (R338) At most one *macro-actual-arg* shall appear corresponding to each optional macro dummy
2 argument.

3 Expansion of a macro is performed by the EXPAND statement. If the EXPAND statement has a label,
4 the label is interpreted after expansion as belonging to the first statement of the expansion.

5 R339 *macro-actual-arg* is [*macro-dummy-name* =] *macro-actual-arg-value*

6 C329 (R339) *macro-dummy-name* shall be the name of a macro dummy argument of the macro being
7 expanded.

8 C330 (R338) The *macro-dummy-name*= shall not be omitted unless it has been omitted from each
9 preceding *macro-actual-arg* in the *expand-stmt*.

10 R340 *macro-actual-arg-value* is *basic-token-sequence*

11 R341 *basic-token-sequence* is *basic-token*
12 or [*basic-token-sequence*] *nested-token-sequence* ■
13 ■ [*basic-token-sequence*]
14 or *basic-token basic-token-sequence*

15 R342 *basic-token* is any lexical token except comma, parentheses, array ■
16 ■ constructor delimiters, and semi-colon.

17 R343 *nested-token-sequence* is ([*arg-token*] ...)
18 or (/ [*arg-token*] ... /)
19 or *lbracket* [*arg-token*] ... *rbracket*

20 R344 *arg-token* is *basic-token*
21 or ,

22 Macro expansion processes any macro declarations of the macro definition, and then expands its macro
23 body block. Any macro expressions in *macro-type-specs* are evaluated and the kinds of the macro
24 variables thereby declared are determined for that particular expansion.

25 Macro expansion of a macro body block processes each macro body construct of the macro body block
26 in turn, starting with the first macro body construct and ending with the last macro body construct.

27 Expansion of a statement within a macro body construct consists of three steps:

- 28 (1) token replacement,
29 (2) token concatenation, and
30 (3) statement-dependent processing.

31 Token replacement replaces each token of a macro body statement or macro expression that is a macro
32 local variable with the value of that variable. In a macro expression, a reference to the PRESENT
33 intrinsic function with a macro dummy argument name as its actual argument is replaced by the token
34 .TRUE. if the specified macro dummy argument is present, and the token .FALSE. if the specified macro
35 dummy argument is not present. Otherwise, the value of a macro dummy argument that is present is
36 the sequence of tokens from the corresponding actual argument. The value of a macro dummy argument
37 that is not present is a zero-length token sequence. The value of an integer macro variable is its minimal-
38 length decimal representation; if negative this will produce two tokens, a minus sign and an unsigned
39 integer literal constant.

40 Token concatenation is performed with the %% operator, which is only permitted inside a macro defini-
41 tion. After expansion, each sequence of single tokens separated by %% operators is replaced by a single
42 token consisting of the concatenated text of the sequence of tokens. The result of a concatenation shall
43 be a valid Fortran token, and may be a different kind of token from one or more of the original sequence

1 of tokens.

NOTE 3.12

For example, the sequence

```
3 %% .14159 %% E %% + %% 0
```

forms the single real literal constant 3.14159E+0.

2 3.5.2.2 Macro body statements

3 Processing a macro body statement produces a whole or partial Fortran statement. A macro body
 4 statement that is either the first macro body statement processed by this macro expansion or the next
 5 macro body statement processed after a macro body statement that did not end with the continuation
 6 generation operator &&, is an initial macro body statement. The next macro body statement processed
 7 after a macro body statement that ends with && is a continuation macro body statement. An initial
 8 macro body statement that does not end with && produces a whole Fortran statement consisting of its
 9 token sequence. All other macro body statements produce partial Fortran statements, and the sequence
 10 of tokens starting with those produced by the initial macro body statement and appending the tokens
 11 produced by each subsequent continuation macro body statement form a Fortran statement. The &&
 12 operators are not included in the token sequence.

13 3.5.2.3 The macro DO construct

14 The macro DO construct specifies the repeated expansion of a macro body block. Processing the macro
 15 DO statement performs the following steps in sequence.

- 16 (1) The initial parameter m_1 , the terminal parameter m_2 , and the incrementation parameter
 17 m_3 are of type integer with the same kind type parameter as the *macro-do-variable-name*.
 18 Their values are given by the first *macro-expr*, the second *macro-expr*, and the third *macro-*
 19 *expr* of the *macro-do-stmt* respectively, including, if necessary, conversion to the kind type
 20 parameter of the *macro-do-variable-name* according to the rules for numeric conversion
 21 (Table 7.12). If the third *macro-expr* does not appear, m_3 has the value 1. The value of m_3
 22 shall not be zero.
- 23 (2) The macro DO variable becomes defined with the value of the initial parameter m_1 .
- 24 (3) The **iteration count** is established and is the value of the expression $(m_2 - m_1 + m_3)/m_3$,
 25 unless that value is negative, in which case the iteration count is 0.

26 After this, the following steps are performed repeatedly until processing of the macro DO construct is
 27 finished.

- 28 (1) The iteration count is tested. If it is zero, the loop terminates and processing of the macro
 29 DO construct is finished.
- 30 (2) If the iteration count is nonzero, the macro body block of the macro DO construct is
 31 expanded.
- 32 (3) The iteration count is decremented by one. The macro DO variable is incremented by the
 33 value of the incrementation parameter m_3 .

34 3.5.2.4 The MACRO IF construct

35 The MACRO IF construct provides conditional expansion of macro body blocks. At most one of the
 36 macro body blocks of the macro IF construct is expanded. The macro conditions of the construct are
 37 evaluated in order until a true value is found or a MACRO ELSE or MACRO END IF statement is
 38 encountered. If a true value or a MACRO ELSE statement is found, the macro body block immediately
 39 following is expanded and this completes the processing of the construct. If none of the evaluated

1 conditions is true and there is no MACRO ELSE statement, the processing of the construct is completed
 2 without expanding any of the macro body blocks within the construct.

3 3.5.2.5 Macro definitions

4 Processing a macro definition defines a new macro. If a macro definition is produced by a macro expansion,
 5 all of the statements of the produced macro definition have token replacement and concatenation
 6 applied to them before the new macro is defined.

7 3.5.2.6 Examples

NOTE 3.13

This is a macro which loops over an array of any rank and processes each array element.

```

DEFINE MACRO loop_over(array,rank,traceinfo)
  MACRO INTEGER :: i
    BLOCK
    MACRO DO i=1,rank
      INTEGER loop_over_temp_%%i
    MACRO END DO
    MACRO DO i=1,rank
      DO loop_over_temp_%%i=1,size(array,i)
    MACRO END DO
      CALL impure_scalar_procedure(array(loop_over_temp_%%i &&
    MACRO DO i=2,rank
      ,loop_over_temp%i &&
    MACRO END DO
      ),traceinfo)
    MACRO DO i=1,rank
      END DO
    MACRO END DO
    END BLOCK
  END MACRO

```

NOTE 3.14

One can effectively pass macro names as macro arguments, since expansion of arguments occurs before analysis of each macro body statement. For example:

```

DEFINE MACRO :: iterator(count,operation)
  MACRO DO i=1,count
    EXPAND operation(i)
  MACRO END DO
END MACRO

DEFINE MACRO :: process_element(j)
  READ *,a(j)
  result(j) = process(a(j))
  IF (j>1) PRINT *,'difference =',result(j)-result(j-1)
END MACRO

EXPAND iterator(17,process_element)

```

This expands into 17 sets of 3 statements:

NOTE 3.14 (cont.)

```
READ *,a(1)
result(1) = process(a(1))
IF (1>1) PRINT *,'difference =',result(1)-result(1-1)
READ *,a(2)
result(2) = process(a(2))
IF (2>1) PRINT *,'difference =',result(2)-result(2-1)
...
READ *,a(17)
result(17) = process(a(17))
IF (17>1) PRINT *,'difference =',result(17)-result(17-1)
```

NOTE 3.15

Using the ability to evaluate initialization expressions under macro control and test them, one can create interfaces and procedures for all kinds of a type, for example:

```
DEFINE MACRO :: i_square_procs()
  MACRO INTEGER i
  MACRO DO i=1,1000
    MACRO IF (SELECTED_INT_KIND(i)>=0 .AND.
              (i==1 .OR. SELECTED_INT_KIND(i)/=SELECTED_INT_KIND(i-1))) THEN
      FUNCTION i_square_range_%i(a) RESULT(r)
        INTEGER(SELECTED_INT_KIND(i)) a,r
        r = a**2
      END FUNCTION
    MACRO END IF
  MACRO END DO
END MACRO
```


1 4 Types

2 4.1 The concept of type

3 Fortran provides an abstract means whereby data may be categorized without relying on a particular
4 physical representation. This abstract means is the concept of type.

5 A type has a name, a set of valid values, a means to denote such values (constants), and a set of
6 operations to manipulate the values.

7 A type is either an intrinsic type or a derived type.

8 This part of ISO/IEC 1539 defines six intrinsic types: integer, real, complex, character, logical, and bits.

9 A derived type is one that is defined by a derived-type definition (4.5.2) or by an intrinsic module. It
10 shall be used only where it is accessible (4.5.2.2). An intrinsic type is always accessible.

11 4.1.1 Set of values

12 For each type, there is a set of valid values. The set of valid values may be completely determined,
13 as is the case for logical and bits, or may be determined by a processor-dependent method, as is the
14 case for integer, character, and real. For complex, the set of valid values consists of the set of all the
15 combinations of the values of the individual components. For derived types, the set of valid values is as
16 defined in 4.5.8.

17 4.1.2 Constants

18 The syntax for literal constants of each intrinsic type is specified in 4.4.

19 The syntax for denoting a value indicates the type, type parameters, and the particular value.

20 A constant value may be given a name (5.3.12, 5.4.10).

21 A structure constructor (4.5.10) may be used to construct a constant value of derived type from an
22 appropriate sequence of initialization expressions (7.1.7). Such a constant value is scalar even though it
23 may have components that are arrays.

24 4.1.3 Operations

25 For each of the intrinsic types, a set of operations and corresponding operators is defined intrinsically.
26 These are described in Clause 7. The intrinsic set may be augmented with operations and operators
27 defined by functions with the OPERATOR interface (12.4.3.2). Operator definitions are described in
28 Clauses 7 and 12.

29 For derived types, there are no intrinsic operations. Operations on derived types may be defined by the
30 program (4.5.11).

31 4.2 Type parameters

32 A type may be parameterized. In this case, the set of values, the syntax for denoting the values, and
33 the set of operations on the values of the type depend on the values of the parameters.

- 1 The intrinsic types are all parameterized. Derived types may be defined to be parameterized.
- 2 A type parameter is either a kind type parameter or a length type parameter. All type parameters are
3 of type integer.
- 4 A kind type parameter may be used in initialization and specification expressions within the derived-type
5 definition (4.5.2) for the type; it participates in generic resolution (12.5.5.2). Each of the intrinsic types
6 has a kind type parameter named KIND, which is used to distinguish multiple representations of the
7 intrinsic type.

NOTE 4.1

The value of a kind type parameter is always known at compile time. Some parameterizations that involve multiple representation forms need to be distinguished at compile time for practical implementation and performance. Examples include the multiple precisions of the intrinsic real type and the possible multiple character sets of the intrinsic character type.

A type parameter of a derived type may be specified to be a kind type parameter in order to allow generic resolution based on the parameter; that is to allow a single generic to include two specific procedures that have interfaces distinguished only by the value of a kind type parameter of a dummy argument. All generic references are resolvable at compile time.

- 8 A length type parameter may be used in specification expressions within the derived-type definition for
9 the type, but it shall not be used in initialization expressions. The intrinsic character type has a length
10 type parameter named LEN, which is the length of the string.

NOTE 4.2

The adjective “length” is used for type parameters other than kind type parameters because they often specify a length, as for intrinsic character type. However, they may be used for other purposes. The important difference from kind type parameters is that their values need not be known at compile time and might change during execution.

- 11 A type parameter value may be specified with a type specification (4.4, 4.5.9).
- 12 R401 *type-param-value* **is** *scalar-int-expr*
13 **or** *
14 **or** :
- 15 C401 (R401) The *type-param-value* for a kind type parameter shall be an initialization expression.
- 16 C402 (R401) A colon may be used as a *type-param-value* only in the declaration of an entity or
17 component that has the POINTER or ALLOCATABLE attribute.
- 18 A **deferred type parameter** is a length type parameter whose value can change during execution of
19 the program. A colon as a *type-param-value* specifies a deferred type parameter.
- 20 The values of the deferred type parameters of an object are determined by successful execution of an
21 ALLOCATE statement (6.3.1), execution of an intrinsic assignment statement (7.4.1.3), execution of a
22 pointer assignment statement (7.4.2), or by argument association (12.5.2).

NOTE 4.3

Deferred type parameters of functions, including function procedure pointers, have no values. Instead, they indicate that those type parameters of the function result will be determined by execution of the function, if it returns an allocated allocatable result or an associated pointer result.

1 An **assumed type parameter** is a length type parameter for a dummy argument that assumes the
 2 type parameter value from the corresponding actual argument; it is also used for an associate name in a
 3 SELECT TYPE construct that assumes the type parameter value from the corresponding selector, and
 4 for a named constant of type character that assumes its length from the *initialization-expr*. An asterisk
 5 as a *type-param-value* specifies an assumed type parameter.

6 4.3 Relationship of types and values to objects

7 The name of a type serves as a type specifier and may be used to declare objects of that type. A
 8 declaration specifies the type of a named object. A data object may be declared explicitly or implicitly.
 9 Data objects may have attributes in addition to their types. Clause 5 describes the way in which a data
 10 object is declared and how its type and other attributes are specified.

11 Scalar data of any intrinsic or derived type may be shaped in a rectangular pattern to compose an array
 12 of the same type and type parameters. An array object has a type and type parameters just as a scalar
 13 object does.

14 A variable is a data object. The type and type parameters of a variable determine which values that
 15 variable may take. Assignment provides one means of defining or redefining the value of a variable of
 16 any type. Assignment is defined intrinsically for all types where the type, type parameters, and shape
 17 of both the variable and the value to be assigned to it are identical. Assignment between objects of
 18 certain differing intrinsic types, type parameters, and shapes is described in Clause 7. A subroutine and
 19 a generic interface (4.5.2, 12.4.3.2) whose generic specifier is ASSIGNMENT (=) define an assignment
 20 that is not defined intrinsically or redefine an intrinsic derived-type assignment (7.4.1.4).

NOTE 4.4

For example, assignment of a real value to an integer variable is defined intrinsically.
--

21 The type of a variable determines the operations that may be used to manipulate the variable.

22 4.3.1 Type specifiers and type compatibility

23 4.3.1.1 General

24 A type is specified by a **type specifier**. In an executable statement, or in an expression within a nonex-
 25 ecutable statement, a *type-spec* is used. In a nonexecutable statement other than within an expression,
 26 a *declaration-type-spec* is used.

27 R402 *type-spec* **is** *intrinsic-type-spec*
 28 **or** *derived-type-spec*

29 C403 (R402) The *derived-type-spec* shall not specify an abstract type (4.5.7).

30 R403 *declaration-type-spec* **is** *intrinsic-type-spec*
 31 **or** TYPE (*intrinsic-type-spec*)
 32 **or** TYPE (*derived-type-spec*)
 33 **or** CLASS (*derived-type-spec*)
 34 **or** CLASS (*)

35 C404 (R403) In a *declaration-type-spec*, every *type-param-value* that is not a colon or an asterisk shall
 36 be a *specification-expr*.

37 C405 (R403) In a *declaration-type-spec* that uses the CLASS keyword, *derived-type-spec* shall specify
 38 an extensible type (4.5.7).

- 1 C406 (R403) The TYPE(*derived-type-spec*) shall not specify an abstract type (4.5.7).
- 2 C407 An entity declared with the CLASS keyword shall be a dummy argument or have the ALLO-
3 CATABLE or POINTER attribute. It shall not have the VALUE attribute.
- 4 An *intrinsic-type-spec* specifies the named intrinsic type and its type parameter values. A *derived-type-*
5 *spec* specifies the named derived type and its type parameter values.

NOTE 4.5

A *type-spec* is used in an array constructor, a SELECT TYPE construct, or an ALLOCATE statement. Elsewhere, a *declaration-type-spec* is used.

6 **4.3.1.2 TYPE**

7 A TYPE type specifier is used to declare entities of an intrinsic or derived type.

8 Where a data entity is declared explicitly using the TYPE type specifier to be of derived type, the
9 specified derived type shall have been defined previously in the scoping unit or be accessible there by
10 use or host association. If the data entity is a function result, the derived type may be specified in
11 the FUNCTION statement provided the derived type is defined within the body of the function or is
12 accessible there by use or host association. If the derived type is specified in the FUNCTION statement
13 and is defined within the body of the function, it is as if the function result variable was declared with
14 that derived type immediately following the *derived-type-def* of the specified derived type.

15 **4.3.1.3 CLASS**

16 A **polymorphic** entity is a data entity that is able to be of differing types during program execution.
17 The type of a data entity at a particular point during execution of a program is its **dynamic type**. The
18 **declared type** of a data entity is the type that it is declared to have, either explicitly or implicitly.

19 A CLASS type specifier is used to declare polymorphic entities. The declared type of a polymorphic
20 entity is the specified type if the CLASS type specifier contains a type name.

21 An entity declared with the CLASS(*) specifier is an **unlimited polymorphic** entity. An unlimited
22 polymorphic entity is not declared to have a type. It is not considered to have the same declared type
23 as any other entity, including another unlimited polymorphic entity.

24 A nonpolymorphic entity is **type compatible** only with entities of the same declared type. A poly-
25 morphic entity that is not an unlimited polymorphic entity is type compatible with entities of the same
26 declared type or any of its extensions. Even though an unlimited polymorphic entity is not considered
27 to have a declared type, it is type compatible with all entities. An entity is type compatible with a type
28 if it is type compatible with entities of that type.

29 Two entities are **type incompatible** if neither is type compatible with the other.

NOTE 4.6

Given

```

TYPE TROOT
...
TYPE, EXTENDS(TROOT) :: TEXTENDED
...
CLASS(TROOT) A
CLASS(TEXTENDED) B
...

```

NOTE 4.6 (cont.)

A is type compatible with B but B is not type compatible with A.

- 1 A polymorphic allocatable object may be allocated to be of any type with which it is type compatible.
 2 A polymorphic pointer or dummy argument may, during program execution, be associated with objects
 3 with which it is type compatible.
- 4 The dynamic type of an allocated allocatable polymorphic object is the type with which it was allocated.
 5 The dynamic type of an associated polymorphic pointer is the dynamic type of its target. The dynamic
 6 type of a nonallocatable nonpointer polymorphic dummy argument is the dynamic type of its associated
 7 actual argument. The dynamic type of an unallocated allocatable or a disassociated pointer is the same
 8 as its declared type. The dynamic type of an entity identified by an associate name (8.1.3) is the dynamic
 9 type of the selector with which it is associated. The dynamic type of an object that is not polymorphic
 10 is its declared type.

NOTE 4.7

Only components of the declared type of a polymorphic object may be designated by component selection (6.1.2).

11 4.4 Intrinsic types

12 4.4.1 Classification and specification

- 13 Each intrinsic type is classified as a numeric type or a nonnumeric type. The **numeric types** are integer,
 14 real, and complex. The nonnumeric intrinsic types are character, logical, and bits.
- 15 The numeric types are provided for numerical computation. The normal operations of arithmetic,
 16 addition (+), subtraction (-), multiplication (*), division (/), exponentiation (**), identity (unary +),
 17 and negation (unary -), are defined intrinsically for the numeric types.

18 R404	<i>intrinsic-type-spec</i>	is INTEGER [<i>kind-selector</i>] or REAL [<i>kind-selector</i>] or DOUBLE PRECISION or COMPLEX [<i>kind-selector</i>] or CHARACTER [<i>char-selector</i>] or LOGICAL [<i>kind-selector</i>] or BITS [<i>kind-selector</i>]
25 R405	<i>kind-selector</i>	is ([KIND =] <i>scalar-int-initialization-expr</i>)

- 26 C408 (R405) The value of *scalar-int-initialization-expr* shall be nonnegative and shall specify a rep-
 27 resentation method that exists on the processor.

28 4.4.2 Integer type

- 29 The set of values for the **integer type** is a subset of the mathematical integers. The processor shall
 30 provide one or more **representation methods** that define sets of values for data of type integer. Each
 31 such method is characterized by a value for a type parameter called the **kind** type parameter; this kind
 32 type parameter is of type default integer. The kind type parameter of a representation method is returned
 33 by the intrinsic inquiry function KIND (13.7.96). The decimal exponent range of a representation method
 34 is returned by the intrinsic function RANGE (13.7.143). The intrinsic function SELECTED_INT_KIND
 35 (13.7.153) returns a kind value based on a specified decimal range requirement. The integer type includes
 36 a zero value, which is considered to be neither negative nor positive. The value of a signed integer zero
 is the same as the value of an unsigned integer zero.

1

2 The processor shall provide at least one representation method with a decimal exponent range greater
3 than or equal to 18.

4 The type specifier for the integer type uses the keyword INTEGER.

5 If the kind type parameter is not specified, the default kind value is KIND (0) and the type specified is
6 **default integer**. The decimal exponent range of default integer shall be at least 5.

7 Any integer value may be represented as a *signed-int-literal-constant*.

8 R406 *signed-int-literal-constant* **is** [*sign*] *int-literal-constant*

9 R407 *int-literal-constant* **is** *digit-string* [*_ kind-param*]

10 R408 *kind-param* **is** *digit-string*

11 **or** *scalar-int-constant-name*

12 R409 *signed-digit-string* **is** [*sign*] *digit-string*

13 R410 *digit-string* **is** *digit* [*digit*] ...

14 R411 *sign* **is** +

15 **or** -

16 C409 (R408) A *scalar-int-constant-name* shall be a named constant of type integer.

17 C410 (R408) The value of *kind-param* shall be nonnegative.

18 C411 (R407) The value of *kind-param* shall specify a representation method that exists on the pro-
19 cessor.

20 The optional kind type parameter following *digit-string* specifies the kind type parameter of the integer
21 constant; if it is not present, the constant is of type default integer.

22 An integer constant is interpreted as a decimal value.

NOTE 4.8

Examples of signed integer literal constants are:

473

+56

-101

21_2

21_SHORT

1976354279568241_8

where SHORT is a scalar integer named constant.

23 4.4.3 Real type

24 The **real type** has values that approximate the mathematical real numbers. The processor shall provide
25 two or more **approximation methods** that define sets of values for data of type real. Each such method
26 has a **representation method** and is characterized by a value for a type parameter called the **kind**
27 type parameter; this kind type parameter is of type default integer. The kind type parameter of an
28 approximation method is returned by the intrinsic inquiry function KIND (13.7.96).

29 The decimal precision, decimal exponent range, and radix of an approximation method are returned by
30 the intrinsic functions PRECISION (13.7.137), RANGE (13.7.143), and RADIX (13.7.140). The intrinsic
31 function SELECTED_REAL_KIND (13.7.154) returns a kind value based on specified precision, range,

1 and radix requirements.

NOTE 4.9

See C.1.1 for remarks concerning selection of approximation methods.

2 The real type includes a zero value. Processors that distinguish between positive and negative zeros
3 shall treat them as mathematically equivalent

- 4 (1) in all relational operations,
5 (2) as actual arguments to intrinsic procedures other than those for which it is explicitly specified
6 that negative zero is distinguished, and
7 (3) as the *scalar-numeric-expr* in an arithmetic IF.

NOTE 4.10

On a processor that can distinguish between 0.0 and -0.0 ,

(X >= 0.0)

evaluates to true if $X = 0.0$ or if $X = -0.0$,

(X < 0.0)

evaluates to false for $X = -0.0$, and

IF (X) 1,2,3

causes a transfer of control to the branch target statement with the statement label “2” for both $X = 0.0$ and $X = -0.0$.

In order to distinguish between 0.0 and -0.0 , a program should use the SIGN function. SIGN(1.0,X) will return -1.0 if $X < 0.0$ or if the processor distinguishes between 0.0 and -0.0 and X has the value -0.0 .

8 The type specifier for the real type uses the keyword REAL. The keyword DOUBLE PRECISION is an
9 alternate specifier for one kind of real type.

10 If the type keyword REAL is specified and the kind type parameter is not specified, the default kind
11 value is KIND (0.0) and the type specified is **default real**. If the type keyword DOUBLE PRECISION
12 is specified, the kind value is KIND (0.0D0) and the type specified is **double precision real**. The
13 decimal precision of the double precision real approximation method shall be greater than that of the
14 default real method.

15 The decimal precision of double precision real shall be at least 10, and its decimal exponent range shall
16 be at least 37. It is recommended that the decimal precision of default real be at least 6, and that its
17 decimal exponent range be at least 37.

- 18 R412 *signed-real-literal-constant* is [*sign*] *real-literal-constant*
19 R413 *real-literal-constant* is *significand* [*exponent-letter exponent*] [*- kind-param*]
20 or *digit-string exponent-letter exponent* [*- kind-param*]
21 R414 *significand* is *digit-string* . [*digit-string*]
22 or . *digit-string*
23 R415 *exponent-letter* is E
24 or D
25 R416 *exponent* is *signed-digit-string*

- 1 C412 (R413) If both *kind-param* and *exponent-letter* are present, *exponent-letter* shall be E.
- 2 C413 (R413) The value of *kind-param* shall specify an approximation method that exists on the
3 processor.
- 4 A real literal constant without a kind type parameter is a default real constant if it is without an
5 exponent part or has exponent letter E, and is a double precision real constant if it has exponent letter
6 D. A real literal constant written with a kind type parameter is a real constant with the specified kind
7 type parameter.
- 8 The exponent represents the power of ten scaling to be applied to the significand or digit string. The
9 meaning of these constants is as in decimal scientific notation.
- 10 The significand may be written with more digits than a processor will use to approximate the value of
11 the constant.

NOTE 4.11

Examples of signed real literal constants are:

```
-12.78
+1.6E3
2.1
-16.E4_8
0.45D-4
10.93E7_QUAD
.123
3E4
```

where QUAD is a scalar integer named constant.

12 4.4.4 Complex type

13 The **complex type** has values that approximate the mathematical complex numbers. The values of a
14 complex type are ordered pairs of real values. The first real value is called the **real part**, and the second
15 real value is called the **imaginary part**.

16 Each approximation method used to represent data entities of type real shall be available for both the
17 real and imaginary parts of a data entity of type complex. A **kind** type parameter may be specified for
18 a complex entity and selects for both parts the real approximation method characterized by this kind
19 type parameter value; this kind type parameter is of type default integer. The kind type parameter of
20 an approximation method is returned by the intrinsic inquiry function KIND (13.7.96).

21 The type specifier for the complex type uses the keyword COMPLEX. There is no keyword for double
22 precision complex. If the type keyword COMPLEX is specified and the kind type parameter is not
23 specified, the default kind value is the same as that for default real, the type of both parts is default
24 real, and the type specified is **default complex**.

25	R417	<i>complex-literal-constant</i>	is	(<i>real-part</i> , <i>imag-part</i>)
26	R418	<i>real-part</i>	is	<i>signed-int-literal-constant</i>
27			or	<i>signed-real-literal-constant</i>
28			or	<i>named-constant</i>
29	R419	<i>imag-part</i>	is	<i>signed-int-literal-constant</i>
30			or	<i>signed-real-literal-constant</i>
31			or	<i>named-constant</i>

- 1 C414 (R417) Each named constant in a complex literal constant shall be of type integer or real.
- 2 If the real part and the imaginary part of a complex literal constant are both real, the kind type
3 parameter value of the complex literal constant is the kind type parameter value of the part with the
4 greater decimal precision; if the precisions are the same, it is the kind type parameter value of one of the
5 parts as determined by the processor. If a part has a kind type parameter value different from that of
6 the complex literal constant, the part is converted to the approximation method of the complex literal
7 constant.
- 8 If both the real and imaginary parts are integer, they are converted to the default real approximation
9 method and the constant is of type default complex. If only one of the parts is an integer, it is converted
10 to the approximation method selected for the part that is real and the kind type parameter value of the
11 complex literal constant is that of the part that is real.

NOTE 4.12

Examples of complex literal constants are:

```
(1.0, -1.0)
(3, 3.1E6)
(4.0_4, 3.6E7_8)
( 0., PI)
```

where PI is a previously declared named real constant.

12 4.4.5 Character type**13 4.4.5.1 Character sets**

14 The **character type** has a set of values composed of character strings. A **character string** is a sequence
15 of characters, numbered from left to right 1, 2, 3, ... up to the number of characters in the string. The
16 number of characters in the string is called the **length** of the string. The length is a type parameter; its
17 kind is processor-dependent and its value is greater than or equal to zero.

18 The processor shall provide one or more **representation methods** that define sets of values for data
19 of type character. Each such method is characterized by a value for a type parameter called the **kind**
20 type parameter; this kind type parameter is of type default integer. The kind type parameter of a rep-
21 resentation method is returned by the intrinsic inquiry function KIND (13.7.96). The intrinsic function
22 SELECTED_CHAR_KIND (13.7.152) returns a kind value based on the name of a character type. Any
23 character of a particular representation method representable in the processor may occur in a character
24 string of that representation method.

25 The character set defined by ISO/IEC 646:1991 (International Reference Version) is referred to as the
26 **ASCII character set** and its corresponding representation method is the **ASCII character type**.
27 The character set defined by ISO/IEC 10646-1:2000 UCS-4 is referred to as the **ISO 10646 character**
28 **set** and its corresponding representation method is the **ISO 10646 character type**.

29 4.4.5.2 Character type specifier

30 The type specifier for the character type uses the keyword CHARACTER.

31 If the kind type parameter is not specified, the default kind value is KIND ('A') and the type specified
32 is **default character**.

33 The default character kind shall support a character set that includes the Fortran character set. By sup-
34 plying nondefault character kinds, the processor may support additional character sets. The characters

1 available in nondefault character kinds are not specified by this standard, except that one character in
 2 each nondefault character set shall be designated as a blank character to be used as a padding character.

3	R420	<i>char-selector</i>	is <i>length-selector</i>
4			or (LEN = <i>type-param-value</i> , ■
5			■ KIND = <i>scalar-int-initialization-expr</i>)
6			or (<i>type-param-value</i> , ■
7			■ [KIND =] <i>scalar-int-initialization-expr</i>)
8			or (KIND = <i>scalar-int-initialization-expr</i> ■
9			■ [, LEN = <i>type-param-value</i>])
10	R421	<i>length-selector</i>	is ([LEN =] <i>type-param-value</i>)
11			or * <i>char-length</i> [,]
12	R422	<i>char-length</i>	is (<i>type-param-value</i>)
13			or <i>scalar-int-literal-constant</i>

14 C415 (R420) The value of *scalar-int-initialization-expr* shall be nonnegative and shall specify a rep-
 15 resentation method that exists on the processor.

16 C416 (R422) The *scalar-int-literal-constant* shall not include a *kind-param*.

17 C417 (R422) A *type-param-value* in a *char-length* shall be a colon, asterisk, or *specification-expr*.

18 C418 (R420 R421 R422) A *type-param-value* of * shall be used only

- 19 (1) to declare a dummy argument,
- 20 (2) to declare a named constant,
- 21 (3) in the *type-spec* of an ALLOCATE statement wherein each *allocate-object* is a dummy
 22 argument of type CHARACTER with an assumed character length,
- 23 (4) in the *type-spec* or *derived-type-spec* of a type guard statement (8.1.9), or
- 24 (5) in an external function, to declare the character length parameter of the function result.

25 C419 A function name shall not be declared with an asterisk *type-param-value* unless it is of type CHAR-
 26 ACTER and is the name of the result of an external function or the name of a dummy function.

27 C420 A function name declared with an asterisk *type-param-value* shall not be an array, a pointer, recursive, or pure.

28 C421 (R421) The optional comma in a *length-selector* is permitted only in a *declaration-type-spec* in a *type-declaration-*
 29 *stmt*.

30 C422 (R421) The optional comma in a *length-selector* is permitted only if no double-colon separator appears in the
 31 *type-declaration-stmt*.

32 C423 (R420) The length specified for a character statement function or for a statement function dummy argument of
 33 type character shall be an initialization expression.

34 The *char-selector* in a CHARACTER *intrinsic-type-spec* and the * *char-length* in an *entity-decl* or in
 35 a *component-decl* of a type definition specify character length. The * *char-length* in an *entity-decl* or
 36 a *component-decl* specifies an individual length and overrides the length specified in the *char-selector*,
 37 if any. If a * *char-length* is not specified in an *entity-decl* or a *component-decl*, the *length-selector* or
 38 *type-param-value* specified in the *char-selector* is the character length. If the length is not specified in a
 39 *char-selector* or a * *char-length*, the length is 1.

40 If the character length parameter value evaluates to a negative value, the length of character entities
 41 declared is zero. A character length parameter value of : indicates a deferred type parameter (4.2). A
 42 *char-length* type parameter value of * has the following meanings.

- 43 (1) If used to declare a dummy argument of a procedure, the dummy argument assumes the
 length of the associated actual argument.

- 1 (2) If used to declare a named constant, the length is that of the constant value.
 2 (3) If used in the *type-spec* of an ALLOCATE statement, each *allocate-object* assumes its length
 3 from the associated actual argument.
 4 (4) If used in the *type-spec* of a type guard statement, the associating entity assumes its length
 5 from the selector.
 6 (5) If used to specify the character length parameter of a function result, any scoping unit invoking the function
 7 shall declare the function name with a character length parameter value other than * or access such a
 8 definition by host or use association. When the function is invoked, the length of the result variable in the
 9 function is assumed from the value of this type parameter.

10 4.4.5.3 Character literal constant

11 A **character literal constant** is written as a sequence of characters, delimited by either apostrophes
 12 or quotation marks.

13 R423 *char-literal-constant* **is** [*kind-param* -] ' [*rep-char*] ... '
 14 **or** [*kind-param* -] " [*rep-char*] ... "

15 C424 (R423) The value of *kind-param* shall specify a representation method that exists on the pro-
 16 cessor.

17 The optional kind type parameter preceding the leading delimiter specifies the kind type parameter of
 18 the character constant; if it is not present, the constant is of type default character.

19 For the type character with kind *kind-param*, if present, and for type default character otherwise, a
 20 **representable character**, *rep-char*, is defined as follows.

- 21 (1) In free source form, it is any graphic character in the processor-dependent character set.
 22 (2) In fixed source form, it is any character in the processor-dependent character set. A processor may restrict
 23 the occurrence of some or all of the control characters.

NOTE 4.13

FORTRAN 77 allowed any character to occur in a character context. This standard allows a source program to contain characters of more than one kind. Some processors may identify characters of nondefault kinds by control characters (called “escape” or “shift” characters). It is difficult, if not impossible, to process, edit, and print files where some occurrences of control characters have their intended meaning and some occurrences might not. Almost all control characters have uses or effects that effectively preclude their use in character contexts and this is why free source form allows only graphic characters as representable characters. Nevertheless, for compatibility with FORTRAN 77, control characters remain permitted in principle in fixed source form.

24 The delimiting apostrophes or quotation marks are not part of the value of the character literal constant.

25 An apostrophe character within a character constant delimited by apostrophes is represented by two
 26 consecutive apostrophes (without intervening blanks); in this case, the two apostrophes are counted as
 27 one character. Similarly, a quotation mark character within a character constant delimited by quotation
 28 marks is represented by two consecutive quotation marks (without intervening blanks) and the two
 29 quotation marks are counted as one character.

30 A zero-length character literal constant is represented by two consecutive apostrophes (without inter-
 31 vening blanks) or two consecutive quotation marks (without intervening blanks) outside of a character
 32 context.

33 The intrinsic operation **concatenation** (//) is defined between two data entities of type character (7.2.3)
 34 with the same kind type parameter.

NOTE 4.14

Examples of character literal constants are:

```
"DON'T"
'DON'T'
```

both of which have the value DON'T and

```
''
```

which has the zero-length character string as its value.

NOTE 4.15

An example of a nondefault character literal constant, where the processor supports the corresponding character set, is:

```
NIHONGO_彼女なしでは何もできない。
```

where NIHONGO is a named constant whose value is the kind type parameter for Nihongo (Japanese) characters. This means “Without her, nothing is possible”.

1 **4.4.5.4 Collating sequence**

2 The processor defines a collating sequence for the character set of each kind of character. A **collating**
3 **sequence** is a one-to-one mapping of the characters into the nonnegative integers such that each charac-
4 ter corresponds to a different nonnegative integer. The intrinsic functions CHAR (13.7.30) and ICHAR
5 (13.7.84) provide conversions between the characters and the integers according to this mapping.

NOTE 4.16

For example:

```
ICHAR ( 'X' )
```

returns the integer value of the character 'X' according to the collating sequence of the processor.

6 The collating sequence of the default character type shall satisfy the following constraints.

7 (1) ICHAR ('A') < ICHAR ('B') < ... < ICHAR ('Z') for the twenty-six upper-case letters.

8 (2) ICHAR ('0') < ICHAR ('1') < ... < ICHAR ('9') for the ten digits.

9 (3) ICHAR (' ') < ICHAR ('0') < ICHAR ('9') < ICHAR ('A') or

10 ICHAR (' ') < ICHAR ('A') < ICHAR ('Z') < ICHAR ('0').

11 (4) ICHAR ('a') < ICHAR ('b') < ... < ICHAR ('z') for the twenty-six lower-case letters.

12 (5) ICHAR (' ') < ICHAR ('0') < ICHAR ('9') < ICHAR ('a') or

13 ICHAR (' ') < ICHAR ('a') < ICHAR ('z') < ICHAR ('0').

14 Except for blank, there are no constraints on the location of the special characters and underscore in
15 the collating sequence, nor is there any specified collating sequence relationship between the upper-case
16 and lower-case letters.

17 The collating sequence for the ASCII character type is as defined by ISO/IEC 646:1991 (International
18 Reference Version); this collating sequence is called the **ASCII collating sequence** in this standard.

19 The collating sequence for the ISO 10646 character type is as defined by ISO/IEC 10646-1:2000.

NOTE 4.17

The intrinsic functions ACHAR (13.7.2) and IACHAR (13.7.77) provide conversion between characters and corresponding integer values according to the ASCII collating sequence.

1 The intrinsic functions LGT, LGE, LLE, and LLT (13.7.101-13.7.104) provide comparisons between
 2 strings based on the ASCII collating sequence. International portability is guaranteed if the set of
 3 characters used is limited to the letters, digits, underscore, and special characters.

4.4.6 Logical type

5 The **logical type** has two values, which represent true and false.

6 The processor shall provide one or more **representation methods** for data of type logical. Each such
 7 method is characterized by a value for a type parameter called the **kind** type parameter; this kind type
 8 parameter is of type default integer. The kind type parameter of a representation method is returned
 9 by the intrinsic inquiry function KIND (13.7.96).

10 The type specifier for the logical type uses the keyword LOGICAL.

11 If the kind type parameter is not specified, the default kind value is KIND (.FALSE.) and the type
 12 specified is **default logical**.

13 R424 *logical-literal-constant* **is** .TRUE. [- *kind-param*]
 14 **or** .FALSE. [- *kind-param*]

15 C425 (R424) The value of *kind-param* shall specify a representation method that exists on the pro-
 16 cessor.

17 The optional kind type parameter following the trailing delimiter specifies the kind type parameter of
 18 the logical constant; if it is not present, the constant is of type default logical.

19 The intrinsic operations defined for data entities of logical type are: negation (.NOT.), conjunction
 20 (.AND.), inclusive disjunction (.OR.), logical equivalence (.EQV.), and logical nonequivalence (.NEQV.)
 21 as described in 7.2.5. There is also a set of intrinsically defined relational operators that compare the
 22 values of data entities of other types and yield a value of type default logical. These operations are
 23 described in 7.2.4.

4.4.7 Bits type

25 The **bits type** has a set of values composed of ordered sequences of bits. The number of bits in
 26 the sequence is specified by the **kind** type parameter, which shall be greater than or equal to zero.
 27 The processor shall provide **representation methods** with kind type parameter values equal to every
 28 nonnegative integer less than or equal to a processor-determined limit. This limit shall be at least as
 29 large as the storage size, expressed in bits, of every supported kind of type integer, real, complex, and
 30 logical. Additional representation methods may be provided.

31 The type specifier for the bits type uses the keyword BITS.

32 If the kind type parameter is not specified for a bits variable, the default kind value is the size of a
 33 numeric storage unit expressed in bits, and the type specified is **default bits**.

34 R425 *boz-literal-constant* **is** *binary-constant* [- *kind-param*]
 35 **or** *octal-constant* [- *kind-param*]
 36 **or** *hex-constant* [- *kind-param*]

37 R426 *binary-constant* **is** B ' *digit* [*digit*] ... '

1 or B " *digit* [*digit*] ... "

2 C426 (R426) *digit* shall have one of the values 0 or 1.

3 R427 *octal-constant* is O ' *digit* [*digit*] ... '

4 or O " *digit* [*digit*] ... "

5 C427 (R427) *digit* shall have one of the values 0 through 7.

6 R428 *hex-constant* is Z ' *hex-digit* [*hex-digit*] ... '

7 or Z " *hex-digit* [*hex-digit*] ... "

8 R429 *hex-digit* is *digit*

9 or A

10 or B

11 or C

12 or D

13 or E

14 or F

15 The *hex-digits* A through F represent the numbers ten through fifteen, respectively; they may be repre-
 16 sented by their lower-case equivalents.

17 If the optional kind type parameter is not specified for a boz literal constant, the kind value is assumed
 18 from the form of the constant. If the constant is a *binary-constant* the kind value is the number
 19 of *digit* characters. If the constant is an *octal-constant* the kind value is three times the number of
 20 *digit* characters. If the constant is a *hex-constant* the kind value is four times the number of *hex-digit*
 21 characters.

NOTE 4.18

Even if a bits value is too large to fit into a single statement as a literal constant, it can be constructed by concatenation of bits named constants.

22 Each digit of an octal constant represents three bits, and each hex digit of a hex constant represents
 23 four bits, according to their numerical representations as binary integers, with leading zero bits where
 24 needed.

25 If a *kind-param* is specified for a boz literal constant and has a value greater than the number of bits
 26 specified by its digits, the constant is padded on the left (13.3) with enough zero bits to create a constant
 27 of kind *kind-param*. If the *kind-param* specified has a value smaller the number of bits specified by its
 28 digits, only the rightmost *kind-param* bits are used to determine the value of the constant.

NOTE 4.19

Though the processor is required to provide bit kinds only up to four times the size of a numeric storage unit, or up to the maximum intrinsic type size (whichever is larger), it is expected that the actual size limit will be much larger, based on system capacity constraints. Use of BITS objects with KIND values equal to small integer multiples of NUMERIC_STORAGE_SIZE should result in more efficient execution.

29 4.5 Derived types

30 4.5.1 Derived type concepts

31 Additional types may be derived from the intrinsic types and other derived types. A type definition is
 32 required to define the name of the type and the names and attributes of its components and type-bound

- 1 procedures.
- 2 A derived type may be parameterized by multiple type parameters, each of which is defined to be either
3 a kind or length type parameter and may have a default value.
- 4 The **ultimate components** of an object of derived type are the components that are of intrinsic type
5 or have the POINTER or ALLOCATABLE attribute, plus the ultimate components of the components
6 of the object that are of derived type and have neither the ALLOCATABLE nor POINTER attribute.

NOTE 4.20

The ultimate components of objects of the derived type `kids` defined below are `name`, `age`, and `other_kids`.

```

type :: person
  character(len=20) :: name
  integer :: age
end type person

type :: kids
  type(person) :: oldest_child
  type(person), allocatable, dimension(:) :: other_kids
end type kids

```

- 7 The **direct components** of an object of derived type are the components of that object, plus the direct
8 components of the components of the object that are of derived type and have neither the ALLOCAT-
9 ABLE nor POINTER attribute.
- 10 By default, no storage sequence is implied by the order of the component definitions. However, a storage
11 order is implied for a sequence type (4.5.2.3). If the derived type has the BIND attribute, the storage
12 sequence is that required by the companion processor (2.5.10, 15.3.4).
- 13 A scalar entity of derived type is a **structure**. If a derived type has the SEQUENCE property, a scalar
14 entity of the type is a **sequence structure**.

4.5.2 Derived-type definition**4.5.2.1 Syntax**

- 17 R430 *derived-type-def* **is** *derived-type-stmt*
18 [*type-param-def-stmt*] ...
19 [*private-or-sequence*] ...
20 [*component-part*]
21 [*type-bound-procedure-part*]
22 *end-type-stmt*
- 23 R431 *derived-type-stmt* **is** TYPE [[, *type-attr-spec-list*] ::] *type-name* ■
24 ■ [(*type-param-name-list*)]
- 25 R432 *type-attr-spec* **is** ABSTRACT
26 **or** *access-spec*
27 **or** BIND (C)
28 **or** EXTENDS (*parent-type-name*)

- 29 C428 (R431) A derived type *type-name* shall not be DOUBLEPRECISION or the same as the name
30 of any intrinsic type defined in this part of ISO/IEC 1539.

- 31 C429 (R431) The same *type-attr-spec* shall not appear more than once in a given *derived-type-stmt*.

- 1 C430 (R432) A *parent-type-name* shall be the name of a previously defined extensible type (4.5.7).
- 2 C431 (R430) If the type definition contains or inherits (4.5.7.2) a deferred binding (4.5.5), ABSTRACT
3 shall appear.
- 4 C432 (R430) If ABSTRACT appears, the type shall be extensible.
- 5 C433 (R430) If EXTENDS appears, SEQUENCE shall not appear.
- 6 C434 (R430) If EXTENDS appears and the type being defined has a co-array ultimate component,
7 its parent type shall have a co-array ultimate component.
- 8 R433 *private-or-sequence* **is** *private-components-stmt*
9 **or** *sequence-stmt*
- 10 C435 (R430) The same *private-or-sequence* shall not appear more than once in a given *derived-type-*
11 *def*.
- 12 R434 *end-type-stmt* **is** END TYPE [*type-name*]
- 13 C436 (R434) If END TYPE is followed by a *type-name*, the *type-name* shall be the same as that in
14 the corresponding *derived-type-stmt*.
- 15 Derived types with the BIND attribute are subject to additional constraints as specified in 15.3.4.

NOTE 4.21

An example of a derived-type definition is:

```
TYPE PERSON
  INTEGER AGE
  CHARACTER (LEN = 50) NAME
END TYPE PERSON
```

An example of declaring a variable CHAIRMAN of type PERSON is:

```
TYPE (PERSON) :: CHAIRMAN
```

16 **4.5.2.2 Accessibility**

17 Types that are defined in a module or accessible in that module by use association have either the
18 PUBLIC or PRIVATE attribute. Types for which an *access-spec* is not explicitly specified in that
19 module have the default accessibility attribute for that module. The default accessibility attribute for a
20 module is PUBLIC unless it has been changed by a PRIVATE statement (5.4.1). Only types that have
21 the PUBLIC attribute in that module are available to be accessed from that module by use association.

22 The accessibility of a type does not affect, and is not affected by, the accessibility of its components and
23 bindings.

24 If a type definition is private, then the type name, and thus the structure constructor (4.5.10) for the
25 type, are accessible only within the module containing the definition, and within its descendants.

NOTE 4.22

An example of a type with a private name is:

```
TYPE, PRIVATE :: AUXILIARY
  LOGICAL :: DIAGNOSTIC
```

NOTE 4.22 (cont.)

```
CHARACTER (LEN = 20) :: MESSAGE
END TYPE AUXILIARY
```

Such a type would be accessible only within the module in which it is defined, and within its descendants.

1 **4.5.2.3 Sequence type**2 R435 *sequence-stmt* **is** SEQUENCE3 C437 (R439) If SEQUENCE appears, each data component shall be declared to be of an intrinsic type
4 or of a sequence derived type.5 C438 (R430) If SEQUENCE appears, a *type-bound-procedure-part* shall not appear.

6 If the **SEQUENCE statement** appears, the type is a **sequence type**. The order of the component
7 definitions in a sequence type specifies a storage sequence for objects of that type. The type is a
8 **numeric sequence type** if there are no type parameters, no pointer or allocatable components, and
9 each component is of type default integer, default real, double precision real, default complex, default
10 logical, default bits, or a numeric sequence type. The type is a **character sequence type** if there
11 are no type parameters, no pointer or allocatable components, and each component is of type default
12 character or a character sequence type.

NOTE 4.23

An example of a numeric sequence type is:

```
TYPE NUMERIC_SEQ
  SEQUENCE
  INTEGER :: INT_VAL
  REAL    :: REAL_VAL
  LOGICAL :: LOG_VAL
END TYPE NUMERIC_SEQ
```

NOTE 4.24

A structure resolves into a sequence of components. Unless the structure includes a SEQUENCE statement, the use of this terminology in no way implies that these components are stored in this, or any other, order. Nor is there any requirement that contiguous storage be used. The sequence merely refers to the fact that in writing the definitions there will necessarily be an order in which the components appear, and this will define a sequence of components. This order is of limited significance because a component of an object of derived type will always be accessed by a component name except in the following contexts: the sequence of expressions in a derived-type value constructor, intrinsic assignment, the data values in namelist input data, and the inclusion of the structure in an input/output list of a formatted data transfer, where it is expanded to this sequence of components. Provided the processor adheres to the defined order in these cases, it is otherwise free to organize the storage of the components for any nonsequence structure in memory as best suited to the particular architecture.

13 **4.5.2.4 Determination of derived types**

14 Derived-type definitions with the same type name may appear in different scoping units, in which case
15 they may be independent and describe different derived types or they may describe the same type.

1 Two data entities have the same type if they are declared with reference to the same derived-type
 2 definition. The definition may be accessed from a module or from a host scoping unit. Data entities in
 3 different scoping units also have the same type if they are declared with reference to different derived-type
 4 definitions that specify the same type name, all have the SEQUENCE property or all have the BIND
 5 attribute, have no components with PRIVATE accessibility, and have type parameters and components
 6 that agree in order, name, and attributes. Otherwise, they are of different derived types. A data entity
 7 declared using a type with the SEQUENCE property or with the BIND attribute is not of the same type
 8 as an entity of a type declared to be PRIVATE or that has any components that are PRIVATE.

NOTE 4.25

An example of declaring two entities with reference to the same derived-type definition is:

```

TYPE POINT
  REAL X, Y
END TYPE POINT
TYPE (POINT) :: X1
CALL SUB (X1)
...
CONTAINS
  SUBROUTINE SUB (A)
    TYPE (POINT) :: A
    ...
  END SUBROUTINE SUB

```

The definition of derived type POINT is known in subroutine SUB by host association. Because the declarations of X1 and A both reference the same derived-type definition, X1 and A have the same type. X1 and A also would have the same type if the derived-type definition were in a module and both SUB and its containing program unit referenced the module.

NOTE 4.26

An example of data entities in different scoping units having the same type is:

```

PROGRAM PGM
  TYPE EMPLOYEE
    SEQUENCE
    INTEGER ID_NUMBER
    CHARACTER (50) NAME
  END TYPE EMPLOYEE
  TYPE (EMPLOYEE) PROGRAMMER
  CALL SUB (PROGRAMMER)
  ...
END PROGRAM PGM
SUBROUTINE SUB (POSITION)
  TYPE EMPLOYEE
    SEQUENCE
    INTEGER ID_NUMBER
    CHARACTER (50) NAME
  END TYPE EMPLOYEE
  TYPE (EMPLOYEE) POSITION
  ...
END SUBROUTINE SUB

```

The actual argument PROGRAMMER and the dummy argument POSITION have the same type

NOTE 4.26 (cont.)

because they are declared with reference to a derived-type definition with the same name, the SEQUENCE property, and components that agree in order, name, and attributes.

Suppose the component name ID_NUMBER was ID_NUM in the subroutine. Because all the component names are not identical to the component names in derived type EMPLOYEE in the main program, the actual argument PROGRAMMER would not be of the same type as the dummy argument POSITION. Thus, the program would not be standard-conforming.

NOTE 4.27

The requirement that the two types have the same name applies to the *type-names* of the respective *derived-type-stmts*, not to local names introduced via renaming in USE statements.

1 **4.5.3 Derived-type parameters**2 **4.5.3.1 Declaration**

3 R436 *type-param-def-stmt* **is** INTEGER [*kind-selector*] , *type-param-attr-spec* :: ■

4 ■ *type-param-decl-list*

5 R437 *type-param-decl* **is** *type-param-name* [= *scalar-int-initialization-expr*]

6 C439 (R436) A *type-param-name* in a *type-param-def-stmt* in a *derived-type-def* shall be one of the
7 *type-param-names* in the *derived-type-stmt* of that *derived-type-def*.

8 C440 (R436) Each *type-param-name* in the *derived-type-stmt* in a *derived-type-def* shall appear as a
9 *type-param-name* in a *type-param-def-stmt* in that *derived-type-def*.

10 R438 *type-param-attr-spec* **is** KIND

11 **or** LEN

12 The derived type is parameterized if the *derived-type-stmt* has any *type-param-names*.

13 Each type parameter is itself of type integer. If its kind selector is omitted, the kind type parameter is
14 default integer.

15 The *type-param-attr-spec* explicitly specifies whether a type parameter is a kind parameter or a length
16 parameter.

17 If a *type-param-decl* has a *scalar-int-initialization-expr*, the type parameter has a default value which
18 is specified by the expression. If necessary, the value is converted according to the rules of intrinsic
19 assignment (7.4.1.3) to a value of the same kind as the type parameter.

20 A type parameter may be used as a primary in a specification expression (7.1.6) in the *derived-type-*
21 *def*. A kind type parameter may also be used as a primary in an initialization expression (7.1.7) in the
22 *derived-type-def*.

NOTE 4.28

The following example uses derived-type parameters.

```

TYPE humongous_matrix(k, d)
  INTEGER, KIND :: k = kind(0.0)
  INTEGER(selected_int_kind(12)), LEN :: d
  !-- Specify a nondefault kind for d.
  REAL(k) :: element(d,d)

```

NOTE 4.28 (cont.)

```
END TYPE
```

In the following example, `dim` is declared to be a kind parameter, allowing generic overloading of procedures distinguished only by `dim`.

```
TYPE general_point(dim)
  INTEGER, KIND :: dim
  REAL :: coordinates(dim)
END TYPE
```

1 **4.5.3.2 Type parameter order**

2 **Type parameter order** is an ordering of the type parameters of a derived type; it is used for derived-
3 type specifiers.

4 The type parameter order of a nonextended type is the order of the type parameter list in the derived-
5 type definition. The type parameter order of an extended type consists of the type parameter order of
6 its parent type followed by any additional type parameters in the order of the type parameter list in the
7 derived-type definition.

NOTE 4.29

Given

```
TYPE :: t1(k1,k2)
  INTEGER,KIND :: k1,k2
  REAL(k1) a(k2)
END TYPE
TYPE,EXTENDS(t1) :: t2(k3)
  INTEGER,KIND :: k3
  LOGICAL(k3) flag
END TYPE
```

the type parameter order for type T1 is K1 then K2, and the type parameter order for type T2 is K1 then K2 then K3.

8 **4.5.4 Components**9 **4.5.4.1 Syntax**

10	R439	<i>component-part</i>	is	[<i>component-def-stmt</i>] ...
11	R440	<i>component-def-stmt</i>	is	<i>data-component-def-stmt</i>
12			or	<i>proc-component-def-stmt</i>
13	R441	<i>data-component-def-stmt</i>	is	<i>declaration-type-spec</i> [[, <i>component-attr-spec-list</i>] ::] ■
14				■ <i>component-decl-list</i>
15	R442	<i>component-attr-spec</i>	is	<i>access-spec</i>
16			or	ALLOCATABLE
17			or	DIMENSION (<i>component-array-spec</i>)
18			or	DIMENSION [(<i>deferred-shape-spec-list</i>)] ■
19				■ <i>lbracket co-array-spec rbracket</i>
20			or	CONTIGUOUS
21			or	POINTER
22	R443	<i>component-decl</i>	is	<i>component-name</i> [(<i>component-array-spec</i>)] ■

- 1 ■ [*lbracket co-array-spec rbracket*] ■
 2 ■ [* *char-length*] [*component-initialization*]
 3 R444 *component-array-spec* **is** *explicit-shape-spec-list*
 4 **or** *deferred-shape-spec-list*
 5
- 6 C441 (R441) No *component-attr-spec* shall appear more than once in a given *component-def-stmt*.
- 7 C442 (R441) If neither the POINTER nor ALLOCATABLE attribute is specified, the *declaration-type-spec*
 8 in the *component-def-stmt* shall specify an intrinsic type or a previously defined derived
 9 type.
- 10 C443 (R441) If the POINTER or ALLOCATABLE attribute is specified, the *declaration-type-spec* in
 11 the *component-def-stmt* shall be CLASS(*) or shall specify an intrinsic type or any accessible
 12 derived type including the type being defined.
- 13 C444 (R441) If the POINTER or ALLOCATABLE attribute is specified, each *component-array-spec*
 14 shall be a *deferred-shape-spec-list*.
- 15 C445 (R441) If a *co-array-spec* appears, it shall be a *deferred-co-shape-spec-list* and the component
 16 shall have the ALLOCATABLE attribute.
- 17 C446 (R441) If a *co-array-spec* appears, the component shall not be of type IMAGE_TEAM (13.8.3.7),
 18 C_PTR, or C_FUNPTR (15.3.3).
- 19 C447 A data component whose type has a co-array ultimate component shall be a nonpointer nonal-
 20 locatable scalar and shall not be a co-array.
- 21 C448 (R441) If neither the POINTER attribute nor the ALLOCATABLE attribute is specified, each
 22 *component-array-spec* shall be an *explicit-shape-spec-list*.
- 23 C449 (R444) Each bound in the *explicit-shape-spec* shall either be an initialization expression or be a
 24 specification expression that does not contain references to specification functions or any object
 25 designators other than named constants or subobjects thereof.

J3 internal note

Unresolved Technical Issue 094

This constraint is inconsistent (as is the one 3 below): it allows

`INTEGER c1(DIGITS(yvariable))`

`INTEGER c2(len_type_param)`

but not

`INTEGER c3(DIGITS(yvariable)+len_type_param)`

It's not quite trivial to word this correctly...

- 26 C450 (R441) A component shall not have both the ALLOCATABLE and the POINTER attribute.
- 27 C451 (R441) If the CONTIGUOUS attribute is specified, the component shall be an array with the
 28 POINTER attribute.
- 29 C452 (R443) The * *char-length* option is permitted only if the component is of type character.
- 30 C453 (R440) Each *type-param-value* within a *component-def-stmt* shall either be a colon, be an ini-
 31 tialization expression, or be a specification expression that contains neither references to speci-
 32 fication functions nor any object designators other than named constants or subobjects thereof.

NOTE 4.30

Because a type parameter is not an object, a *type-param-value* or a bound in an *explicit-shape-spec* may contain a *type-param-name*.

1 R445 *proc-component-def-stmt* is PROCEDURE ([*proc-interface*]) , ■
 2 ■ *proc-component-attr-spec-list* :: *proc-decl-list*

NOTE 4.31

See 12.4.3.5 for definitions of *proc-interface* and *proc-decl*.

3 R446 *proc-component-attr-spec* is POINTER
 4 or PASS [(*arg-name*)]
 5 or NOPASS
 6 or *access-spec*

7 C454 (R445) The same *proc-component-attr-spec* shall not appear more than once in a given *proc-*
 8 *component-def-stmt*.

9 C455 (R445) POINTER shall appear in each *proc-component-attr-spec-list*.

10 C456 (R445) If the procedure pointer component has an implicit interface or has no arguments,
 11 NOPASS shall be specified.

12 C457 (R445) If PASS (*arg-name*) appears, the interface shall have a dummy argument named *arg-*
 13 *name*.

14 C458 (R445) PASS and NOPASS shall not both appear in the same *proc-component-attr-spec-list*.

15 4.5.4.2 Array components

16 A data component is an array if its *component-decl* contains a *component-array-spec* or its *data-compo-*
 17 *nent-def-stmt* contains the DIMENSION clause with a *component-array-spec*. If the *component-decl*
 18 contains a *component-array-spec*, it specifies the array rank, and if the array is explicit shape (5.3.7.2),
 19 the array bounds; otherwise, the *component-array-spec* in the DIMENSION clause specifies the array
 20 rank, and if the array is explicit shape, the array bounds.

21 A data component is a co-array if its *component-decl* contains a *co-array-spec* or its *data-component-def-*
 22 *stmt* contains a DIMENSION clause with a *co-array-spec*. If the *component-decl* contains a *co-array-spec*
 23 it specifies the co-rank; otherwise, the *co-array-spec* in the DIMENSION clause specifies the co-rank.

NOTE 4.32

An example of a derived type definition with an array component is:

```
TYPE LINE
  REAL, DIMENSION (2, 2) :: COORD      !
                                     ! COORD(:,1) has the value of (/X1, Y1/)
                                     ! COORD(:,2) has the value of (/X2, Y2/)
  REAL                          :: WIDTH ! Line width in centimeters
  INTEGER                        :: PATTERN ! 1 for solid, 2 for dash, 3 for dot
END TYPE LINE
```

An example of declaring a variable LINE_SEGMENT to be of the type LINE is:

```
TYPE (LINE)      :: LINE_SEGMENT
```


NOTE 4.32 (cont.)

The scalar variable `LINE_SEGMENT` has a component that is an array. In this case, the array is a subobject of a scalar. The double colon in the definition for `COORD` is required; the double colon in the definition for `WIDTH` and `PATTERN` is optional.

NOTE 4.33

An example of a derived type definition with an allocatable component is:

```

TYPE STACK
  INTEGER                :: INDEX
  INTEGER, ALLOCATABLE :: CONTENTS (:)
END TYPE STACK

```

For each scalar variable of type `STACK`, the shape of the component `CONTENTS` is determined by execution of an `ALLOCATE` statement or assignment statement, or by argument association.

NOTE 4.34

Default initialization of an explicit-shape array component may be specified by an initialization expression consisting of an array constructor (4.7), or of a single scalar that becomes the value of each array element.

1 **4.5.4.3 Pointer components**

- 2 A component is a pointer (2.4.7) if its *component-attr-spec-list* contains the `POINTER` attribute. A
 3 pointer component may be a data pointer or a procedure pointer.

NOTE 4.35

An example of a derived type definition with a pointer component is:

```

TYPE REFERENCE
  INTEGER                :: VOLUME, YEAR, PAGE
  CHARACTER (LEN = 50)   :: TITLE
  PROCEDURE (printer_interface), POINTER :: PRINT => NULL()
  CHARACTER, DIMENSION (:), POINTER    :: SYNOPSIS
END TYPE REFERENCE

```

Any object of type `REFERENCE` will have the four nonpointer components `VOLUME`, `YEAR`, `PAGE`, and `TITLE`, the procedure pointer `PRINT`, which has an explicit interface the same as `printer_interface`, plus a pointer to an array of characters holding `SYNOPSIS`. The size of this target array will be determined by the length of the abstract. The space for the target may be allocated (6.3.1) or the pointer component may be associated with a target by a pointer assignment statement (7.4.2).

4 **4.5.4.4 The passed-object dummy argument**

- 5 A **passed-object dummy argument** is a distinguished dummy argument of a procedure pointer
 6 component or type-bound procedure. It affects procedure overriding (4.5.7.3) and argument association
 7 (12.5.2.2).
 8 If `NOPASS` is specified, the procedure pointer component or type-bound procedure has no passed-object
 9 dummy argument.

- 1 If neither PASS nor NOPASS is specified or PASS is specified without *arg-name*, the first dummy argu-
 2 ment of a procedure pointer component or type-bound procedure is its passed-object dummy argument.
- 3 If PASS (*arg-name*) is specified, the dummy argument named *arg-name* is the passed-object dummy
 4 argument of the procedure pointer component or named type-bound procedure.
- 5 C459 The passed-object dummy argument shall be a scalar, nonpointer, nonallocatable dummy data
 6 object with the same declared type as the type being defined; all of its length type parameters
 7 shall be assumed; it shall be polymorphic (4.3.1.3) if and only if the type being defined is
 8 extensible (4.5.7). It shall not have the VALUE attribute.

NOTE 4.36

If a procedure is bound to several types as a type-bound procedure, different dummy arguments might be the passed-object dummy argument in different contexts.

9 **4.5.4.5 Default initialization for components**

10 Default initialization provides a means of automatically initializing pointer components to be disassoci-
 11 ated or associated with specific targets, and nonpointer nonallocatable components to have a particular
 12 value. Allocatable components are always initialized to not allocated.

13 A pointer variable or component is **data-pointer-initialization compatible** with a target if the pointer
 14 is type compatible with the target, they have the same rank, and the values of corresponding nondeferred
 15 type parameters are specified by initialization expressions that have the same value.

16 R447 *component-initialization* **is** = *initialization-expr*
 17 **or** => *null-init*
 18 **or** => *initial-data-target*

19 R448 *initial-data-target* **is** *designator*

20 C460 (R441) If *component-initialization* appears, a double-colon separator shall appear before the
 21 *component-decl-list*.

22 C461 (R441) If *component-initialization* appears, every type parameter and array bound of the com-
 23 ponent shall be a colon or initialization expression.

24 C462 (R441) If => appears in *component-initialization*, POINTER shall appear in the *component-*
 25 *attr-spec-list*. If = appears in *component-initialization*, neither POINTER nor ALLOCATABLE
 26 shall appear in the *component-attr-spec-list*.

27 C463 (R447) If *initial-data-target* appears, *component-name* shall be data-pointer-initialization com-
 28 patible with it.

29 C464 (R448) The *designator* shall designate a variable that is an initialization target. Every subscript,
 30 section subscript, substring starting point, and substring ending point in *designator* shall be an
 31 initialization expression.

32 If *null-init* appears for a pointer component, that component in any object of the type has an initial
 33 association status of disassociated (16.5.2.2) or becomes disassociated as specified in 16.5.2.2.2.

34 A variable is an **initialization target** if it has the TARGET attribute, either has the SAVE attribute
 35 or is declared in the main program, and does not have the ALLOCATABLE attribute.

36 If *initial-data-target* appears for a data pointer component, that component in any object of the type is
 37 initially associated with the target or becomes associated with the target as specified in 16.5.2.2.1.

- 1 If *initial-proc-target* (12.4.3.5) appears in *proc-decl* for a procedure pointer component, that component
 2 in any object of the type is initially associated with the target or becomes associated with the target as
 3 specified in 16.5.2.2.1.
- 4 If *initialization-expr* appears for a nonpointer component, that component in any object of the type
 5 is initially defined (16.6.3) or becomes defined as specified in 16.6.5 with the value determined from
 6 *initialization-expr*. If necessary, the value is converted according to the rules of intrinsic assignment
 7 (7.4.1.3) to a value that agrees in type, type parameters, and shape with the component. If the compo-
 8 nent is of a type for which default initialization is specified for a component, the default initialization
 9 specified by *initialization-expr* overrides the default initialization specified for that component. When
 10 one initialization **overrides** another it is as if only the overriding initialization were specified (see Note
 11 4.38). Explicit initialization in a type declaration statement (5.2) overrides default initialization (see
 12 Note 4.37). Unlike explicit initialization, default initialization does not imply that the object has the
 13 SAVE attribute.
- 14 A subcomponent (6.1.2) is **default-initialized** if the type of the object of which it is a component
 15 specifies default initialization for that component, and the subcomponent is not a subobject of an object
 16 that is default-initialized or explicitly initialized.

NOTE 4.37

It is not required that initialization be specified for each component of a derived type. For example:

```
TYPE DATE
  INTEGER DAY
  CHARACTER (LEN = 5) MONTH
  INTEGER :: YEAR = 1994      ! Partial default initialization
END TYPE DATE
```

In the following example, the default initial value for the YEAR component of TODAY is overridden by explicit initialization in the type declaration statement:

```
TYPE (DATE), PARAMETER :: TODAY = DATE (21, "Feb.", 1995)
```

NOTE 4.38

The default initial value of a component of derived type may be overridden by default initialization specified in the definition of the type. Continuing the example of Note 4.37:

```
TYPE SINGLE_SCORE
  TYPE (DATE) :: PLAY_DAY = TODAY
  INTEGER SCORE
  TYPE (SINGLE_SCORE), POINTER :: NEXT => NULL ( )
END TYPE SINGLE_SCORE
TYPE (SINGLE_SCORE) SETUP
```

The PLAY_DAY component of SETUP receives its initial value from TODAY, overriding the initialization for the YEAR component.

NOTE 4.39

Arrays of structures may be declared with elements that are partially or totally initialized by default. Continuing the example of Note 4.38 :

```
TYPE MEMBER (NAME_LEN)
```

NOTE 4.39 (cont.)

```

INTEGER, LEN :: NAME_LEN
CHARACTER (LEN = NAME_LEN) NAME = ''
INTEGER :: TEAM_NO, HANDICAP = 0
TYPE (SINGLE_SCORE), POINTER :: HISTORY => NULL ( )
END TYPE MEMBER
TYPE (MEMBER(9)) LEAGUE (36)          ! Array of partially initialized elements
TYPE (MEMBER(9)) :: ORGANIZER = MEMBER ("I. Manage",1,5,NULL ( ))

```

ORGANIZER is explicitly initialized, overriding the default initialization for an object of type MEMBER.

Allocated objects may also be initialized partially or totally. For example:

```

ALLOCATE (ORGANIZER % HISTORY)      ! A partially initialized object of type
                                   ! SINGLE_SCORE is created.

```

NOTE 4.40

A pointer component of a derived type may have as its target an object of that derived type. The type definition may specify that in objects declared to be of this type, such a pointer is default initialized to disassociated. For example:

```

TYPE NODE
  INTEGER          :: VALUE = 0
  TYPE (NODE), POINTER :: NEXT_NODE => NULL ( )
END TYPE

```

A type such as this may be used to construct linked lists of objects of type NODE. See C.1.4 for an example. Linked lists can also be constructed using allocatable components.

NOTE 4.41

A pointer component of a derived type may be default-initialized to have an initial target.

```

TYPE NODE
  INTEGER          :: VALUE = 0
  TYPE (NODE), POINTER :: NEXT_NODE => SENTINEL
END TYPE

TYPE(NODE), SAVE, TARGET :: SENTINEL

```

1 4.5.4.6 Component order

2 **Component order** is an ordering of the nonparent components of a derived type; it is used for intrinsic
3 formatted input/output and structure constructors (where component keywords are not used). Parent
4 components are excluded from the component order of an extended type (4.5.7).

5 The component order of a nonextended type is the order of the declarations of the components in the
6 derived-type definition. The component order of an extended type consists of the component order of
7 its parent type followed by any additional components in the order of their declarations in the extended
8 derived-type definition.

NOTE 4.42

Given the same type definitions as in Note 4.29, the component order of type T1 is just A (there is only one component), and the component order of type T2 is A then FLAG. The parent component (T1) does not participate in the component order.

1 **4.5.4.7 Component accessibility**2 R449 *private-components-stmt* is PRIVATE3 C465 (R449) A *private-components-stmt* is permitted only if the type definition is within the speci-
4 cation part of a module.5 The default accessibility for the components that are declared in a type's *component-part* is private
6 if the type definition contains a *private-components-stmt*, and public otherwise. The accessibility of a
7 component may be explicitly declared by an *access-spec*; otherwise its accessibility is the default for the
8 type definition in which it is declared.9 If a component is private, that component name is accessible only within the module containing the
10 definition, and within its descendants.**NOTE 4.43**

Type parameters are not components. They are effectively always public.

NOTE 4.44

The accessibility of the components of a type is independent of the accessibility of the type name. It is possible to have all four combinations: a public type name with a public component, a private type name with a private component, a public type name with a private component, and a private type name with a public component.

NOTE 4.45

An example of a type with private components is:

```
TYPE POINT
  PRIVATE
  REAL :: X, Y
END TYPE POINT
```

Such a type definition is accessible in any scoping unit accessing the module via a USE statement; however, the components X and Y are accessible only within the module, and within its descendants.

NOTE 4.46

The following example illustrates the use of an individual component *access-spec* to override the default accessibility:

```
TYPE MIXED
  PRIVATE
  INTEGER :: I
  INTEGER, PUBLIC :: J
END TYPE MIXED

TYPE (MIXED) :: M
```

NOTE 4.46 (cont.)

The component M%J is accessible in any scoping unit where M is accessible; M%I is accessible only within the module containing the TYPE MIXED definition, and within its descendants.

1 **4.5.5 Type-bound procedures**

- 2 R450 *type-bound-procedure-part* **is** *contains-stmt*
 3 [*binding-private-stmt*]
 4 [*proc-binding-stmt*] ...
- 5 R451 *binding-private-stmt* **is** PRIVATE
- 6 C466 (R450) A *binding-private-stmt* is permitted only if the type definition is within the specification
 7 part of a module.
- 8 R452 *proc-binding-stmt* **is** *specific-binding*
 9 **or** *generic-binding*
 10 **or** *final-binding*
- 11 R453 *specific-binding* **is** PROCEDURE [(*interface-name*)] ■
 12 ■ [[, *binding-attr-list*] ::] ■
 13 ■ *binding-name* [=> *procedure-name*]
- 14 C467 (R453) If => *procedure-name* appears, the double-colon separator shall appear.
- 15 C468 (R453) If => *procedure-name* appears, *interface-name* shall not appear.
- 16 C469 (R453) The *procedure-name* shall be the name of an accessible module procedure or an external
 17 procedure that has an explicit interface.
- 18 If neither => *procedure-name* nor *interface-name* appears, it is as though => *procedure-name* had
 19 appeared with a procedure name the same as the binding name.
- 20 R454 *generic-binding* **is** GENERIC ■
 21 ■ [, *access-spec*] :: *generic-spec* => *binding-name-list*
- 22 C470 (R454) Within the *specification-part* of a module, each *generic-binding* shall specify, either
 23 implicitly or explicitly, the same accessibility as every other *generic-binding* with that *generic-*
 24 *spec* in the same derived type.
- 25 C471 (R454) Each *binding-name* in *binding-name-list* shall be the name of a specific binding of the
 26 type.
- 27 C472 (R454) If *generic-spec* is not *generic-name*, each of its specific bindings shall have a passed-object
 28 dummy argument (4.5.4.4).
- 29 C473 (R454) If *generic-spec* is OPERATOR (*defined-operator*), the interface of each binding shall
 30 be as specified in 12.4.3.3.1.
- 31 C474 (R454) If *generic-spec* is ASSIGNMENT (=), the interface of each binding shall be as specified
 32 in 12.4.3.3.2.
- 33 C475 (R454) If *generic-spec* is *dtio-generic-spec*, the interface of each binding shall be as specified in
 34 9.5.4.7. The type of the *dtv* argument shall be *type-name*.
- 35 R455 *binding-attr* **is** PASS [(*arg-name*)]
 36 **or** NOPASS
 37 **or** NON_OVERRIDABLE

- 1 **or** DEFERRED
 2 **or** *access-spec*
- 3 C476 (R455) The same *binding-attr* shall not appear more than once in a given *binding-attr-list*.
- 4 C477 (R453) If the interface of the binding has no dummy argument of the type being defined,
 5 NOPASS shall appear.
- 6 C478 (R453) If PASS (*arg-name*) appears, the interface of the binding shall have a dummy argument
 7 named *arg-name*.
- 8 C479 (R455) PASS and NOPASS shall not both appear in the same *binding-attr-list*.
- 9 C480 (R455) NON_OVERRIDABLE and DEFERRED shall not both appear in the same *binding-*
 10 *attr-list*.
- 11 C481 (R455) DEFERRED shall appear if and only if *interface-name* appears.
- 12 C482 (R453) An overriding binding (4.5.7.3) shall have the DEFERRED attribute only if the binding
 13 it overrides is deferred.
- 14 C483 (R453) A binding shall not override an inherited binding (4.5.7.2) that has the NON_OVER-
 15 RIDABLE attribute.
- 16 Each binding in a *proc-binding-stmt* specifies a **type-bound procedure**. A type-bound procedure
 17 may have a passed-object dummy argument (4.5.4.4). A *generic-binding* specifies a type-bound generic
 18 interface for its specific bindings. A binding that specifies the DEFERRED attribute is a **deferred**
 19 **binding**. A deferred binding shall appear only in the definition of an abstract type.
- 20 A type-bound procedure may be identified by a **binding name** in the scope of the type definition.
 21 This name is the *binding-name* for a specific binding, and the *generic-name* for a generic binding whose
 22 *generic-spec* is *generic-name*. A final binding, or a generic binding whose *generic-spec* is not *generic-*
 23 *name*, has no binding name.
- 24 The interface of a specific binding is that of the procedure specified by *procedure-name* or the interface
 25 specified by *interface-name*.

NOTE 4.47

An example of a type and a type-bound procedure is:

```

TYPE POINT
  REAL :: X, Y
CONTAINS
  PROCEDURE, PASS :: LENGTH => POINT_LENGTH
END TYPE POINT
...
```

and in the *module-subprogram-part* of the same module:

```

REAL FUNCTION POINT_LENGTH (A, B)
  CLASS (POINT), INTENT (IN) :: A, B
  POINT_LENGTH = SQRT ( (A%X - B%X)**2 + (A%Y - B%Y)**2 )
END FUNCTION POINT_LENGTH
```

- 26 The same *generic-spec* may be used in several *generic-bindings* within a single derived-type definition.
 27 Each additional *generic-binding* with the same *generic-spec* extends the generic interface.

NOTE 4.48

Unlike the situation with generic procedure names, a generic type-bound procedure name is not permitted to be the same as a specific type-bound procedure name in the same type (16.3).

1 The default accessibility for the procedure bindings of a type is private if the type definition contains a
 2 *binding-private-stmt*, and public otherwise. The accessibility of a procedure binding may be explicitly
 3 declared by an *access-spec*; otherwise its accessibility is the default for the type definition in which it is
 4 declared.

5 A public type-bound procedure is accessible via any accessible object of the type. A private type-bound
 6 procedure is accessible only within the module containing the type definition, and within its descendants.

NOTE 4.49

The accessibility of a type-bound procedure is not affected by a PRIVATE statement in the *component-part*; the accessibility of a data component is not affected by a PRIVATE statement in the *type-bound-procedure-part*.

7 **4.5.6 Final subroutines**8 **4.5.6.1 Declaration**

9 R456 *final-binding* is FINAL [::] *final-subroutine-name-list*

10 C484 (R456) A *final-subroutine-name* shall be the name of a module procedure with exactly one
 11 dummy argument. That argument shall be nonoptional and shall be a nonpointer, nonallocat-
 12 able, nonpolymorphic variable of the derived type being defined. All length type parameters of
 13 the dummy argument shall be assumed. The dummy argument shall not be INTENT(OUT).

14 C485 (R456) A *final-subroutine-name* shall not be one previously specified as a final subroutine for
 15 that type.

16 C486 (R456) A final subroutine shall not have a dummy argument with the same kind type parameters
 17 and rank as the dummy argument of another final subroutine of that type.

18 The FINAL keyword specifies a list of **final subroutines**. A final subroutine might be executed when
 19 a data entity of that type is finalized (4.5.6.2).

20 A derived type is **finalizable** if it has any final subroutines or if it has any nonpointer, nonallocatable
 21 component whose type is finalizable. A nonpointer data entity is finalizable if its type is finalizable.

NOTE 4.50

Final subroutines are effectively always “accessible”. They are called for entity finalization regardless of the accessibility of the type, its other type-bound procedures, or the subroutine name itself.

NOTE 4.51

Final subroutines are not inherited through type extension and cannot be overridden. The final subroutines of the parent type are called after any additional final subroutines of an extended type are called.

22 **4.5.6.2 The finalization process**

23 Only finalizable entities are finalized. When an entity is **finalized**, the following steps are carried out in sequence.

- 1 (1) If the dynamic type of the entity has a final subroutine whose dummy argument has the
2 same kind type parameters and rank as the entity being finalized, it is called with the entity
3 as an actual argument. Otherwise, if there is an elemental final subroutine whose dummy
4 argument has the same kind type parameters as the entity being finalized, it is called with
5 the entity as an actual argument. Otherwise, no subroutine is called at this point.
- 6 (2) Each finalizable component that appears in the type definition is finalized. If the entity
7 being finalized is an array, each finalizable component of each element of that entity is
8 finalized separately.
- 9 (3) If the entity is of extended type and the parent type is finalizable, the parent component is
10 finalized.

11 If several entities are to be finalized as a consequence of an event specified in 4.5.6.3, the order in which
12 they are finalized is processor-dependent. A final subroutine shall not reference or define an object that
13 has already been finalized.

14 If an object is not finalized, it retains its definition status and does not become undefined.

15 4.5.6.3 When finalization occurs

16 When a pointer is deallocated its target is finalized. When an allocatable entity is deallocated, it is
17 finalized.

18 A nonpointer, nonallocatable object that is not a dummy argument or function result is finalized immedi-
19 ately before it would become undefined due to execution of a RETURN or END statement (16.6.6, item
20 (3)). If the object is defined in a module or submodule, and there are no longer any active procedures
21 referencing the module or submodule, it is processor-dependent whether it is finalized.

22 A nonpointer nonallocatable local variable of a BLOCK construct is finalized immediately before it
23 would become undefined due to termination of the BLOCK construct (16.6.6, item (20)).

24 If an executable construct references a function, the result is finalized after execution of the innermost
25 executable construct containing the reference.

26 If an executable construct references a structure constructor or array constructor, the entity created by
27 the constructor is finalized after execution of the innermost executable construct containing the reference.

28 If a specification expression in a scoping unit references a function, the result is finalized before execution
29 of the executable constructs in the scoping unit.

30 If a specification expression in a scoping unit references a structure constructor or array constructor, the
31 entity created by the constructor is finalized before execution of the executable constructs in the scoping
32 unit.

33 When a procedure is invoked, a nonpointer, nonallocatable object that is an actual argument associated
34 with an INTENT(OUT) dummy argument is finalized.

35 When an intrinsic assignment statement is executed, the variable is finalized after evaluation of *expr*
36 and before the definition of the variable.

NOTE 4.52

If finalization is used for storage management, it often needs to be combined with defined assign- ment.

37 If an object is allocated via pointer allocation and later becomes unreachable due to all pointers to that
38 object having their pointer association status changed, it is processor dependent whether it is finalized.
39 If it is finalized, it is processor dependent as to when the final subroutines are called.

1 4.5.6.4 Entities that are not finalized

- 2 If image execution is terminated, either by an error (e.g. an allocation failure) or by execution of a STOP
3 or END PROGRAM statement, entities existing immediately prior to termination are not finalized.

NOTE 4.53

A nonpointer, nonallocatable object that has the SAVE attribute or that occurs in the main program is never finalized as a direct consequence of the execution of a RETURN or END statement.

A variable in a module or submodule is not finalized if it retains its definition status and value, even when there is no active procedure referencing the module or submodule.

4 4.5.7 Type extension

5 4.5.7.1 Concepts

- 6 A nonsequence derived type that does not have the BIND attribute is an **extensible type**.
7 An extensible type that does not have the EXTENDS attribute is a **base type**. A type that has the
8 EXTENDS attribute is an **extended type**. The **parent type** of an extended type is the type named
9 in the EXTENDS attribute specification.

NOTE 4.54

The name of the parent type might be a local name introduced via renaming in a USE statement.

- 10 A base type is an **extension type** of itself only. An extended type is an extension of itself and of all
11 types for which its parent type is an extension.
12 An **abstract type** is a type that has the ABSTRACT attribute.

NOTE 4.55

A deferred binding (4.5.5) defers the implementation of a type-bound procedure to extensions of the type; it may appear only in an abstract type. The dynamic type of an object cannot be abstract; therefore, a deferred binding cannot be invoked. An extension of an abstract type need not be abstract if it has no deferred bindings. A short example of an abstract type is:

```
TYPE, ABSTRACT :: FILE_HANDLE
CONTAINS
  PROCEDURE(OPEN_FILE), DEFERRED, PASS(HANDLE) :: OPEN
  ...
END TYPE
```

For a more elaborate example see C.1.3.

13 4.5.7.2 Inheritance

- 14 An extended type includes all of the type parameters, all of the components, and the nonoverridden
15 (4.5.7.3) nonfinal procedure bindings of its parent type. These are **inherited** by the extended type from
16 the parent type. They retain all of the attributes that they had in the parent type. Additional type
17 parameters, components, and procedure bindings may be declared in the derived-type definition of the
18 extended type.

NOTE 4.56

Inaccessible components and bindings of the parent type are also inherited, but they remain inaccessible in the extended type. Inaccessible entities occur if the type being extended is accessed via use association and has a private entity.

NOTE 4.57

A base type is not required to have any components, bindings, or parameters; an extended type is not required to have more components, bindings, or parameters than its parent type.

- 1 An extended type has a scalar, nonpointer, nonallocatable, **parent component** with the type and
 2 type parameters of the parent type. The name of this component is the parent type name. It has the
 3 accessibility of the parent type. Components of the parent component are **inheritance associated**
 4 (16.5.4) with the corresponding components inherited from the parent type. An **ancestor component**
 5 of a type is the parent component of the type or an ancestor component of the parent component.

NOTE 4.58

A component or type parameter declared in an extended type shall not have the same name as any accessible component or type parameter of its parent type.

NOTE 4.59

Examples:

```

TYPE POINT                                ! A base type
  REAL :: X, Y
END TYPE POINT

TYPE, EXTENDS(POINT) :: COLOR_POINT      ! An extension of TYPE(POINT)
  ! Components X and Y, and component name POINT, inherited from parent
  INTEGER :: COLOR
END TYPE COLOR_POINT

```

6 4.5.7.3 Type-bound procedure overriding

- 7 If a nongeneric binding specified in a type definition has the same binding name as a binding from the
 8 parent type then the binding specified in the type definition **overrides** the one from the parent type.
- 9 The overriding binding and the overridden binding shall satisfy the following conditions.
- 10 (1) Either both shall have a passed-object dummy argument or neither shall.
 - 11 (2) If the overridden binding is pure then the overriding binding shall also be pure.
 - 12 (3) Either both shall be elemental or neither shall.
 - 13 (4) They shall have the same number of dummy arguments.
 - 14 (5) Passed-object dummy arguments, if any, shall correspond by name and position.
 - 15 (6) Dummy arguments that correspond by position shall have the same names and characteris-
 16 tics, except for the type of the passed-object dummy arguments.
 - 17 (7) Either both shall be subroutines or both shall be functions having the same result charac-
 18 teristics (12.3.3).
 - 19 (8) If the overridden binding is PUBLIC then the overriding binding shall not be PRIVATE.

NOTE 4.60

The following is an example of procedure overriding, expanding on the example in Note 4.47.

```
TYPE, EXTENDS (POINT) :: POINT_3D
  REAL :: Z
CONTAINS
  PROCEDURE, PASS :: LENGTH => POINT_3D_LENGTH
END TYPE POINT_3D
...
```

and in the *module-subprogram-part* of the same module:

```
REAL FUNCTION POINT_3D_LENGTH ( A, B )
  CLASS (POINT_3D), INTENT (IN) :: A
  CLASS (POINT), INTENT (IN) :: B
  SELECT TYPE(B)
    CLASS IS(POINT_3D)
      POINT_3D_LENGTH = SQRT( (A%X-B%X)**2 + (A%Y-B%Y)**2 + (A%Z-B%Z)**2 )
      RETURN
  END SELECT
  PRINT *, 'In POINT_3D_LENGTH, dynamic type of argument is incorrect.'
  STOP
END FUNCTION POINT_3D_LENGTH
```

- 1 If a generic binding specified in a type definition has the same *generic-spec* as an inherited binding, it
- 2 extends the generic interface and shall satisfy the requirements specified in 12.4.3.3.4.
- 3 If a generic binding in a type definition has the same *dtio-generic-spec* as one inherited from the parent, it
- 4 extends the type-bound generic interface for *dtio-generic-spec* and shall satisfy the requirements specified
- 5 in 12.4.3.3.4.
- 6 A binding of a type and a binding of an extension of that type correspond if the latter binding is the
- 7 same binding as the former, overrides a corresponding binding, or is an inherited corresponding binding.

8 **4.5.8 Derived-type values**

9 The **component value** of

- 10 (1) a pointer component is its pointer association,
- 11 (2) an allocatable component is its allocation status and, if it is allocated, its dynamic type and
- 12 type parameters, bounds and value, and
- 13 (3) a nonpointer nonallocatable component is its value.

14 The set of values of a particular derived type consists of all possible sequences of the component values

15 of its components.

16 **4.5.9 Derived-type specifier**

17 A derived-type specifier is used in several contexts to specify a particular derived type and type param-

18 eters.

- 19 R457 *derived-type-spec* **is** *type-name* [(*type-param-spec-list*)]
- 20 R458 *type-param-spec* **is** [*keyword* =] *type-param-value*

C487 (R457) *type-name* shall be the name of an accessible derived type.

- 1 C488 (R457) *type-param-spec-list* shall appear only if the type is parameterized.
- 2 C489 (R457) There shall be at most one *type-param-spec* corresponding to each parameter of the type.
3 If a type parameter does not have a default value, there shall be a *type-param-spec* corresponding
4 to that type parameter.
- 5 C490 (R458) The *keyword=* may be omitted from a *type-param-spec* only if the *keyword=* has been
6 omitted from each preceding *type-param-spec* in the *type-param-spec-list*.
- 7 C491 (R458) Each *keyword* shall be the name of a parameter of the type.
- 8 C492 (R458) An asterisk may be used as a *type-param-value* in a *type-param-spec* only in the decla-
9 ration of a dummy argument or associate name or in the allocation of a dummy argument.
- 10 Type parameter values that do not have type parameter keywords specified correspond to type param-
11 eters in type parameter order (4.5.3.2). If a type parameter keyword is present, the value is assigned to
12 the type parameter named by the keyword. If necessary, the value is converted according to the rules of
13 intrinsic assignment (7.4.1.3) to a value of the same kind as the type parameter.
- 14 The value of a type parameter for which no *type-param-value* has been specified is its default value.

15 4.5.10 Construction of derived-type values

16 A derived-type definition implicitly defines a corresponding **structure constructor** that allows con-
17 struction of values of that derived type. The type and type parameters of a constructed value are
18 specified by a derived type specifier.

- | | | | |
|---------|------------------------------|-----------|---|
| 19 R459 | <i>structure-constructor</i> | is | <i>derived-type-spec</i> ([<i>component-spec-list</i>]) |
| 20 R460 | <i>component-spec</i> | is | [<i>keyword =</i>] <i>component-data-source</i> |
| 21 R461 | <i>component-data-source</i> | is | <i>expr</i> |
| 22 | | or | <i>data-target</i> |
| 23 | | or | <i>proc-target</i> |

- 24 C493 (R459) The *derived-type-spec* shall not specify an abstract type (4.5.7).
- 25 C494 (R459) At most one *component-spec* shall be provided for a component.
- 26 C495 (R459) If a *component-spec* is provided for an ancestor component, a *component-spec* shall not
27 be provided for any component that is inheritance associated with a subcomponent of that
28 ancestor component.
- 29 C496 (R459) A *component-spec* shall be provided for a nonallocatable component unless it has default
30 initialization or is inheritance associated with a subcomponent of another component for which
31 a *component-spec* is provided.
- 32 C497 (R460) The *keyword=* may be omitted from a *component-spec* only if the *keyword=* has been
33 omitted from each preceding *component-spec* in the constructor.
- 34 C498 (R460) Each *keyword* shall be the name of a component of the type.
- 35 C499 (R459) The type name and all components of the type for which a *component-spec* appears shall
36 be accessible in the scoping unit containing the structure constructor.
- 37 C4100 (R459) If *derived-type-spec* is a type name that is the same as a generic name, the *component-*
38 *spec-list* shall not be a valid *actual-arg-spec-list* for a function reference that is resolvable as a
39 generic reference to that name (12.5.5.2).
- 40 C4101 (R461) A *data-target* shall correspond to a nonprocedure pointer component; a *proc-target* shall

- 1 correspond to a procedure pointer component.
 2 C4102 (R461) A *data-target* shall have the same rank as its corresponding component.

NOTE 4.61

The form 'name(...)' is interpreted as a generic *function-reference* if possible; it is interpreted as a *structure-constructor* only if it cannot be interpreted as a generic *function-reference*.

- 3 In the absence of a component keyword, each *component-data-source* is assigned to the corresponding
 4 component in component order (4.5.4.6). If a component keyword appears, the *expr* is assigned to the
 5 component named by the keyword. For a nonpointer component, the declared type and type parameters
 6 of the component and *expr* shall conform in the same way as for a *variable* and *expr* in an intrinsic
 7 assignment statement (7.4.1.2), as specified in Table 7.11. If necessary, each value of intrinsic type is
 8 converted according to the rules of intrinsic assignment (7.4.1.3) to a value that agrees in type and type
 9 parameters with the corresponding component of the derived type. For a nonpointer nonallocatable
 10 component, the shape of the expression shall conform with the shape of the component.
 11 If a component with default initialization has no corresponding *component-data-source*, then the default
 12 initialization is applied to that component. If an allocatable component has no corresponding *component-*
 13 *data-source*, then that component has an allocation status of unallocated.

NOTE 4.62

Because no parent components appear in the defined component ordering, a value for a parent component can be specified only with a component keyword. Examples of equivalent values using types defined in Note 4.59:

```
! Create values with components x = 1.0, y = 2.0, color = 3.
TYPE(PPOINT) :: PV = PPOINT(1.0, 2.0)      ! Assume components of TYPE(PPOINT)
                                           ! are accessible here.
...
COLOR_POINT( point=point(1,2), color=3)    ! Value for parent component
COLOR_POINT( point=PV, color=3)            ! Available even if TYPE(point)
                                           ! has private components
COLOR_POINT( 1, 2, 3)                      ! All components of TYPE(point)
                                           ! need to be accessible.
```

- 14 A structure constructor shall not appear before the referenced type is defined.

NOTE 4.63

This example illustrates a derived-type constant expression using a derived type defined in Note 4.21:

```
PERSON (21, 'JOHN SMITH')
```

This could also be written as

```
PERSON (NAME = 'JOHN SMITH', AGE = 21)
```

NOTE 4.64

An example constructor using the derived type GENERAL_POINT defined in Note 4.28 is

```
general_point(dim=3) ( (/ 1., 2., 3. /) )
```

- 1 For a pointer component, the corresponding *component-data-source* shall be an allowable *data-target* or
 2 *proc-target* for such a pointer in a pointer assignment statement (7.4.2). If the component data source is
 3 a pointer, the association of the component is that of the pointer; otherwise, the component is pointer
 4 associated with the component data source.

NOTE 4.65

For example, if the variable TEXT were declared (5.2) to be

```
CHARACTER, DIMENSION (1:400), TARGET :: TEXT
```

and BIBLIO were declared using the derived-type definition REFERENCE in Note 4.35

```
TYPE (REFERENCE) :: BIBLIO
```

the statement

```
BIBLIO = REFERENCE (1, 1987, 1, "This is the title of the referenced &  

&paper", TEXT)
```

is valid and associates the pointer component SYNOPSIS of the object BIBLIO with the target object TEXT.

- 5 If a component of a derived type is allocatable, the corresponding constructor expression shall either be
 6 a reference to the intrinsic function NULL with no arguments, an allocatable entity of the same rank,
 7 or shall evaluate to an entity of the same rank. If the expression is a reference to the intrinsic function
 8 NULL, the corresponding component of the constructor has a status of unallocated. If the expression
 9 is an allocatable entity, the corresponding component of the constructor has the same allocation status
 10 as that allocatable entity and, if it is allocated, the same dynamic type, bounds, and value; if a length
 11 parameter of the component is deferred, its value is the same as the corresponding parameter of the
 12 expression. Otherwise the corresponding component of the constructor has an allocation status of
 13 allocated and has the same bounds and value as the expression.

NOTE 4.66

When the constructor is an actual argument, the allocation status of the allocatable component is available through the associated dummy argument.

14 4.5.11 Derived-type operations and assignment

- 15 Intrinsic assignment of derived-type entities is described in 7.4.1. This part of ISO/IEC 1539 does
 16 not specify any intrinsic operations on derived-type entities. Any operation on derived-type entities
 17 or defined assignment (7.4.1.4) for derived-type entities shall be defined explicitly by a function or a
 18 subroutine, and a generic interface (4.5.2, 12.4.3.2).

19 4.6 Enumerations and enumerators

- 20 An enumeration is a set of enumerators. An enumerator is a named integer constant. An enumeration
 21 definition specifies the enumeration and its set of enumerators of the corresponding integer kind.

```
22 R462  enum-def          is  enum-def-stmt  

23                                     enumerator-def-stmt  

24                                     [ enumerator-def-stmt ] ...  

25                                     end-enum-stmt
```

1 R463 *enum-def-stmt* is ENUM, BIND(C)
 2 R464 *enumerator-def-stmt* is ENUMERATOR [::] *enumerator-list*
 3 R465 *enumerator* is *named-constant* [= *scalar-int-initialization-expr*]
 4 R466 *end-enum-stmt* is END ENUM

5 C4103 (R464) If = appears in an *enumerator*, a double-colon separator shall appear before the *enu-*
 6 *merator-list*.

7 For an enumeration, the kind is selected such that an integer type with that kind is interoperable (15.3.2)
 8 with the corresponding C enumeration type. The corresponding C enumeration type is the type that
 9 would be declared by a C enumeration specifier (6.7.2.2 of the C International Standard) that specified
 10 C enumeration constants with the same values as those specified by the *enum-def*, in the same order as
 11 specified by the *enum-def*.

12 The companion processor (2.5.10) shall be one that uses the same representation for the types declared
 13 by all C enumeration specifiers that specify the same values in the same order.

NOTE 4.67

If a companion processor uses an unsigned type to represent a given enumeration type, the Fortran processor will use the signed integer type of the same width for the enumeration, even though some of the values of the enumerators cannot be represented in this signed integer type. The values of any such enumerators will be interoperable with the values declared in the C enumeration.

NOTE 4.68

The C International Standard guarantees the enumeration constants fit in a C int (6.7.2.2 of the C International Standard). Therefore, the Fortran processor can evaluate all enumerator values using the integer type with kind parameter C.INT, and then determine the kind parameter of the integer type that is interoperable with the corresponding C enumerated type.

NOTE 4.69

The C International Standard specifies that two enumeration types are compatible only if they specify enumeration constants with the same names and same values in the same order. This part of ISO/IEC 1539 further requires that a C processor that is to be a companion processor of a Fortran processor use the same representation for two enumeration types if they both specify enumeration constants with the same values in the same order, even if the names are different.

14 An enumerator is treated as if it were explicitly declared with the PARAMETER attribute. The enu-
 15 merator is defined in accordance with the rules of intrinsic assignment (7.4) with the value determined
 16 as follows.

- 17 (1) If *scalar-int-initialization-expr* is specified, the value of the enumerator is the result of
 18 *scalar-int-initialization-expr*.
 19 (2) If *scalar-int-initialization-expr* is not specified and the enumerator is the first enumerator
 20 in *enum-def*, the enumerator has the value 0.
 21 (3) If *scalar-int-initialization-expr* is not specified and the enumerator is not the first enumer-
 22 ator in *enum-def*, its value is the result of adding 1 to the value of the enumerator that
 23 immediately precedes it in the *enum-def*.

NOTE 4.70

Example of an enumeration definition:

```
ENUM, BIND(C)
  ENUMERATOR :: RED = 4, BLUE = 9
```


NOTE 4.70 (cont.)

```

ENUMERATOR YELLOW
END ENUM

```

The kind type parameter for this enumeration is processor dependent, but the processor is required to select a kind sufficient to represent the values 4, 9, and 10, which are the values of its enumerators. The following declaration might be equivalent to the above enumeration definition.

```

INTEGER(SELECTED_INT_KIND(2)), PARAMETER :: RED = 4, BLUE = 9, YELLOW = 10

```

An entity of the same kind type parameter value can be declared using the intrinsic function KIND with one of the enumerators as its argument, for example

```

INTEGER(KIND(RED)) :: X

```

NOTE 4.71

There is no difference in the effect of declaring the enumerators in multiple ENUMERATOR statements or in a single ENUMERATOR statement. The order in which the enumerators in an enumeration definition are declared is significant, but the number of ENUMERATOR statements is not.

1 4.7 Construction of array values

2 An **array constructor** is defined as a sequence of scalar values and is interpreted as a rank-one array
 3 where the element values are those specified in the sequence.

4	R467	<i>array-constructor</i>	is (<i>/ ac-spec /</i>)
5			or <i>lbracket ac-spec rbracket</i>
6	R468	<i>ac-spec</i>	is <i>type-spec ::</i>
7			or [<i>type-spec ::</i>] <i>ac-value-list</i>
8	R469	<i>lbracket</i>	is [
9	R470	<i>rbracket</i>	is]
10	R471	<i>ac-value</i>	is <i>expr</i>
11			or <i>ac-implied-do</i>
12	R472	<i>ac-implied-do</i>	is (<i>ac-value-list</i> , <i>ac-implied-do-control</i>)
13	R473	<i>ac-implied-do-control</i>	is <i>ac-do-variable</i> = <i>scalar-int-expr</i> , <i>scalar-int-expr</i> ■
14			■ [, <i>scalar-int-expr</i>]
15	R474	<i>ac-do-variable</i>	is <i>do-variable</i>

16 C4104 (R468) If *type-spec* is omitted, each *ac-value* expression in the *array-constructor* shall have the
 17 same type and kind type parameters.

18 C4105 (R468) If *type-spec* specifies an intrinsic type, each *ac-value* expression in the *array-constructor*
 19 shall be of an intrinsic type that is in type conformance with a variable of type *type-spec* as
 20 specified in Table 7.11.

21 C4106 (R468) If *type-spec* specifies a derived type, all *ac-value* expressions in the *array-constructor*
 22 shall be of that derived type and shall have the same kind type parameter values as specified by
 23 *type-spec*.

24 C4107 (R472) The *ac-do-variable* of an *ac-implied-do* that is in another *ac-implied-do* shall not appear
 25 as the *ac-do-variable* of the containing *ac-implied-do*.

- 1 If *type-spec* is omitted, each *ac-value* expression in the array constructor shall have the same length type
 2 parameters; in this case, the type and type parameters of the array constructor are those of the *ac-value*
 3 expressions.
- 4 If *type-spec* appears, it specifies the type and type parameters of the array constructor. Each *ac-value*
 5 expression in the *array-constructor* shall be compatible with intrinsic assignment to a variable of this
 6 type and type parameters. Each value is converted to the type parameters of the *array-constructor* in
 7 accordance with the rules of intrinsic assignment (7.4.1.3).
- 8 The character length of an *ac-value* in an *ac-implied-do* whose iteration count is zero shall not depend
 9 on the value of the *ac-do-variable* and shall not depend on the value of an expression that is not an
 10 initialization expression.
- 11 If an *ac-value* is a scalar expression, its value specifies an element of the array constructor. If an *ac-*
 12 *value* is an array expression, the values of the elements of the expression, in array element order (6.2.2.2),
 13 specify the corresponding sequence of elements of the array constructor. If an *ac-value* is an *ac-implied-*
 14 *do*, it is expanded to form a sequence of elements under the control of the *ac-do-variable*, as in the DO
 15 construct (8.1.7.5).
- 16 For an *ac-implied-do*, the loop initialization and execution is the same as for a DO construct.
- 17 An empty sequence forms a zero-sized rank-one array.

NOTE 4.72

A one-dimensional array may be reshaped into any allowable array shape using the RESHAPE intrinsic function (13.7.146). An example is:

```
X = (/ 3.2, 4.01, 6.5 /)
Y = RESHAPE (SOURCE = [ 2.0, [ 4.5, 4.5 ], X ], SHAPE = [ 3, 2 ])
```

This results in Y having the 3 × 2 array of values:

```
2.0  3.2
4.5  4.01
4.5  6.5
```

NOTE 4.73

Examples of array constructors containing an implied DO are:

```
(/ (I, I = 1, 1075) /)
```

and

```
[ 3.6, (3.6 / I, I = 1, N) ]
```

NOTE 4.74

Using the type definition for PERSON in Note 4.21, an example of the construction of a derived-type array value is:

```
(/ PERSON (40, 'SMITH'), PERSON (20, 'JONES') /)
```

NOTE 4.75

Using the type definition for LINE in Note 4.32, an example of the construction of a derived-type scalar value with a rank-2 array component is:

```
LINE (RESHAPE ( (/ 0.0, 0.0, 1.0, 2.0 /), (/ 2, 2 /) ), 0.1, 1)
```

The RESHAPE intrinsic function is used to construct a value that represents a solid line from (0, 0) to (1, 2) of width 0.1 centimeters.

NOTE 4.76

Examples of zero-size array constructors are:

```
(/ INTEGER :: /)  
(/ ( I, I = 1, 0) /)
```

NOTE 4.77

An example of an array constructor that specifies a length type parameter:

```
(/ CHARACTER(LEN=7) :: 'Takata', 'Tanaka', 'Hayashi' /)
```

In this constructor, without the type specification, it would have been necessary to specify all of the constants with the same character length.

1 5 Attribute declarations and specifications

2 5.1 General

3 Every data object has a type and rank and may have type parameters and other attributes that determine
4 the uses of the object. Collectively, these properties are the **attribute** of the object. The type of a
5 named data object is either specified explicitly in a type declaration statement or determined implicitly
6 by the first letter of its name (5.5). All of its attributes may be specified in a type declaration statement
7 or individually in separate specification statements.

8 A function has a type and rank and may have type parameters and other attributes that determine the
9 uses of the function. The type, rank, and type parameters are the same as those of its result variable.

10 A subroutine does not have a type, rank, or type parameters, but may have other attributes that
11 determine the uses of the subroutine.

12 5.2 Type declaration statements

13 5.2.1 Syntax

14 R501 *type-declaration-stmt* **is** *declaration-type-spec* [[, *attr-spec*] ... ::] *entity-decl-list*

15 The type declaration statement specifies the type of the entities in the entity declaration list. The type
16 and type parameters are those specified by *declaration-type-spec*, except that the character length type
17 parameter may be overridden for an entity by the appearance of * *char-length* in its *entity-decl*.

18 R502 *attr-spec* **is** *access-spec*
19 **or** ALLOCATABLE
20 **or** ASYNCHRONOUS
21 **or** CONTIGUOUS
22 **or** dimension-spec
23 **or** EXTERNAL
24 **or** INTENT (*intent-spec*)
25 **or** INTRINSIC
26 **or** *language-binding-spec*
27 **or** OPTIONAL
28 **or** PARAMETER
29 **or** POINTER
30 **or** PROTECTED
31 **or** SAVE
32 **or** TARGET
33 **or** VALUE
34 **or** VOLATILE
35

36 C501 (R501) The same *attr-spec* shall not appear more than once in a given *type-declaration-stmt*.

37 C502 (R501) If a *language-binding-spec* with a NAME= specifier appears, the *entity-decl-list* shall
38 consist of a single *entity-decl*.

39 C503 (R501) If a *language-binding-spec* is specified, the *entity-decl-list* shall not contain any procedure

1 names.

2 The type declaration statement also specifies the attributes whose keywords appear in the *attr-spec*,
3 except that the DIMENSION attribute may be specified or overridden for an entity by the appearance
4 of *array-spec* in its *entity-decl*.

5 R503 *entity-decl* **is** *object-name* [(*array-spec*)] ■
6 ■ [*lbracket co-array-spec rbracket*] ■
7 ■ [* *char-length*] [*initialization*]
8 **or** *function-name* [* *char-length*]

9 C504 (R503) If the entity is not of type character, * *char-length* shall not appear.

10 C505 (R501) If *initialization* appears, a double-colon separator shall appear before the *entity-decl-list*.

11 C506 (R503) An *initialization* shall not appear if *object-name* is a dummy argument, a function result,
12 an object in a named common block unless the type declaration is in a block data program unit,
13 an object in blank common, an allocatable variable, an external function, an intrinsic function,
14 or an automatic object.

15 C507 (R503) An *initialization* shall appear if the entity is a named constant (5.3.12).

16 C508 (R503) The *function-name* shall be the name of an external function, an intrinsic function, a
17 function dummy procedure, a procedure pointer, or a statement function.

18 R504 *object-name* **is** *name*

19 C509 (R504) The *object-name* shall be the name of a data object.

20 R505 *initialization* **is** = *initialization-expr*
21 **or** => *null-init*
22 **or** => *initial-data-target*

23 R506 *null-init* **is** *function-reference*

24 C510 (R503) If => appears in *initialization*, the entity shall have the POINTER attribute. If =
25 appears in *initialization*, the entity shall not have the POINTER attribute.

26 C511 (R503) If *initial-data-target* appears, *object-name* shall be data-pointer-initialization compatible
27 with it (4.5.4.5).

28 C512 (R506) The *function-reference* shall be a reference to the NULL intrinsic function with no
29 arguments.

30 A name that identifies a specific intrinsic function in a scoping unit has a type as specified in 13.6. An
31 explicit type declaration statement is not required; however, it is permitted. Specifying a type for a
32 generic intrinsic function name in a type declaration statement is not sufficient, by itself, to remove the
33 generic properties from that function.

34 5.2.2 Automatic data objects

35 An **automatic data object** is a nondummy data object with a type parameter or array bound that
36 depends on the value of a *specification-expr* that is not an initialization expression.

37 C513 An automatic object shall not be a local variable of a main program, module, or submodule.

NOTE 5.1

An automatic object shall not have the SAVE attribute and shall not appear in a common block.

1 If a type parameter in a *declaration-type-spec* or in a *char-length* in an *entity-decl* is defined by an
 2 expression that is not an initialization expression, the type parameter value is established on entry to
 3 the procedure and is not affected by any redefinition or undefinition of the variables in the expression
 4 during execution of the procedure.

5.2.3 Initialization

6 The appearance of *initialization* in an *entity-decl* for an entity without the PARAMETER attribute
 7 specifies that the entity is a variable with **explicit initialization**. Explicit initialization alternatively
 8 may be specified in a DATA statement unless the variable is of a derived type for which default initial-
 9 ization is specified. If *initialization* is *=initialization-expr*, the variable is initially defined with the value
 10 specified by the *initialization-expr*; if necessary, the value is converted according to the rules of intrinsic
 11 assignment (7.4.1.3) to a value that agrees in type, type parameters, and shape with the variable. A
 12 variable, or part of a variable, shall not be explicitly initialized more than once in a program. If the
 13 variable is an array, it shall have its shape specified in either the type declaration statement or a previous
 14 attribute specification statement in the same scoping unit.

15 If *null-init* appears, the initial association status of the object is disassociated. If *initial-data-target*
 16 appears, the object is initially associated with the target.

17 Explicit initialization of a variable that is not in a common block implies the SAVE attribute, which
 18 may be confirmed by explicit specification.

5.2.4 Examples of type declaration statements**NOTE 5.2**

```

REAL A (10)
LOGICAL, DIMENSION (5, 5) :: MASK1, MASK2
COMPLEX :: CUBE_ROOT = (-0.5, 0.866)
INTEGER, PARAMETER :: SHORT = SELECTED_INT_KIND (4)
INTEGER (SHORT) K      ! Range at least -9999 to 9999.
REAL (KIND (0.0D0)) A
REAL (KIND = 2) B
COMPLEX (KIND = KIND (0.0D0)) :: C
CHARACTER (LEN = 10, KIND = 2) A
CHARACTER B, C *20
TYPE (PERSON) :: CHAIRMAN
TYPE(NODE), POINTER :: HEAD => NULL ( )
TYPE (humongous_matrix (k=8, d=1000)) :: mat

```

(The last line above uses a type definition from Note 4.28.)

5.3 Attributes**5.3.1 Constraints**

22 An attribute may be explicitly specified by an *attr-spec* in a type declaration statement or by an attribute
 specification statement (5.4). The following constraints apply to attributes.

1

2 C514 An entity shall not be explicitly given any attribute more than once in a scoping unit.

3 C515 An *array-spec* for a function result that does not have the ALLOCATABLE or POINTER
4 attribute shall be an *explicit-shape-spec-list*.5 C516 The ALLOCATABLE, POINTER, or OPTIONAL attribute shall not be specified for a dummy
6 argument of a procedure that has a *proc-language-binding-spec*.7 **5.3.2 Accessibility attribute**8 The **accessibility attribute** specifies the accessibility of an entity via a particular identifier.9 R507 *access-spec* **is** PUBLIC
10 **or** PRIVATE11 C517 (R507) An *access-spec* shall appear only in the *specification-part* of a module.12 Identifiers that are specified in a module or accessible in that module by use association have either
13 the PUBLIC or PRIVATE attribute. Identifiers for which an *access-spec* is not explicitly specified in
14 that module have the default accessibility attribute for that module. The default accessibility attribute
15 for a module is PUBLIC unless it has been changed by a PRIVATE statement (5.4.1). Only identifiers
16 that have the PUBLIC attribute in that module are available to be accessed from that module by use
17 association.**NOTE 5.3**

In order for an identifier to be accessed by use association, it must have the PUBLIC attribute in the module from which it is accessed. It can nonetheless have the PRIVATE attribute in a module in which it is accessed by use association, and therefore not be available for use association from a module where it is PRIVATE.

NOTE 5.4

An example of an accessibility specification is:

REAL, PRIVATE :: X, Y, Z

18 **5.3.3 ALLOCATABLE attribute**19 An entity with the **ALLOCATABLE attribute** is a variable for which space is allocated by an AL-
20 LOCATE statement (6.3.1) or by an intrinsic assignment statement (7.4.1.3).21 **5.3.4 ASYNCHRONOUS attribute**22 An entity with the **ASYNCHRONOUS attribute** is a variable that may be subject to asynchronous
23 input/output.

24 The base object of a variable shall have the ASYNCHRONOUS attribute in a scoping unit if

- 25 (1) the variable appears in an executable statement or specification expression in that scoping
-
- 26 unit and
-
- 27 (2) any statement of the scoping unit is executed while the variable is a pending I/O storage
-
- 28 sequence affector (9.5.2.5).

29 Use of a variable in an asynchronous input/output statement can imply the ASYNCHRONOUS attribute;

- 1 see subclause (9.5.2.5).
- 2 An object may have the ASYNCHRONOUS attribute in a particular scoping unit without necessarily
 3 having it in other scoping units (11.2.2, 16.5.1.4). If an object has the ASYNCHRONOUS attribute,
 4 then all of its subobjects also have the ASYNCHRONOUS attribute.

NOTE 5.5

The ASYNCHRONOUS attribute specifies the variables that might be associated with a pending input/output storage sequence (the actual memory locations on which asynchronous input/output is being performed) while the scoping unit is in execution. This information could be used by the compiler to disable certain code motion optimizations.

The ASYNCHRONOUS attribute is similar to the VOLATILE attribute. It is intended to facilitate traditional code motion optimizations in the presence of asynchronous input/output.

5.3.5 BIND attribute for data entities

- 6 The BIND attribute for a variable or common block specifies that it is capable of interoperating with a
 7 C variable that has external linkage (15.4).
- 8 R508 *language-binding-spec* is BIND (C [, NAME = *scalar-char-initialization-expr*])
- 9 C518 An entity with the BIND attribute shall be a common block, variable, or procedure.
- 10 C519 A variable with the BIND attribute shall be declared in the specification part of a module.
- 11 C520 A variable with the BIND attribute shall be interoperable (15.3).
- 12 C521 Each variable of a common block with the BIND attribute shall be interoperable.
- 13 C522 (R508) The *scalar-char-initialization-expr* shall be of default character kind. If the value of the
 14 *scalar-char-initialization-expr* after discarding leading and trailing blanks has nonzero length,
 15 it shall be valid as an identifier on the companion processor.

NOTE 5.6

The C International Standard provides a facility for creating C identifiers whose characters are not restricted to the C basic character set. Such a C identifier is referred to as a universal character name (6.4.3 of the C International Standard). The name of such a C identifier might include characters that are not part of the representation method used by the processor for type default character. If so, the C entity cannot be referenced from Fortran.

- 16 The BIND attribute for a variable or common block implies the SAVE attribute, which may be confirmed
 17 by explicit specification.

5.3.6 CONTIGUOUS attribute

- 19 C523 An entity that has the CONTIGUOUS attribute shall be an array pointer or an assumed-shape
 20 array.
- 21 The **CONTIGUOUS attribute** specifies that an assumed-shape array can only be argument associated
 22 with a contiguous object, or that an array pointer can only be pointer associated with a contiguous target.
- 23 An object is **contiguous** if it is
- 24 (1) an object with the CONTIGUOUS attribute,
 25 (2) a scalar object,

- 1 (3) a nonpointer whole array that is not assumed-shape,
 2 (4) an array allocated by an ALLOCATE statement,
 3 (5) an assumed-shape array that is argument associated with an array that is contiguous,
 4 (6) a pointer associated with a contiguous target,
 5 (7) an array with at most one element, or
 6 (8) a nonzero-sized array section (6.2.2) with the following properties:
 7 (a) Its base object is contiguous.
 8 (b) It does not have a vector subscript.
 9 (c) The elements of the section, in array element order, are a subset of the base object
 10 elements that are consecutive in array element order.
 11 (d) If the array is of type character and a *substring-range* appears, the *substring-range*
 12 specifies all of the characters of the *parent-string* (6.1.1).
 13 (e) Only its final *part-ref* has nonzero rank.
 14 (f) It is not the real or imaginary part (6.1.3) of an array of type complex.

J3 internal note

Unresolved Technical Issue 011

The above list (CONTIGUOUS definition) does not conform to the ISO guidelines.

15 An object is not contiguous if it is an array subobject, and

- 16 (1) the object has two or more elements,
 17 (2) the elements of the object in array element order are not consecutive in the elements of the
 18 base object,
 19 (3) the object is not of type character with length zero, and
 20 (4) the object is not of a derived type that has no ultimate components other than zero-sized
 21 arrays and characters with length zero.

22 It is processor-dependent whether any other object is contiguous.

NOTE 5.7

If a derived type has only one component that is not zero-sized, it is processor-dependent whether a structure component of a contiguous array of that type is contiguous. That is, the derived type might contain padding on some processors.

NOTE 5.8

The CONTIGUOUS attribute allows a processor to enable optimizations that depend on the memory layout of the object occupying a contiguous block of memory. Examples of CONTIGUOUS attribute specifications are:

```
REAL, POINTER, CONTIGUOUS      :: SPTR(:)
REAL, CONTIGUOUS, DIMENSION(:, :) :: D
```

23 5.3.7 DIMENSION attribute**24 5.3.7.1 General**

25 The **DIMENSION attribute** specifies that an entity is an array, a co-array, or both. If an *array-spec*
 26 appears, it is an array. If a *co-array-spec* appears, it is a co-array.

27 For an array, its *array-spec* specifies its rank or rank and shape. For a co-array, its *co-array-spec* specifies

1 its co-rank or co-rank and co-bounds.

NOTE 5.9

Unless it is a dummy argument, a co-array has the same bounds and co-bounds on every image.

See Note 12.31 for further discussion of the bounds and co-bounds of dummy co-arrays.

2 R509 *dimension-spec* **is** DIMENSION (*array-spec*)
3 **or** DIMENSION [(*array-spec*)] *lbracket co-array-spec rbracket*

4 C524 (R501) A co-array that has the ALLOCATABLE attribute shall be specified with a *co-array-spec*
5 that is a *deferred-co-shape-spec-list*.

6 C525 A co-array shall be a component or a variable that is not a function result.

7 C526 A co-array shall not be of type IMAGE_TEAM (13.8.3.7), C_PTR, or C_FUNPTR (15.3.3).

8 C527 An entity whose type has a co-array ultimate component shall be a nonpointer nonallocatable
9 scalar, shall not be a co-array, and shall not be a function result.

NOTE 5.10

A co-array is permitted to be of a derived type with pointer or allocatable components. The target of such a pointer component is always on the same image.

An allocatable component of a co-array need be allocated on an image only if it is referenced or defined.

10 C528 The SAVE attribute shall be specified for a co-array unless it is a dummy argument, declared
11 in the main program, or allocatable.

NOTE 5.11

This requirement for the SAVE attribute has the effect that automatic co-arrays are not permitted; for example, the co-array WORK in the following code fragment is not valid.

```
SUBROUTINE SOLVE3(N,A,B)
INTEGER :: N
REAL    :: A(N)[*], B(N)
REAL    :: WORK(N)[*]    ! Not permitted
```

If this were permitted, it would require an implicit synchronization on entry to the procedure.

Explicit-shape co-arrays that are declared in a subprogram and are not dummy arguments are required to have the SAVE attribute because otherwise they might be implemented as if they were automatic co-arrays.

12 R510 *array-spec* **is** *explicit-shape-spec-list*
13 **or** *assumed-shape-spec-list*
14 **or** *deferred-shape-spec-list*
15 **or** *assumed-size-spec*
16 **or** *implied-shape-spec-list*

17 R511 *co-array-spec* **is** *deferred-co-shape-spec-list*
18 **or** *explicit-co-shape-spec*

C529 The sum of the rank and co-rank of an entity shall not exceed fifteen.

NOTE 5.12

Examples of DIMENSION attribute specifications are:

```

SUBROUTINE EX (N, A, B)
  REAL, DIMENSION (N, 10) :: W           ! Automatic explicit-shape array
  REAL A (:), B (0:)                    ! Assumed-shape arrays
  REAL, POINTER :: D (:, :)             ! Array pointer
  REAL, DIMENSION (:), POINTER :: P     ! Array pointer
  REAL, ALLOCATABLE, DIMENSION (:) :: E ! Allocatable array
  REAL, PARAMETER :: V(0:*) = [0.1, 1.1] ! Implied-shape array

```

1 **5.3.7.2 Explicit-shape array**

2 An **explicit-shape array** is a named array that is declared with an *explicit-shape-spec-list*. This specifies
3 explicit values for the bounds in each dimension of the array.

4 R512 *explicit-shape-spec* **is** [*lower-bound* :] *upper-bound*

5 R513 *lower-bound* **is** *specification-expr*

6 R514 *upper-bound* **is** *specification-expr*

7 C530 (R512) An *explicit-shape-spec* whose bounds are not initialization expressions shall appear only
8 in a subprogram or interface body.

9 An explicit-shape array that is a local variable of a subprogram or BLOCK construct may have bounds
10 that are not initialization expressions. The bounds, and hence shape, are determined at entry to a
11 procedure defined by the subprogram, or on execution of the BLOCK statement, by evaluating the
12 bounds expressions. The bounds of such an array are unaffected by the redefinition or undefinition of
13 any variable during execution of the procedure or BLOCK construct.

14 The values of each *lower-bound* and *upper-bound* determine the bounds of the array along a particular
15 dimension and hence the extent of the array in that dimension. The value of a lower bound or an upper
16 bound may be positive, negative, or zero. The subscript range of the array in that dimension is the set
17 of integer values between and including the lower and upper bounds, provided the upper bound is not
18 less than the lower bound. If the upper bound is less than the lower bound, the range is empty, the
19 extent in that dimension is zero, and the array is of zero size. If the *lower-bound* is omitted, the default
20 value is 1. The rank is equal to the number of *explicit-shape-specs*.

21 **5.3.7.3 Assumed-shape array**

22 An **assumed-shape array** is a nonallocatable nonpointer dummy argument array that takes its shape
23 from the associated actual argument array.

24 R515 *assumed-shape-spec* **is** [*lower-bound*] :

25 The rank is equal to the number of colons in the *assumed-shape-spec-list*.

26 The extent of a dimension of an assumed-shape array dummy argument is the extent of the corresponding
27 dimension of the associated actual argument array. If the lower bound value is d and the extent of the
28 corresponding dimension of the associated actual argument array is s , then the value of the upper bound
29 is $s + d - 1$. If *lower-bound* appears it specifies the lower bound; otherwise the lower bound is 1.

30 **5.3.7.4 Deferred-shape array**

31 A **deferred-shape array** is an allocatable array or an array pointer.

1 An **allocatable array** is an array that has the ALLOCATABLE attribute and a specified rank, but its
2 bounds, and hence shape, are determined by allocation or argument association.

3 An **array pointer** is an array with the POINTER attribute and a specified rank. Its bounds, and hence
4 shape, are determined when it is associated with a target.

5 R516 *deferred-shape-spec* **is** :

6 C531 An array that has the POINTER or ALLOCATABLE attribute shall have an *array-spec* that is
7 a *deferred-shape-spec-list*.

8 The rank is equal to the number of colons in the *deferred-shape-spec-list*.

9 The size, bounds, and shape of an unallocated allocatable array or a disassociated array pointer are
10 undefined. No part of such an array shall be referenced or defined; however, the array may appear as an
11 argument to an intrinsic inquiry function as specified in 13.1.

12 The bounds of each dimension of an allocated allocatable array are those specified when the array is
13 allocated.

14 The bounds of each dimension of an associated array pointer may be specified in two ways:

- 15 (1) in an ALLOCATE statement (6.3.1) when the target is allocated;
- 16 (2) by pointer assignment (7.4.2).

17 The bounds of the array pointer or allocatable array are unaffected by any subsequent redefinition or
18 undefinition of variables on which the bounds' expressions depend.

19 5.3.7.5 Assumed-size array

20 An **assumed-size array** is a dummy argument array whose size is assumed from that of an associated
21 actual argument. The rank and extents may differ for the actual and dummy arrays; only the size of the
22 actual array is assumed by the dummy array. An assumed-size array is declared with an *assumed-size-*
23 *spec*.

24 R517 *assumed-size-spec* **is** [*explicit-shape-spec* ,]... [*lower-bound* :] *

25 C532 An *assumed-size-spec* shall not appear except as the declaration of the array bounds of a dummy
26 data argument.

27 C533 An assumed-size array with INTENT (OUT) shall not be polymorphic, of a finalizable type, of
28 a type with an ultimate allocatable component, or of a type for which default initialization is
29 specified.

30 The size of an assumed-size array is determined as follows.

- 31 (1) If the actual argument associated with the assumed-size dummy array is an array of any
32 type other than default character, the size is that of the actual array.
- 33 (2) If the actual argument associated with the assumed-size dummy array is an array element
34 of any type other than default character with a subscript order value of r (6.2.2.2) in an
35 array of size x , the size of the dummy array is $x - r + 1$.
- 36 (3) If the actual argument is a default character array, default character array element, or a
37 default character array element substring (6.1.1), and if it begins at character storage unit t
38 of an array with c character storage units, the size of the dummy array is $\text{MAX}(\text{INT}((c -$
39 $t + 1)/e), 0)$, where e is the length of an element in the dummy character array.
- 40 (4) If the actual argument is of type default character and is a scalar that is not an array element
41 or array element substring designator, the size of the dummy array is $\text{MAX}(\text{INT}(l/e), 0)$,

1 where e is the length of an element in the dummy character array and l is the length of the
2 actual argument.

3 The rank is equal to one plus the number of *explicit-shape-specs*.

4 An assumed-size array has no upper bound in its last dimension and therefore has no extent in its last
5 dimension and no shape. An assumed-size array name shall not be written as a whole array reference
6 except as an actual argument in a procedure reference for which the shape is not required.

7 If a list of *explicit-shape-specs* appears, it specifies the bounds of the first rank – 1 dimensions. If *lower-*
8 *bound* appears it specifies the lower bound of the last dimension; otherwise that lower bound is 1. An
9 assumed-size array may be subscripted or sectioned (6.2.2.3). The upper bound shall not be omitted
10 from a subscript triplet in the last dimension.

11 If an assumed-size array has bounds that are not initialization expressions, the bounds are determined
12 at entry to the procedure. The bounds of such an array are unaffected by the redefinition or undefinition
13 of any variable during execution of the procedure.

14 5.3.7.6 Implied-shape array

15 An **implied-shape array** is a named constant that takes its shape from the *initialization-expr* in its
16 declaration. An implied-shape array is declared with an *implied-shape-spec-list*.

17 R518 *implied-shape-spec* **is** [*lower-bound* :] *

18 C534 An implied-shape array shall be a named constant.

19 The rank of an implied-shape array is the number of *implied-shape-specs* in the *implied-shape-spec-list*.

20 The extent of each dimension of an implied-shape array is the same as the extent of the corresponding
21 dimension of the *initialization-expr*. The lower bound of each dimension is *lower-bound*, if it appears,
22 and 1 otherwise; the upper bound is one less than the sum of the lower bound and the extent.

23 5.3.8 EXTERNAL attribute

24 The **EXTERNAL attribute** specifies that an entity is an external procedure, dummy procedure,
25 procedure pointer, or block data subprogram.

26 C535 An entity shall not have both the EXTERNAL attribute and the INTRINSIC attribute.

27 In an external subprogram, the EXTERNAL attribute shall not be specified for a procedure defined by
28 the subprogram.

29 If an external procedure or dummy procedure is used as an actual argument or is the target of a procedure
30 pointer assignment, it shall be declared to have the EXTERNAL attribute.

31 A procedure that has both the EXTERNAL and POINTER attributes is a procedure pointer.

32 5.3.8.1 Allocatable co-array

33 A co-array that has the ALLOCATABLE attribute has a specified co-rank, but its **co-bounds** are
34 determined by allocation or argument association.

35 R519 *deferred-co-shape-spec* **is** :

36 C536 A co-array that has the ALLOCATABLE attribute shall have a *co-array-spec* that is a *deferred-*
37 *co-shape-spec-list*.

1 The co-rank of an allocatable co-array is equal to the number of colons in its *deferred-co-shape-spec-list*.

2 The co-bounds of an unallocated allocatable co-array are undefined. No part of such a co-array shall be
3 referenced or defined; however, the co-array may appear as an argument to an intrinsic inquiry function
4 as specified in 13.1.

5 The co-bounds of an allocated allocatable co-array are those specified when the co-array is allocated.

6 The co-bounds of an allocatable co-array are unaffected by any subsequent redefinition or undefinition
7 of the variables on which the bounds' expressions depend.

8 5.3.8.2 Explicit-co-shape co-array

9 An **explicit-co-shape co-array** is a named co-array that has its co-rank and **co-bounds** declared by
10 an *explicit-co-shape-spec*.

11 R520 *explicit-co-shape-spec* **is** [[*lower-co-bound* :] *upper-co-bound*,]... [*lower-co-bound* :] *

13 C537 A co-array that does not have the ALLOCATABLE attribute shall have a *co-array-spec* that is
14 an *explicit-co-shape-spec*.

15 The co-rank is equal to one plus the number of *upper-co-bounds*.

16 R521 *lower-co-bound* **is** *specification-expr*

17 R522 *upper-co-bound* **is** *specification-expr*

18 C538 (R520) A *lower-co-bound* or *upper-co-bound* that is not an initialization expression shall appear
19 only in a subprogram or interface body.

20 If an explicit-co-shape co-array has co-bounds that are not initialization expressions, the co-bounds are
21 determined at entry to the procedure by evaluating the co-bounds expressions. The co-bounds of such
22 a co-array are unaffected by the redefinition or undefinition of any variable during execution of the
23 procedure.

24 The values of each *lower-co-bound* and *upper-co-bound* determine the co-bounds of the array along a
25 particular co-dimension. The subscript range of the array in that co-dimension is the set of integer values
26 between and including the lower and upper co-bounds. If the lower co-bound is omitted, the default
27 value is 1. The upper co-bound shall not be less than the lower co-bound.

28 5.3.9 INTENT attribute

29 The **INTENT attribute** specifies the intended use of a dummy argument. An INTENT (IN) dummy
30 argument is suitable for receiving data from the invoking scoping unit, an INTENT (OUT) dummy
31 argument is suitable for returning data to the invoking scoping unit, and an INTENT (INOUT) dummy
32 argument is suitable for use both to receive data from and to return data to the invoking scoping unit.

33 R523 *intent-spec* **is** IN
34 **or** OUT
35 **or** INOUT

36 C539 An entity with the INTENT attribute shall be a dummy data object or a dummy procedure
37 pointer.

38 C540 (R523) A nonpointer object with the INTENT (IN) attribute shall not appear in a variable
39 definition context (16.6.7).

1 C541 A pointer with the INTENT (IN) attribute shall not appear in a pointer association context
2 (16.6.8).

3 The INTENT (IN) attribute for a nonpointer dummy argument specifies that it shall neither be de-
4 fined nor become undefined during the execution of the procedure. The INTENT (IN) attribute for a
5 pointer dummy argument specifies that during the execution of the procedure its association shall not
6 be changed except that it may become undefined if the target is deallocated other than through the
7 pointer (16.5.2.2.3).

8 The INTENT (OUT) attribute for a nonpointer dummy argument specifies that the dummy argument
9 becomes undefined on invocation of the procedure, except for any subcomponents that are default-
10 initialized (4.5.4.5). Any actual argument that becomes associated with such a dummy argument shall
11 be definable. The INTENT (OUT) attribute for a pointer dummy argument specifies that on invocation
12 of the procedure the pointer association status of the dummy argument becomes undefined. Any actual
13 argument associated with such a pointer dummy shall be a pointer variable.

NOTE 5.13

If the actual argument is finalizable it will be finalized before undefinition and default initialization of the dummy argument (4.5.6).

14 The INTENT (INOUT) attribute for a nonpointer dummy argument specifies that any actual argument
15 associated with the dummy argument shall be definable. The INTENT (INOUT) attribute for a pointer
16 dummy argument specifies that any actual argument associated with the dummy argument shall be a
17 pointer variable.

NOTE 5.14

The INTENT attribute for an allocatable dummy argument applies to both the allocation status and the definition status. An actual argument associated with an INTENT(OUT) allocatable dummy argument is deallocated on procedure invocation (6.3.3.1).

18 If no INTENT attribute is specified for a dummy argument, its use is subject to the limitations of the
19 argument associated entity (12.5.2).

NOTE 5.15

An example of INTENT specification is:

```
SUBROUTINE MOVE (FROM, TO)
  USE PERSON_MODULE
  TYPE (PERSON), INTENT (IN) :: FROM
  TYPE (PERSON), INTENT (OUT) :: TO
```

20 If an object has an INTENT attribute, then all of its subobjects have the same INTENT attribute.

NOTE 5.16

If a dummy argument is a derived-type object with a pointer component, then the pointer as a pointer is a subobject of the dummy argument, but the target of the pointer is not. Therefore, the restrictions on subobjects of the dummy object apply to the pointer in contexts where it is used as a pointer, but not in contexts where it is dereferenced to indicate its target. For example, if X is a dummy argument of derived type with an integer pointer component P, and X has INTENT(IN), then the statement

```
X%P => NEW_TARGET
```


NOTE 5.16 (cont.)

is prohibited, but

X%P = 0

is allowed (provided that X%P is associated with a definable target).

Similarly, the INTENT restrictions on pointer dummy arguments apply only to the association of the dummy argument; they do not restrict the operations allowed on its target.

NOTE 5.17

Argument intent specifications serve several purposes in addition to documenting the intended use of dummy arguments. A processor can check whether an INTENT (IN) dummy argument is used in a way that could redefine it. A slightly more sophisticated processor could check to see whether an INTENT (OUT) dummy argument could possibly be referenced before it is defined. If the procedure's interface is explicit, the processor can also verify that actual arguments corresponding to INTENT (OUT) or INTENT (INOUT) dummy arguments are definable. A more sophisticated processor could use this information to optimize the translation of the referencing scoping unit by taking advantage of the fact that actual arguments corresponding to INTENT (IN) dummy arguments will not be changed and that any prior value of an actual argument corresponding to an INTENT (OUT) dummy argument will not be referenced and could thus be discarded.

INTENT (OUT) means that the value of the argument after invoking the procedure is entirely the result of executing that procedure. If an argument should retain its current value rather than being redefined, INTENT (INOUT) should be used rather than INTENT (OUT), even if there is no explicit reference to the value of the dummy argument. Because an INTENT(OUT) variable is considered undefined on entry to the procedure, any default initialization specified for its type will be applied.

INTENT (INOUT) is not equivalent to omitting the INTENT attribute. The actual argument corresponding to an INTENT (INOUT) dummy argument is always required to be definable, while an argument corresponding to a dummy argument without an INTENT attribute need be definable only if the dummy argument is actually redefined.

1 **5.3.10 INTRINSIC attribute**

2 The **INTRINSIC attribute** specifies that the entity is an intrinsic procedure. It may be a generic
3 procedure (13.5), a specific procedure (13.6), or both.

4 If the specific name of an intrinsic procedure (13.6) is used as an actual argument, the name shall be
5 explicitly specified to have the INTRINSIC attribute. An intrinsic procedure whose specific name is
6 marked with a bullet (●) in 13.6 shall not be used as an actual argument.

7 C542 If the name of a generic intrinsic procedure is explicitly declared to have the INTRINSIC at-
8 tribute, and it is also the generic name of one or more generic interfaces (12.4.3.2) accessible in
9 the same scoping unit, the procedures in the interfaces and the specific intrinsic procedures shall
10 all be functions or all be subroutines, and the characteristics of the specific intrinsic procedures
11 and the procedures in the interfaces shall differ as specified in 12.4.3.3.4.

12 **5.3.11 OPTIONAL attribute**

13 The **OPTIONAL attribute** specifies that the dummy argument need not be associated with an actual
14 argument in a reference to the procedure (12.5.2.13). The PRESENT intrinsic function can be used to

1 determine whether an optional dummy argument is associated with an actual argument.

2 C543 An entity with the OPTIONAL attribute shall be a dummy argument.

3 **5.3.12 PARAMETER attribute**

4 The **PARAMETER attribute** specifies that an entity is a named constant. The entity has the value
5 specified by its *initialization-expr*, converted, if necessary, to the type, type parameters and shape of the
6 entity.

7 C544 An entity with the PARAMETER attribute shall not be a variable, a co-array, or a procedure.

8 A named constant shall not be referenced unless it has been defined previously in the same statement,
9 defined in a prior statement, or made accessible by use or host association.

NOTE 5.18

Examples of declarations with a PARAMETER attribute are:

```
REAL, PARAMETER :: ONE = 1.0, Y = 4.1 / 3.0
INTEGER, DIMENSION (3), PARAMETER :: ORDER = (/ 1, 2, 3 /)
TYPE(NODE), PARAMETER :: DEFAULT = NODE(0, NULL ( ))
```

10 **5.3.13 POINTER attribute**

11 Entities with the **POINTER attribute** can be associated with different data objects or procedures
12 during execution of a program. A pointer is either a data pointer or a procedure pointer. Procedure
13 pointers are described in 12.4.3.5.

14 C545 An entity with the POINTER attribute shall not have the ALLOCATABLE, INTRINSIC, or
15 TARGET attribute.

16 C546 A procedure with the POINTER attribute shall have the EXTERNAL attribute.

17 C547 A co-array shall not have the POINTER attribute.

18 A data pointer shall not be referenced unless it is pointer associated with a target object that is defined.

19 A data pointer shall not be defined unless it is pointer associated with a target object that is definable.

20 If a data pointer is associated, the values of its deferred type parameters are the same as the values of
21 the corresponding type parameters of its target.

22 A procedure pointer shall not be referenced unless it is pointer associated with a target procedure.

NOTE 5.19

Examples of POINTER attribute specifications are:

```
TYPE (NODE), POINTER :: CURRENT, TAIL
REAL, DIMENSION (:, :), POINTER :: IN, OUT, SWAP
```

For a more elaborate example see C.2.1.

23 **5.3.14 PROTECTED attribute**

24 The **PROTECTED attribute** imposes limitations on the usage of module entities.

- 1 C548 The PROTECTED attribute shall be specified only in the specification part of a module.
- 2 C549 An entity with the PROTECTED attribute shall be a procedure pointer or variable.
- 3 C550 An entity with the PROTECTED attribute shall not be in a common block.
- 4 C551 A nonpointer object that has the PROTECTED attribute and is accessed by use association
5 shall not appear in a variable definition context (16.6.7) or as the *data-target* or *proc-target* in
6 a *pointer-assignment-stmt*.
- 7 C552 A pointer that has the PROTECTED attribute and is accessed by use association shall not
8 appear in a pointer association context (16.6.8).
- 9 Other than within the module in which an entity is given the PROTECTED attribute, or within any of
10 its descendants,
- 11 (1) if it is a nonpointer object, it is not definable, and
12 (2) if it is a pointer, its association status shall not be changed except that it may become
13 undefined if its target is deallocated other than through the pointer (16.5.2.2.3) or if its
14 target becomes undefined by execution of a RETURN or END statement.
- 15 If an object has the PROTECTED attribute, all of its subobjects have the PROTECTED attribute.

NOTE 5.20

An example of the PROTECTED attribute:

```

MODULE temperature
  REAL, PROTECTED :: temp_c, temp_f
CONTAINS
  SUBROUTINE set_temperature_c(c)
    REAL, INTENT(IN) :: c
    temp_c = c
    temp_f = temp_c*(9.0/5.0) + 32
  END SUBROUTINE
END MODULE

```

The PROTECTED attribute ensures that the variables `temp_c` and `temp_f` cannot be modified other than via the `set_temperature_c` procedure, thus keeping them consistent with each other.

16 5.3.15 SAVE attribute

- 17 The SAVE attribute specifies that a variable retains its association status, allocation status, definition
18 status, and value after execution of a RETURN or END statement unless it is a pointer and its target
19 becomes undefined (16.5.2.2.3(4)). If it is a local variable of a subprogram it is shared by all instances
20 (12.6.2.3) of the subprogram.
- 21 Giving a common block the SAVE attribute confers the attribute on all variables in the common block.
- 22 C553 An entity with the SAVE attribute shall be a common block, variable, or procedure pointer.
- 23 C554 The SAVE attribute shall not be specified for a dummy argument, a function result, an automatic
24 data object, or an object that is in a common block.
- 25 A **saved** entity is an entity that has the SAVE attribute. An **unsaved** entity is an entity that does not
26 have the SAVE attribute.
- 27 The SAVE attribute has no effect on entities declared in a main program. If a common block has the

- 1 SAVE attribute in any scoping unit that is not a main program, it shall have the SAVE attribute in
2 every scoping unit that is not a main program.

3 **5.3.16 TARGET attribute**

- 4 The **TARGET attribute** specifies that a data object may have a pointer associated with it (7.4.2).
5 An object without the TARGET attribute shall not have an accessible pointer associated with it.
6 C555 An entity with the TARGET attribute shall be a variable.
7 C556 An entity with the TARGET attribute shall not have the POINTER attribute.

NOTE 5.21

In addition to variables explicitly declared to have the TARGET attribute, the objects created by allocation of pointers (6.3.1.2) have the TARGET attribute.

- 8 If an object has the TARGET attribute, then all of its nonpointer subobjects also have the TARGET
9 attribute.

NOTE 5.22

Examples of TARGET attribute specifications are:

```
TYPE (NODE), TARGET :: HEAD
REAL, DIMENSION (1000, 1000), TARGET :: A, B
```

For a more elaborate example see C.2.2.

NOTE 5.23

Every object designator that starts from a target object will have either the TARGET or POINTER attribute. If pointers are involved, the designator might not necessarily be a subobject of the original target object, but because pointers may point only to targets, there is no way to end up at a nonpointer that is not a target.

10 **5.3.17 VALUE attribute**

- 11 The **VALUE attribute** specifies a type of argument association (12.5.2.5) for a dummy argument.
12 C557 An entity with the VALUE attribute shall be a scalar dummy data object.
13 C558 An entity with the VALUE attribute shall not have the ALLOCATABLE, INTENT(INOUT),
14 INTENT(OUT), POINTER, or VOLATILE attributes.
15 C559 If an entity has the VALUE attribute, any length type parameter value in its declaration shall
16 be omitted or specified by an initialization expression.

17 **5.3.18 VOLATILE attribute**

- 18 The **VOLATILE attribute** specifies that an object may be referenced, defined, or become undefined,
19 by means not specified by the program, or by another image without synchronization.
20 C560 An entity with the VOLATILE attribute shall be a variable that is not an INTENT(IN) dummy
21 argument.
22 An object may have the VOLATILE attribute in a particular scoping unit without necessarily having

- 1 it in other scoping units (11.2.2, 16.5.1.4). If an object has the VOLATILE attribute, then all of its
 2 subobjects also have the VOLATILE attribute.

NOTE 5.24

The Fortran processor should use the most recent definition of a volatile object when a value is required. Likewise, it should make the most recent Fortran definition available. It is the programmer's responsibility to manage any interaction with non-Fortran processes.

- 3 A pointer with the VOLATILE attribute may additionally have its association status, dynamic type and
 4 type parameters, and array bounds changed by means not specified by the program.

NOTE 5.25

If the target of a pointer is referenced, defined, or becomes undefined, by means not specified by the program, while the pointer is associated with the target, then the pointer shall have the VOLATILE attribute. Usually a pointer should have the VOLATILE attribute if its target has the VOLATILE attribute. Similarly, all members of an EQUIVALENCE group should have the VOLATILE attribute if one member has the VOLATILE attribute.

- 5 An allocatable object with the VOLATILE attribute may additionally have its allocation status, dynamic
 6 type and type parameters, and array bounds changed by means not specified by the program.

7 **5.4 Attribute specification statements**

8 **5.4.1 Accessibility statement**

9 R524 *access-stmt* **is** *access-spec* [[*::*] *access-id-list*]
 10 R525 *access-id* **is** *use-name*
 11 **or** *generic-spec*

12 C561 (R524) An *access-stmt* shall appear only in the *specification-part* of a module. Only one ac-
 13 cessibility statement with an omitted *access-id-list* is permitted in the *specification-part* of a
 14 module.

15 C562 (R525) Each *use-name* shall be the name of a named variable, procedure, derived type, named
 16 constant, namelist group, or macro.

17 An *access-stmt* with an *access-id-list* specifies the accessibility attribute (5.3.2), PUBLIC or PRIVATE,
 18 of each *access-id* in the list. An *access-stmt* without an *access-id* list specifies the default accessibility
 19 that applies to all potentially accessible identifiers in the *specification-part* of the module. The statement

20 PUBLIC

21 specifies a default of public accessibility. The statement

22 PRIVATE

23 specifies a default of private accessibility. If no such statement appears in a module, the default is public
 24 accessibility.

NOTE 5.26

Examples of accessibility statements are:

```
MODULE EX
PRIVATE
```

NOTE 5.26 (cont.)

PUBLIC :: A, B, C, ASSIGNMENT (=), OPERATOR (+)

1 5.4.2 ALLOCATABLE statement

2 R526 *allocatable-stmt* **is** ALLOCATABLE [::] *allocatable-decl-list*
3 R527 *allocatable-decl* **is** *object-name* [(*array-spec*)] ■
4 ■ [*lbracket co-array-spec rbracket*]

5 This statement specifies the ALLOCATABLE attribute (5.3.3) for a list of objects.

NOTE 5.27

An example of an ALLOCATABLE statement is:
--

REAL A, B (:), SCALAR ALLOCATABLE :: A (:, :), B, SCALAR

6 5.4.3 ASYNCHRONOUS statement

7 R528 *asynchronous-stmt* **is** ASYNCHRONOUS [::] *object-name-list*

8 The ASYNCHRONOUS statement specifies the ASYNCHRONOUS attribute (5.3.4) for a list of objects.

9 5.4.4 BIND statement

10 R529 *bind-stmt* **is** *language-binding-spec* [::] *bind-entity-list*
11 R530 *bind-entity* **is** *entity-name*
12 **or** / *common-block-name* /

13 C563 (R529) If the *language-binding-spec* has a NAME= specifier, the *bind-entity-list* shall consist of
14 a single *bind-entity*.

15 The BIND statement specifies the BIND attribute (5.3.5) for a list of variables and common blocks.

16 5.4.5 CONTIGUOUS statement

17 R531 *contiguous-stmt* **is** CONTIGUOUS [::] *object-name-list*

18 The CONTIGUOUS statement specifies the CONTIGUOUS attribute (5.3.6) for a list of objects.

19 5.4.6 DATA statement

20 R532 *data-stmt* **is** DATA *data-stmt-set* [[,] *data-stmt-set*] ...

21 This statement specifies explicit initialization (5.2.3).

22 A variable, or part of a variable, shall not be explicitly initialized more than once in a program. If a
23 nonpointer object has been specified with default initialization in a type definition, it shall not appear
24 in a *data-stmt-object-list*.

25 A variable that appears in a DATA statement and has not been typed previously may appear in a
26 subsequent type declaration only if that declaration confirms the implicit typing. An array name,
27 array section, or array element that appears in a DATA statement shall have had its array properties
28 established by a previous specification statement.

1 Except for variables in named common blocks, a named variable has the SAVE attribute if any part of
 2 it is initialized in a DATA statement, and this may be confirmed by explicit specification.

3	R533	<i>data-stmt-set</i>	is	<i>data-stmt-object-list</i> / <i>data-stmt-value-list</i> /
4	R534	<i>data-stmt-object</i>	is	<i>variable</i>
5			or	<i>data-implied-do</i>
6	R535	<i>data-implied-do</i>	is	(<i>data-i-do-object-list</i> , <i>data-i-do-variable</i> = ■
7				■ <i>scalar-int-initialization-expr</i> , ■
8				■ <i>scalar-int-initialization-expr</i> ■
9				■ [, <i>scalar-int-initialization-expr</i>])
10	R536	<i>data-i-do-object</i>	is	<i>array-element</i>
11			or	<i>scalar-structure-component</i>
12			or	<i>data-implied-do</i>
13	R537	<i>data-i-do-variable</i>	is	<i>do-variable</i>

14 C564 A *data-stmt-object* or *data-i-do-object* shall not be a co-indexed variable.

15 C565 (R534) In a *variable* that is a *data-stmt-object*, any subscript, section subscript, substring start-
 16 ing point, and substring ending point shall be an initialization expression.

17 C566 (R534) A variable whose designator appears as a *data-stmt-object* or a *data-i-do-object* shall
 18 not be a dummy argument, made accessible by use association or host association, in a named
 19 common block unless the DATA statement is in a block data program unit, in a blank common
 20 block, a function name, a function result name, an automatic object, or an allocatable variable.

21 C567 (R534) A *data-i-do-object* or a *variable* that appears as a *data-stmt-object* shall not be an object
 22 designator in which a pointer appears other than as the entire rightmost *part-ref*.

23 C568 (R536) The *array-element* shall be a variable.

24 C569 (R536) The *scalar-structure-component* shall be a variable.

25 C570 (R536) The *scalar-structure-component* shall contain at least one *part-ref* that contains a *sub-*
 26 *script-list*.

27 C571 (R536) In an *array-element* or *scalar-structure-component* that is a *data-i-do-object*, any sub-
 28 script shall be an initialization expression, and any primary within that subscript that is a
 29 *data-i-do-variable* shall be a DO variable of this *data-implied-do* or of a containing *data-implied-*
 30 *do*.

31	R538	<i>data-stmt-value</i>	is	[<i>data-stmt-repeat</i> *] <i>data-stmt-constant</i>
----	------	------------------------	-----------	---

32	R539	<i>data-stmt-repeat</i>	is	<i>scalar-int-constant</i>
----	------	-------------------------	-----------	----------------------------

33			or	<i>scalar-int-constant-subobject</i>
----	--	--	-----------	--------------------------------------

34 C572 (R539) The *data-stmt-repeat* shall be positive or zero. If the *data-stmt-repeat* is a named con-
 35 stant, it shall have been declared previously in the scoping unit or made accessible by use
 36 association or host association.

37	R540	<i>data-stmt-constant</i>	is	<i>scalar-constant</i>
----	------	---------------------------	-----------	------------------------

38			or	<i>scalar-constant-subobject</i>
----	--	--	-----------	----------------------------------

39			or	<i>signed-int-literal-constant</i>
----	--	--	-----------	------------------------------------

40			or	<i>signed-real-literal-constant</i>
----	--	--	-----------	-------------------------------------

41			or	<i>null-init</i>
----	--	--	-----------	------------------

42			or	<i>initial-data-target</i>
----	--	--	-----------	----------------------------

43			or	<i>structure-constructor</i>
----	--	--	-----------	------------------------------

44 C573 (R540) If a DATA statement constant value is a named constant or a structure constructor, the

- 1 named constant or derived type shall have been declared previously in the scoping unit or made
2 accessible by use or host association.
- 3 C574 (R540) If a *data-stmt-constant* is a *structure-constructor*, it shall be an initialization expression.
- 4 R541 *int-constant-subobject* **is** *constant-subobject*
- 5 C575 (R541) *int-constant-subobject* shall be of type integer.
- 6 R542 *constant-subobject* **is** *designator*
- 7 C576 (R542) *constant-subobject* shall be a subobject of a constant.
- 8 C577 (R542) Any subscript, substring starting point, or substring ending point shall be an initializa-
9 tion expression.
- 10 The *data-stmt-object-list* is expanded to form a sequence of pointers and scalar variables, referred to
11 as “sequence of variables” in subsequent text. A nonpointer array whose unqualified name appears
12 as a *data-stmt-object* or *data-i-do-object* is equivalent to a complete sequence of its array elements in
13 array element order (6.2.2.2). An array section is equivalent to the sequence of its array elements in
14 array element order. A *data-implied-do* is expanded to form a sequence of array elements and structure
15 components, under the control of the *data-i-do-variable*, as in the DO construct (8.1.7.5).
- 16 The *data-stmt-value-list* is expanded to form a sequence of *data-stmt-constants*. A *data-stmt-repeat*
17 indicates the number of times the following *data-stmt-constant* is to be included in the sequence; omission
18 of a *data-stmt-repeat* has the effect of a repeat factor of 1.
- 19 A zero-sized array or a *data-implied-do* with an iteration count of zero contributes no variables to the
20 expanded sequence of variables, but a zero-length scalar character variable does contribute a variable
21 to the expanded sequence. A *data-stmt-constant* with a repeat factor of zero contributes no *data-stmt-*
22 *constants* to the expanded sequence of scalar *data-stmt-constants*.
- 23 The expanded sequences of variables and *data-stmt-constants* are in one-to-one correspondence. Each
24 *data-stmt-constant* specifies the initial value, initial data target, or *null-init* for the corresponding vari-
25 able. The lengths of the two expanded sequences shall be the same.
- 26 A *data-stmt-constant* shall be *null-init* or *initial-data-target* if and only if the corresponding *data-stmt-*
27 *object* has the POINTER attribute. If *data-stmt-constant* is *null-init*, the initial association status of
28 the corresponding data statement object is disassociated. If *data-stmt-constant* is *initial-data-target* the
29 corresponding data statement object shall be data-pointer-initialization compatible with the initial data
30 target; the data statement object is initially associated with the target.
- 31 A *data-stmt-constant* other than *null-init* or *initial-data-target* shall be compatible with its corresponding
32 variable according to the rules of intrinsic assignment (7.4.1.2). The variable is initially defined with
33 the value specified by the *data-stmt-constant*; if necessary, the value is converted according to the rules
34 of intrinsic assignment (7.4.1.3) to a value that agrees in type, type parameters, and shape with the
35 variable.

NOTE 5.28

Examples of DATA statements are:

```

CHARACTER (LEN = 10) NAME
INTEGER, DIMENSION (0:9) :: MILES
REAL, DIMENSION (100, 100) :: SKEW
TYPE (NODE), POINTER :: HEAD_OF_LIST
TYPE (PERSON) MYNAME, YOURNAME
DATA NAME / 'JOHN DOE' /, MILES / 10 * 0 /

```


NOTE 5.28 (cont.)

```
DATA ((SKEW (K, J), J = 1, K), K = 1, 100) / 5050 * 0.0 /
DATA ((SKEW (K, J), J = K + 1, 100), K = 1, 99) / 4950 * 1.0 /
DATA HEAD_OF_LIST / NULL() /
DATA MYNAME / PERSON (21, 'JOHN SMITH') /
DATA YOURNAME % AGE, YOURNAME % NAME / 35, 'FRED BROWN' /
```

The character variable NAME is initialized with the value JOHN DOE with padding on the right because the length of the constant is less than the length of the variable. All ten elements of the integer array MILES are initialized to zero. The two-dimensional array SKEW is initialized so that the lower triangle of SKEW is zero and the strict upper triangle is one. The structures MYNAME and YOURNAME are declared using the derived type PERSON from Note 4.21. The pointer HEAD_OF_LIST is declared using the derived type NODE from Note 4.40; it is initially disassociated. MYNAME is initialized by a structure constructor. YOURNAME is initialized by supplying a separate value for each component.

1 5.4.7 DIMENSION statement

- 2 R543 *dimension-stmt* **is** DIMENSION [::] *dimension-decl-list*
3 R544 *dimension-decl* **is** *array-name* (*array-spec*)
4 **or** *co-name* [(*array-spec*)] *lbracket co-array-spec rbracket*

- 5 This statement specifies the DIMENSION attribute (5.3.7) for a list of objects.

NOTE 5.29

An example of a DIMENSION statement is:

```
DIMENSION A (10), B (10, 70), C (:)
```

6 5.4.8 INTENT statement

- 7 R545 *intent-stmt* **is** INTENT (*intent-spec*) [::] *dummy-arg-name-list*

- 8 This statement specifies the INTENT attribute (5.3.9) for the dummy arguments in the list.

NOTE 5.30

An example of an INTENT statement is:

```
SUBROUTINE EX (A, B)
  INTENT (INOUT) :: A, B
```

9 5.4.9 OPTIONAL statement

- 10 R546 *optional-stmt* **is** OPTIONAL [::] *dummy-arg-name-list*

- 11 This statement specifies the OPTIONAL attribute (5.3.11) for the dummy arguments in the list.

NOTE 5.31

An example of an OPTIONAL statement is:

```
SUBROUTINE EX (A, B)
  OPTIONAL :: B
```

1 5.4.10 PARAMETER statement

2 The **PARAMETER statement** specifies the PARAMETER attribute (5.3.12) and the values for the
3 named constants in the list.

4 R547 *parameter-stmt* **is** PARAMETER (*named-constant-def-list*)
5 R548 *named-constant-def* **is** *named-constant* = *initialization-expr*

6 If a named constant is defined by a PARAMETER statement, it shall not be subsequently declared to
7 have a type or type parameter value that differs from the type and type parameters it would have if
8 declared implicitly (5.5). A named array constant defined by a PARAMETER statement shall have its
9 shape specified in a prior specification statement.

10 The value of each named constant is that specified by the corresponding initialization expression; if
11 necessary, the value is converted according to the rules of intrinsic assignment (7.4.1.3) to a value that
12 agrees in type, type parameters, and shape with the named constant.

NOTE 5.32

An example of a PARAMETER statement is:

```
PARAMETER (MODULUS = MOD (28, 3), NUMBER_OF_SENATORS = 100)
```

13 5.4.11 POINTER statement

14 R549 *pointer-stmt* **is** POINTER [::] *pointer-decl-list*
15 R550 *pointer-decl* **is** *object-name* [(*deferred-shape-spec-list*)]
16 **or** *proc-entity-name*

17 This statement specifies the POINTER attribute (5.3.13) for a list of entities.

NOTE 5.33

An example of a POINTER statement is:

```
TYPE (NODE) :: CURRENT  
POINTER :: CURRENT, A (:, :)
```

18 5.4.12 PROTECTED statement

19 R551 *protected-stmt* **is** PROTECTED [::] *entity-name-list*

20 The **PROTECTED statement** specifies the PROTECTED attribute (5.3.14) for a list of entities.

21 5.4.13 SAVE statement

22 R552 *save-stmt* **is** SAVE [[::] *saved-entity-list*]
23 R553 *saved-entity* **is** *object-name*
24 **or** *proc-pointer-name*
25 **or** / *common-block-name* /
26 R554 *proc-pointer-name* **is** *name*

27 C578 (R552) If a SAVE statement with an omitted saved entity list appears in a scoping unit, no
28 other appearance of the SAVE *attr-spec* or SAVE statement is permitted in that scoping unit.

29 A SAVE statement with a saved entity list specifies the SAVE attribute (5.3.15) for a list of entities. A

- 1 SAVE statement without a saved entity list is treated as though it contained the names of all allowed
2 items in the same scoping unit.

NOTE 5.34

An example of a SAVE statement is:

```
SAVE A, B, C, / BLOCKA /, D
```

3 **5.4.14 TARGET statement**

- 4 R555 *target-stmt* is TARGET [::] *target-decl-list*
5 R556 *target-decl* is *object-name* [(*array-spec*)] ■
6 ■ [*lbracket co-array-spec rbracket*]

- 7 This statement specifies the TARGET attribute (5.3.16) for a list of objects.

NOTE 5.35

An example of a TARGET statement is:

```
TARGET :: A (1000, 1000), B
```

8 **5.4.15 VALUE statement**

- 9 R557 *value-stmt* is VALUE [::] *dummy-arg-name-list*

- 10 The VALUE statement specifies the VALUE attribute (5.3.17) for a list of dummy arguments.

11 **5.4.16 VOLATILE statement**

- 12 R558 *volatile-stmt* is VOLATILE [::] *object-name-list*

- 13 The VOLATILE statement specifies the VOLATILE attribute (5.3.18) for a list of objects.

14 **5.5 IMPLICIT statement**

- 15 In a scoping unit, an **IMPLICIT statement** specifies a type, and possibly type parameters, for all
16 implicitly typed data entities whose names begin with one of the letters specified in the statement.
17 Alternatively, it may indicate that no implicit typing rules are to apply in a particular scoping unit.

- 18 R559 *implicit-stmt* is IMPLICIT *implicit-spec-list*
19 or IMPLICIT NONE
20 R560 *implicit-spec* is *declaration-type-spec* (*letter-spec-list*)
21 R561 *letter-spec* is *letter* [– *letter*]

- 22 C579 (R559) If IMPLICIT NONE is specified in a scoping unit, it shall precede any PARAMETER
23 statements that appear in the scoping unit and there shall be no other IMPLICIT statements
24 in the scoping unit.

- 25 C580 (R561) If the minus and second *letter* appear, the second letter shall follow the first letter
26 alphabetically.

- 27 A *letter-spec* consisting of two *letters* separated by a minus is equivalent to writing a list containing all
28 of the letters in alphabetical order in the alphabetic sequence from the first letter through the second

- 1 letter. For example, A–C is equivalent to A, B, C. The same letter shall not appear as a single letter, or
 2 be included in a range of letters, more than once in all of the IMPLICIT statements in a scoping unit.
- 3 In each scoping unit, there is a mapping, which may be null, between each of the letters A, B, ..., Z
 4 and a type (and type parameters). An IMPLICIT statement specifies the mapping for the letters in
 5 its *letter-spec-list*. IMPLICIT NONE specifies the null mapping for all the letters. If a mapping is not
 6 specified for a letter, the default for a program unit or an interface body is default integer if the letter
 7 is I, J, ..., or N and default real otherwise, and the default for an internal or module procedure is the
 8 mapping in the host scoping unit.
- 9 Any data entity that is not explicitly declared by a type declaration statement, is not an intrinsic
 10 function, and is not made accessible by use association or host association is declared implicitly to be of
 11 the type (and type parameters) mapped from the first letter of its name, provided the mapping is not
 12 null. The mapping for the first letter of the data entity shall either have been established by a prior
 13 IMPLICIT statement or be the default mapping for the letter. The mapping may be to a derived type
 14 that is inaccessible in the local scope if the derived type is accessible in the host scoping unit. The data
 15 entity is treated as if it were declared in an explicit type declaration in the outermost scoping unit in
 16 which it appears. An explicit type specification in a FUNCTION statement overrides an IMPLICIT
 17 statement for the name of the result variable of that function subprogram.

NOTE 5.36

The following are examples of the use of IMPLICIT statements:

```

MODULE EXAMPLE_MODULE
  IMPLICIT NONE
  ...
  INTERFACE
    FUNCTION FUN (I)      ! Not all data entities need to
      INTEGER FUN        ! be declared explicitly
    END FUNCTION FUN
  END INTERFACE
CONTAINS
  FUNCTION JFUN (J)      ! All data entities need to
    INTEGER JFUN, J     ! be declared explicitly.
    ...
  END FUNCTION JFUN
END MODULE EXAMPLE_MODULE
SUBROUTINE SUB
  IMPLICIT COMPLEX (C)
  C = (3.0, 2.0)        ! C is implicitly declared COMPLEX
  ...
CONTAINS
  SUBROUTINE SUB1
    IMPLICIT INTEGER (A, C)
    C = (0.0, 0.0)      ! C is host associated and of
                       ! type complex
    Z = 1.0             ! Z is implicitly declared REAL
    A = 2               ! A is implicitly declared INTEGER
    CC = 1              ! CC is implicitly declared INTEGER
    ...
  END SUBROUTINE SUB1
  SUBROUTINE SUB2
    Z = 2.0             ! Z is implicitly declared REAL and
                       ! is different from the variable of

```

NOTE 5.36 (cont.)

```

                                ! the same name in SUB1
...
END SUBROUTINE SUB2
SUBROUTINE SUB3
  USE EXAMPLE_MODULE ! Accesses integer function FUN
                    ! by use association
  Q = FUN (K)        ! Q is implicitly declared REAL and
                    ! K is implicitly declared INTEGER
...
END SUBROUTINE SUB3
END SUBROUTINE SUB

```

NOTE 5.37

The following is an example of a mapping to a derived type that is inaccessible in the local scope:

```

PROGRAM MAIN
  IMPLICIT TYPE(BLOB) (A)
  TYPE BLOB
    INTEGER :: I
  END TYPE BLOB
  TYPE(BLOB) :: B
  CALL STEVE
CONTAINS
  SUBROUTINE STEVE
    INTEGER :: BLOB
    ..
    AA = B
    ..
  END SUBROUTINE STEVE
END PROGRAM MAIN

```

In the subroutine STEVE, it is not possible to explicitly declare a variable to be of type BLOB because BLOB has been given a different meaning, but implicit mapping for the letter A still maps to type BLOB, so AA is of type BLOB.

1 5.6 NAMELIST statement

2 A **NAMELIST statement** specifies a group of named data objects, which may be referred to by a
 3 single name for the purpose of data transfer (9.5, 10.11).

4 R562 *namelist-stmt* **is** **NAMELIST** ■
 5 ■ / *namelist-group-name* / *namelist-group-object-list* ■
 6 ■ [[,] / *namelist-group-name* / ■
 7 ■ *namelist-group-object-list*] ...

8 C581 (R562) The *namelist-group-name* shall not be a name accessed by use association.

9 R563 *namelist-group-object* **is** *variable-name*

10 C582 (R563) A *namelist-group-object* shall not be an assumed-size array.

11 C583 (R562) A *namelist-group-object* shall not have the PRIVATE attribute if the *namelist-group-*
 12 *name* has the PUBLIC attribute.

- 1 The order in which the variables are specified in the NAMELIST statement determines the order in
 2 which the values appear on output.
- 3 Any *namelist-group-name* may occur more than once in the NAMELIST statements in a scoping unit.
 4 The *namelist-group-object-list* following each successive appearance of the same *namelist-group-name* in
 5 a scoping unit is treated as a continuation of the list for that *namelist-group-name*.
- 6 A namelist group object may be a member of more than one namelist group.
- 7 A namelist group object shall either be accessed by use or host association or shall have its type, type
 8 parameters, and shape specified by previous specification statements or the procedure heading in the
 9 same scoping unit or by the implicit typing rules in effect for the scoping unit. If a namelist group object
 10 is typed by the implicit typing rules, its appearance in any subsequent type declaration statement shall
 11 confirm the implied type and type parameters.

NOTE 5.38

An example of a NAMELIST statement is:

```
NAMELIST /NLIST/ A, B, C
```

12 5.7 Storage association of data objects

13 5.7.1 EQUIVALENCE statement

14 5.7.1.1 General

- 15 An **EQUIVALENCE statement** is used to specify the sharing of storage units by two or more objects
 16 in a scoping unit. This causes storage association (16.5.3) of the objects that share the storage units.

NOTE 5.39

The co-size of a co-array is not a constant, therefore co-arrays are not allowed in COMMON or EQUIVALENCE.

- 17 If the equivalenced objects have differing type or type parameters, the EQUIVALENCE statement does
 18 not cause type conversion or imply mathematical equivalence. If a scalar and an array are equivalenced,
 19 the scalar does not have array properties and the array does not have the properties of a scalar.

- | | | | | |
|----|------|---------------------------|-----------|--|
| 20 | R564 | <i>equivalence-stmt</i> | is | EQUIVALENCE <i>equivalence-set-list</i> |
| 21 | R565 | <i>equivalence-set</i> | is | (<i>equivalence-object</i> , <i>equivalence-object-list</i>) |
| 22 | R566 | <i>equivalence-object</i> | is | <i>variable-name</i> |
| 23 | | | or | <i>array-element</i> |
| 24 | | | or | <i>substring</i> |

- 25 C584 (R566) An *equivalence-object* shall not be a designator with a base object that is a dummy
 26 argument, a pointer, an allocatable variable, a derived-type object that has an allocatable ulti-
 27 mate component, an object of a nonsequence derived type, an object of a derived type that has
 28 a pointer at any level of component selection, an automatic object, a function name, an entry
 29 name, a result name, a variable with the BIND attribute, a variable in a common block that
 30 has the BIND attribute, or a named constant.

- 31 C585 (R566) An *equivalence-object* shall not be a designator that has more than one *part-ref*.

- 32 C586 (R566) An *equivalence-object* shall not be a co-array or a subobject thereof.

- 1 C587 (R566) An *equivalence-object* shall not have the TARGET attribute.
- 2 C588 (R566) Each subscript or substring range expression in an *equivalence-object* shall be an integer
3 initialization expression (7.1.7).
- 4 C589 (R565) If an *equivalence-object* is of type default integer, default real, double precision real,
5 default complex, default logical, default bits, or numeric sequence type, all of the objects in the
6 equivalence set shall be of these types.
- 7 C590 (R565) If an *equivalence-object* is of type default character or character sequence type, all of the
8 objects in the equivalence set shall be of these types.
- 9 C591 (R565) If an *equivalence-object* is of a sequence derived type that is not a numeric sequence or
10 character sequence type, all of the objects in the equivalence set shall be of the same type with
11 the same type parameter values.
- 12 C592 (R565) If an *equivalence-object* is of an intrinsic type other than default integer, default real,
13 double precision real, default complex, default logical, or default character, all of the objects in
14 the equivalence set shall be of the same type with the same kind type parameter value.
- 15 C593 (R566) If an *equivalence-object* has the PROTECTED attribute, all of the objects in the equiv-
16 alence set shall have the PROTECTED attribute.
- 17 C594 (R566) The name of an *equivalence-object* shall not be a name made accessible by use association.
- 18 C595 (R566) A *substring* shall not have length zero.

NOTE 5.40

The EQUIVALENCE statement allows the equivalencing of sequence structures and the equivalencing of objects of intrinsic type with nondefault type parameters, but there are strict rules regarding the appearance of these objects in an EQUIVALENCE statement.

A structure that appears in an EQUIVALENCE statement shall be a sequence structure. If a sequence structure is not of numeric sequence type or of character sequence type, it shall be equivalenced only to objects of the same type with the same type parameter values.

A structure of a numeric sequence type shall be equivalenced only to another structure of a numeric sequence type, an object of default integer type, default real type, double precision real type, default complex type, default logical type, or default bits such that components of the structure ultimately become associated only with objects of these types.

A structure of a character sequence type shall be equivalenced only to an object of default character type or another structure of a character sequence type.

An object of intrinsic type with nondefault kind type parameters shall not be equivalenced to objects of different type or kind type parameters.

Further rules on the interaction of EQUIVALENCE statements and default initialization are given in 16.5.3.4.

19 **5.7.1.2 Equivalence association**

- 20 An EQUIVALENCE statement specifies that the storage sequences (16.5.3.2) of the data objects specified
21 in an *equivalence-set* are storage associated. All of the nonzero-sized sequences in the *equivalence-set*, if
22 any, have the same first storage unit, and all of the zero-sized sequences in the *equivalence-set*, if any,
23 are storage associated with one another and with the first storage unit of any nonzero-sized sequences.
24 This causes the storage association of the data objects in the *equivalence-set* and may cause storage

1 association of other data objects.

2 5.7.1.3 Equivalence of default character objects

3 A data object of type default character shall not be equivalenced to an object that is not of type default
4 character and not of a character sequence type. The lengths of the equivalenced character objects need
5 not be the same.

6 An EQUIVALENCE statement specifies that the storage sequences of all the default character data
7 objects specified in an *equivalence-set* are storage associated. All of the nonzero-sized sequences in the
8 *equivalence-set*, if any, have the same first character storage unit, and all of the zero-sized sequences in
9 the *equivalence-set*, if any, are storage associated with one another and with the first character storage
10 unit of any nonzero-sized sequences. This causes the storage association of the data objects in the
11 *equivalence-set* and may cause storage association of other data objects.

NOTE 5.41

For example, using the declarations:

```
CHARACTER (LEN = 4) :: A, B
CHARACTER (LEN = 3) :: C (2)
EQUIVALENCE (A, C (1)), (B, C (2))
```

the association of A, B, and C can be illustrated graphically as:

1	2	3	4	5	6	7
---	--- A	---	---			
			---	--- B	---	---
---	C(1)	---	---	C(2)	---	

12 5.7.1.4 Array names and array element designators

13 For a nonzero-sized array, the use of the array name unqualified by a subscript list as an *equivalence-*
14 *object* has the same effect as using an array element designator that identifies the first element of the
15 array.

16 5.7.1.5 Restrictions on EQUIVALENCE statements

17 An EQUIVALENCE statement shall not specify that the same storage unit is to occur more than once
18 in a storage sequence.

NOTE 5.42

For example:

```
REAL, DIMENSION (2) :: A
REAL :: B
EQUIVALENCE (A (1), B), (A (2), B) ! Not standard-conforming
```

is prohibited, because it would specify the same storage unit for A (1) and A (2).

19 An EQUIVALENCE statement shall not specify that consecutive storage units are to be nonconsecutive.

NOTE 5.43

For example, the following is prohibited:

```
REAL A (2)
DOUBLE PRECISION D (2)
EQUIVALENCE (A (1), D (1)), (A (2), D (2)) ! Not standard-conforming
```

1 5.7.2 COMMON statement

2 5.7.2.1 General

3 The **COMMON statement** specifies blocks of physical storage, called **common blocks**, that can be
4 accessed by any of the scoping units in a program. Thus, the COMMON statement provides a global
5 data facility based on storage association (16.5.3).

6 The common blocks specified by the COMMON statement may be named and are called **named com-**
7 **mon blocks**, or may be unnamed and are called **blank common**.

```
8 R567  common-stmt           is  COMMON ■
9                                     ■ [ / [ common-block-name ] / ] common-block-object-list ■
10                                    ■ [ [ , ] / [ common-block-name ] / ■
11                                    ■ common-block-object-list ] ...
12 R568  common-block-object   is  variable-name [ ( array-spec ) ]
13                                     or  proc-pointer-name
```

14 C596 (R568) An *array-spec* in a *common-block-object* shall be an *explicit-shape-spec-list*.

15 C597 (R568) Only one appearance of a given *variable-name* or *proc-pointer-name* is permitted in all
16 *common-block-object-lists* within a scoping unit.

17 C598 (R568) A *common-block-object* shall not be a dummy argument, an allocatable variable, a
18 derived-type object with an ultimate component that is allocatable, an automatic object, a
19 function name, an entry name, a variable with the BIND attribute, a co-array, or a result name.

20 C599 (R568) If a *common-block-object* is of a derived type, it shall be a sequence type (4.5.2.3) or a
21 type with the BIND attribute and it shall have no default initialization.

22 C5100 (R568) A *variable-name* or *proc-pointer-name* shall not be a name made accessible by use
23 association.

24 In each COMMON statement, the data objects whose names appear in a common block object list
25 following a common block name are declared to be in that common block. If the first common block
26 name is omitted, all data objects whose names appear in the first common block object list are specified to
27 be in blank common. Alternatively, the appearance of two slashes with no common block name between
28 them declares the data objects whose names appear in the common block object list that follows to be
29 in blank common.

30 Any common block name or an omitted common block name for blank common may occur more than
31 once in one or more COMMON statements in a scoping unit. The common block list following each
32 successive appearance of the same common block name in a scoping unit is treated as a continuation of
33 the list for that common block name. Similarly, each blank common block object list in a scoping unit
34 is treated as a continuation of blank common.

35 The form *variable-name* (*array-spec*) specifies the DIMENSION attribute for that variable.

36 If derived-type objects of numeric sequence type (4.5.2) or character sequence type (4.5.2) appear in

1 common, it is as if the individual components were enumerated directly in the common list.

NOTE 5.44

Examples of COMMON statements are:

```
COMMON /BLOCKA/ A, B, D (10, 30)
```

```
COMMON I, J, K
```

2 5.7.2.2 Common block storage sequence

3 For each common block in a scoping unit, a **common block storage sequence** is formed as follows:

- 4 (1) A storage sequence is formed consisting of the sequence of storage units in the storage
5 sequences (16.5.3.2) of all data objects in the common block object lists for the common
6 block. The order of the storage sequences is the same as the order of the appearance of the
7 common block object lists in the scoping unit.
- 8 (2) The storage sequence formed in (1) is extended to include all storage units of any storage
9 sequence associated with it by equivalence association. The sequence shall be extended only
10 by adding storage units beyond the last storage unit. Data objects associated with an entity
11 in a common block are considered to be in that common block.

12 Only COMMON statements and EQUIVALENCE statements appearing in the scoping unit contribute
13 to common block storage sequences formed in that scoping unit.

14 5.7.2.3 Size of a common block

15 The **size of a common block** is the size of its common block storage sequence, including any extensions
16 of the sequence resulting from equivalence association.

17 5.7.2.4 Common association

18 Within a program, the common block storage sequences of all nonzero-sized common blocks with the
19 same name have the same first storage unit, and the common block storage sequences of all zero-sized
20 common blocks with the same name are storage associated with one another. Within a program, the
21 common block storage sequences of all nonzero-sized blank common blocks have the same first storage
22 unit and the storage sequences of all zero-sized blank common blocks are associated with one another and
23 with the first storage unit of any nonzero-sized blank common blocks. This results in the association of
24 objects in different scoping units. Use association or host association may cause these associated objects
25 to be accessible in the same scoping unit.

26 A nonpointer object of default integer type, default real type, double precision real type, default complex
27 type, default logical type, or numeric sequence type shall be associated only with nonpointer objects of
28 these types.

29 A nonpointer object of type default character or character sequence type shall be associated only with
30 nonpointer objects of these types.

31 A nonpointer object of a derived type that is not a numeric sequence or character sequence type shall
32 be associated only with nonpointer objects of the same type with the same type parameter values.

33 A nonpointer object of intrinsic type other than default integer, default real, double precision real,
34 default complex, default logical, or default character shall be associated only with nonpointer objects of
35 the same type and type parameters.

36 A data pointer shall be storage associated only with data pointers of the same type and rank. Data
37 pointers that are storage associated shall have deferred the same type parameters; corresponding non-

1 deferred type parameters shall have the same value. A procedure pointer shall be storage associated
 2 only with another procedure pointer; either both interfaces shall be explicit or both interfaces shall be
 3 implicit. If the interfaces are explicit, the characteristics shall be the same. If the interfaces are implicit,
 4 either both shall be subroutines or both shall be functions with the same type and type parameters.

5 An object with the TARGET attribute shall be storage associated only with another object that has
 6 the TARGET attribute and the same type and type parameters.

NOTE 5.45

A common block is permitted to contain sequences of different storage units, provided each scoping unit that accesses the common block specifies an identical sequence of storage units for the common block. For example, this allows a single common block to contain both numeric and character storage units.

Association in different scoping units between objects of default type, objects of double precision real type, and sequence structures is permitted according to the rules for equivalence objects (5.7.1).

7 **5.7.2.5 Differences between named common and blank common**

8 A blank common block has the same properties as a named common block, except for the following.

- 9 (1) Execution of a RETURN or END statement may cause data objects in a named common
 10 block to become undefined unless the common block has the SAVE attribute, but never
 11 causes data objects in blank common to become undefined (16.6.6).
 12 (2) Named common blocks of the same name shall be of the same size in all scoping units of a
 13 program in which they appear, but blank common blocks may be of different sizes.
 14 (3) A data object in a named common block may be initially defined by means of a DATA
 15 statement or type declaration statement in a block data program unit (11.3), but objects in
 16 blank common shall not be initially defined.

17 **5.7.3 Restrictions on common and equivalence**

18 An EQUIVALENCE statement shall not cause the storage sequences of two different common blocks to
 19 be associated.

20 Equivalence association shall not cause a derived-type object with default initialization to be associated
 21 with an object in a common block.

22 Equivalence association shall not cause a common block storage sequence to be extended by adding
 23 storage units preceding the first storage unit of the first object specified in a COMMON statement for
 24 the common block.

NOTE 5.46

For example, the following is not permitted:

```
COMMON /X/ A
REAL B (2)
EQUIVALENCE (A, B (2))    ! Not standard-conforming
```


1 6 Use of data objects

2 The appearance of a data object designator in a context that requires its value is termed a reference. A
 3 reference is permitted only if the data object is defined. A reference to a pointer is permitted only if the
 4 pointer is associated with a target object that is defined. A data object becomes defined with a value
 5 when events described in 16.6.5 occur.

6 R601 *variable* **is** *designator*
 7 **or** *expr*

8 C601 (R601) *designator* shall not be a constant or a subobject of a constant.

9 C602 (R601) *expr* shall be a reference to a function that has a pointer result.

10 A variable is either the data object denoted by *designator* or the target of *expr*.

11 R602 *variable-name* **is** *name*

12 C603 (R602) *variable-name* shall be the name of a variable.

13 R603 *designator* **is** *object-name*
 14 **or** *array-element*
 15 **or** *array-section*
 16 **or** *complex-part-designator*
 17 **or** *structure-component*
 18 **or** *substring*

19 R604 *logical-variable* **is** *variable*

20 C604 (R604) *logical-variable* shall be of type logical.

21 R605 *default-logical-variable* **is** *variable*

22 C605 (R605) *default-logical-variable* shall be of type default logical.

23 R606 *char-variable* **is** *variable*

24 C606 (R606) *char-variable* shall be of type character.

25 R607 *default-char-variable* **is** *variable*

26 C607 (R607) *default-char-variable* shall be of type default character.

27 R608 *int-variable* **is** *variable*

28 C608 (R608) *int-variable* shall be of type integer.

NOTE 6.1

For example, given the declarations:

```
CHARACTER (10) A, B (10)
TYPE (PERSON) P ! See Note 4.21
```

then A, B, B (1), B (1:5), P % AGE, and A (1:1) are all variables.

1 A constant (3.2.2) is a literal constant or a named constant. A literal constant is a scalar denoted by a
 2 syntactic form, which indicates its type, type parameters, and value. A named constant is a constant
 3 that has a name; the name has the PARAMETER attribute (5.3.12, 5.4.10). A reference to a constant
 4 is always permitted; redefinition of a constant is never permitted.

5 6.1 Scalars

6 A **scalar** (2.4.4) is a data entity that can be represented by a single value of the type and that is not an
 7 array (6.2). Its value, if defined, is a single element from the set of values that characterize its type.

NOTE 6.2

A scalar object of derived type has a single value that consists of the values of its components (4.5.8).

8 A scalar has rank zero.

9 6.1.1 Substrings

10 A **substring** is a contiguous portion of a character string (4.4.5). The following rules define the forms
 11 of a substring:

12 R609	<i>substring</i>	is	<i>parent-string</i> (<i>substring-range</i>)
13 R610	<i>parent-string</i>	is	<i>scalar-variable-name</i>
14		or	<i>array-element</i>
15		or	<i>scalar-structure-component</i>
16		or	<i>scalar-constant</i>
17 R611	<i>substring-range</i>	is	[<i>scalar-int-expr</i>] : [<i>scalar-int-expr</i>]

18 C609 (R610) *parent-string* shall be of type character.

19 The value of the first *scalar-int-expr* in *substring-range* is called the **starting point** and the value of
 20 the second one is called the **ending point**. The length of a substring is the number of characters in the
 21 substring and is $\text{MAX}(l - f + 1, 0)$, where f and l are the starting and ending points, respectively.

22 Let the characters in the parent string be numbered 1, 2, 3, ..., n , where n is the length of the parent
 23 string. Then the characters in the substring are those from the parent string from the starting point and
 24 proceeding in sequence up to and including the ending point. Both the starting point and the ending
 25 point shall be within the range 1, 2, ..., n unless the starting point exceeds the ending point, in which
 26 case the substring has length zero. If the starting point is not specified, the default value is 1. If the
 27 ending point is not specified, the default value is n .

28 If the parent is a variable, the substring is also a variable.

NOTE 6.3

Examples of character substrings are:

B(1)(1:5)	array element as parent string
P%NAME(1:1)	structure component as parent string
ID(4:9)	scalar variable name as parent string
'0123456789'(N:N)	character constant as parent string

29 6.1.2 Structure components

- 1 A **structure component** is part of an object of derived type; it may be referenced by an object
2 designator. A structure component may be a scalar or an array.
- 3 R612 *data-ref* **is** *part-ref* [% *part-ref*] ...
4 R613 *part-ref* **is** *part-name* [(*section-subscript-list*)] [*image-selector*]
- 5 C610 (R612) Each *part-name* except the rightmost shall be of derived type.
- 6 C611 (R612) Each *part-name* except the leftmost shall be the name of a component of the declared
7 type of the preceding *part-name*.
- 8 C612 (R612) If the rightmost *part-name* is of abstract type, *data-ref* shall be polymorphic.
- 9 C613 (R612) The leftmost *part-name* shall be the name of a data object.
- 10 C614 (R613) If a *section-subscript-list* appears, the number of *section-subscripts* shall equal the rank
11 of *part-name*.
- 12 C615 (R613) If *image-selector* appears and *part-name* is an array, *section-subscript-list* shall appear.
- 13 C616 (R612) A *data-ref* that is a co-indexed object shall not be of a type that has a pointer ultimate
14 component.
- 15 C617 (R612) If *image-selector* appears, *data-ref* shall not be, or have a direct component, of type
16 IMAGE_TEAM (13.8.3.7), C_PTR, or C_FUNPTR (15.3.3).

NOTE 6.4

This restriction is needed to avoid a disallowed pointer assignment for a component, such as

```
Z[P] = Z ! Not allowed if Z has a pointer component
Z = Z[P] ! Not allowed if Z has a pointer component
```

J3 internal note**Unresolved Technical Issue 020**

Is “Z[P] = Z” allowed if Z has an allocatable component? This would imply remote (re)allocation, which I was told was bad when I asked why we were prohibiting “array[i]” forcing it to be “array(:)[i]”. Also, constraint (C633a) in ALLOCATE prevents explicit remote allocation. If the answer is that remote reallocation is fine in this instance, we need to make that consistent (including allowing explicit remote allocation). Or if it is bad, we need to remove it consistently. Subgroup say they want to allow this, but require allocatable components in Z[P] to have the same shape as those in Z, removing the auto realloc. The editor disagrees that this is an obviously “good” technical fix; see discussion under assignment.

- 17 The rank of a *part-ref* of the form *part-name* is the rank of *part-name*. The rank of a *part-ref* that has
18 a section subscript list is the number of subscript triplets and vector subscripts in the list.
- 19 C618 (R612) There shall not be more than one *part-ref* with nonzero rank. A *part-name* to the right
20 of a *part-ref* with nonzero rank shall not have the ALLOCATABLE or POINTER attribute.
- 21 The rank of a *data-ref* is the rank of the *part-ref* with nonzero rank, if any; otherwise, the rank is zero.
22 The **base object** of a *data-ref* is the data object whose name is the leftmost part name.
- 23 The type and type parameters, if any, of a *data-ref* are those of the rightmost part name.
- 24 A *data-ref* with more than one *part-ref* is a subobject of its base object if none of the *part-names*,
25 except for possibly the rightmost, are pointers. If the rightmost *part-name* is the only pointer, then the

- 1 *data-ref* is a subobject of its base object in contexts that pertain to its pointer association status but
 2 not in any other contexts.

NOTE 6.5

If X is an object of derived type with a pointer component P, then the pointer X%P is a subobject of X when considered as a pointer – that is in contexts where it is not dereferenced.

However the target of X%P is not a subobject of X. Thus, in contexts where X%P is dereferenced to refer to the target, it is not a subobject of X.

- 3 R614 *structure-component* is *data-ref*
- 4 C619 (R614) There shall be more than one *part-ref* and the rightmost *part-ref* shall be of the form
 5 *part-name*.
- 6 A structure component shall be neither referenced nor defined before the declaration of the base object.
 7 A structure component is a pointer only if the rightmost part name is defined to have the POINTER
 8 attribute.

NOTE 6.6

Examples of structure components are:

SCALAR_PARENT%SCALAR_FIELD	scalar component of scalar parent
ARRAY_PARENT(J)%SCALAR_FIELD	component of array element parent
ARRAY_PARENT(1:N)%SCALAR_FIELD	component of array section parent

For a more elaborate example see C.3.1.

NOTE 6.7

The syntax rules are structured such that a *data-ref* that ends in a component name without a following subscript list is a structure component, even when other component names in the *data-ref* are followed by a subscript list. A *data-ref* that ends in a component name with a following subscript list is either an array element or an array section. A *data-ref* of nonzero rank that ends with a *substring-range* is an array section. A *data-ref* of zero rank that ends with a *substring-range* is a substring.

- 9 A **subcomponent** of an object of derived type is a component of that object or of a subobject of that
 10 object.
- 11 A *data-ref* shall not be a co-indexed object that has a pointer subcomponent. A *data-ref* that is a
 12 co-indexed object shall not be, or have a subcomponent, of type IMAGE_TEAM (13.8.3.7), C_PTR, or
 13 C_FUNPTR (15.3.3).

J3 internal note

Unresolved Technical Issue 089

I'm just so sure the sky is going to fall in if the user dares to do this, but why are we adding uncheckable requirements? (Well, ok it's theoretically checkable at runtime only.)

And why is this requirement placed here? None of the other stuff which uses the term “sub-component” is here, this is just the definition. It probably belongs in a separate subclause for co-indexed objects... which might clean up some of the other mess that is being made of clause 6.

We *REALLY* ought not to be placing type requirements on runtime values. That is *UNACCEPTABLE*. I am not kidding – note the careful design of the object-oriented stuff to avoid such requirements. If the co-array folk don't have the time to design the facility properly, omit it, don't go and make a complete pig's ear of the standard.

It is not beyond the wit of man to design this to be compile-time checkable.

Back to the original question – why is the sky going to fall in? I do not believe that the consequences of allowing this justify the complexity of these arbitrary restrictions. *Furthermore*, the restrictions are ineffective. There are standard-conforming ways (using TRANSFER) to “hide” the type of the value and restore it later, on another image.

1 **6.1.3 Complex parts**

2 A **complex part designator** is used to designate the real or imaginary part of a complex data object,
3 independently of the other part.

4 R615 *complex-part-designator* **is** *designator* % RE
5 **or** *designator* % IM

6 C620 (R615) The *designator* shall be of complex type.

7 If *complex-part-designator* is *designator*%RE it designates the real part of *designator*. If it is *designa-*
8 *tor*%IM it designates the imaginary part of *designator*. The type of a *complex-part-designator* is real,
9 and its kind and shape are those of the *designator*.

NOTE 6.8

The following are examples of complex part designators:

```

impedance%re      !-- Same value as REAL(impedance)
fft%im           !-- Same value as AIMAG(fft)
x%im = 0.0       !-- Sets the imaginary part of X to zero

```

10 **6.1.4 Type parameter inquiry**

11 A **type parameter inquiry** is used to inquire about a type parameter of a data object. It applies to
12 both intrinsic and derived types.

13 R616 *type-param-inquiry* **is** *designator* % *type-param-name*

14 C621 (R616) The *type-param-name* shall be the name of a type parameter of the declared type of the
15 object designated by the *designator*.

16 A deferred type parameter of a pointer that is not associated or of an unallocated allocatable variable
17 shall not be inquired about.

NOTE 6.9

A *type-param-inquiry* has a syntax like that of a structure component reference, but it does not have the same semantics. It is not a variable and thus can never be assigned to. It may be used only as a primary in an expression. It is scalar even if *designator* is an array.

The intrinsic type parameters can also be inquired about by using the intrinsic functions KIND and LEN.

NOTE 6.10

The following are examples of type parameter inquiries:

```

a%kind      !-- A is real.  Same value as KIND(a).
s%len       !-- S is character.  Same value as LEN(s).
b(10)%kind  !-- Inquiry about an array element.
p%dim       !-- P is of the derived type general_point.

```

See Note 4.28 for the definition of the `general_point` type used in the last example above.

1 6.2 Arrays

2 An **array** is a set of scalar data, all of the same type and type parameters, whose individual elements
3 are arranged in a rectangular pattern. The scalar data that make up an array are the **array elements**.

4 No order of reference to the elements of an array is indicated by the appearance of the array designator,
5 except where array element ordering (6.2.2.2) is specified.

6 6.2.1 Whole arrays

7 A **whole array** is a named array, which may be either a named constant (5.3.12, 5.4.10) or a variable;
8 no subscript list is appended to the name.

9 The appearance of a whole array variable in an executable construct specifies all the elements of the
10 array (2.4.5). An assumed-size array is permitted to appear as a whole array in an executable construct
11 only as an actual argument in a procedure reference that does not require the shape.

12 The appearance of a whole array name in a nonexecutable statement specifies the entire array except
13 for the appearance of a whole array name in an equivalence set (5.7.1.4).

14 6.2.2 Array elements and array sections

15 R617 *array-element* **is** *data-ref*

16 C622 (R617) Every *part-ref* shall have rank zero and the last *part-ref* shall contain a *subscript-list*.

17 R618 *array-section* **is** *data-ref* [(*substring-range*)]
18 **or** *complex-part-designator*

19 C623 (R618) Exactly one *part-ref* shall have nonzero rank, and either the final *part-ref* shall have a
20 *section-subscript-list* with nonzero rank, the *array-section* is a *complex-part-designator* that is
21 an array, or another *part-ref* shall have nonzero rank.

22 C624 (R618) If a *substring-range* appears, the rightmost *part-name* shall be of type character.

23 R619 *subscript* **is** *scalar-int-expr*

- 1 R620 *section-subscript* **is** *subscript*
 2 **or** *subscript-triplet*
 3 **or** *vector-subscript*
 4 R621 *subscript-triplet* **is** [*subscript*] : [*subscript*] [: *stride*]
 5 R622 *stride* **is** *scalar-int-expr*
 6 R623 *vector-subscript* **is** *int-expr*
- 7 C625 (R623) A *vector-subscript* shall be an integer array expression of rank one.
- 8 C626 (R621) The second subscript shall not be omitted from a *subscript-triplet* in the last dimension
 9 of an assumed-size array.
- 10 An array element is a scalar. An array section is an array. If a *substring-range* is present in an *array-*
 11 *section*, each element is the designated substring of the corresponding element of the array section.

NOTE 6.11

For example, with the declarations:

```
REAL A (10, 10)
CHARACTER (LEN = 10) B (5, 5, 5)
```

A (1, 2) is an array element, A (1:N:2, M) is a rank-one array section, and B (:, :, :) (2:3) is an array of shape (5, 5, 5) whose elements are substrings of length 2 of the corresponding elements of B.

NOTE 6.12

Unless otherwise specified, an array element or array section does not have an attribute of the whole array. In particular, an array element or an array section does not have the POINTER or ALLOCATABLE attribute.

NOTE 6.13

Examples of array elements and array sections are:

```
ARRAY_A(1:N:2)%ARRAY_B(I, J)%STRING(K)(:)    array section
SCALAR_PARENT%ARRAY_FIELD(J)                array element
SCALAR_PARENT%ARRAY_FIELD(1:N)              array section
SCALAR_PARENT%ARRAY_FIELD(1:N)%SCALAR_FIELD array section
```

12 **6.2.2.1 Array elements**

- 13 The value of a subscript in an array element shall be within the bounds for that dimension.

14 **6.2.2.2 Array element order**

- 15 The elements of an array form a sequence known as the **array element order**. The position of an array
 16 element in this sequence is determined by the subscript order value of the subscript list designating the
 17 element. The subscript order value is computed from the formulas in Table 6.1.

Table 6.1: **Subscript order value**

Rank	Subscript bounds	Subscript list	Subscript order value
1	$j_1:k_1$	s_1	$1 + (s_1 - j_1)$
2	$j_1:k_1, j_2:k_2$	s_1, s_2	$1 + (s_1 - j_1) + (s_2 - j_2) \times d_1$

Rank	Subscript bounds	Subscript list	Subscript order value
3	$j_1:k_1, j_2:k_2, j_3:k_3$	s_1, s_2, s_3	$1 + (s_1 - j_1)$ $+(s_2 - j_2) \times d_1$ $+(s_3 - j_3) \times d_2 \times d_1$
.	.	.	.
.	.	.	.
.	.	.	.
15	$j_1:k_1, \dots, j_{15}:k_{15}$	s_1, \dots, s_{15}	$1 + (s_1 - j_1)$ $+(s_2 - j_2) \times d_1$ $+(s_3 - j_3) \times d_2 \times d_1$ $+\dots$ $+(s_{15} - j_{15}) \times d_{14}$ $\times d_{13} \times \dots \times d_1$
Notes for Table 6.1:			
1) $d_i = \max(k_i - j_i + 1, 0)$ is the size of the i th dimension.			
2) If the size of the array is nonzero, $j_i \leq s_i \leq k_i$ for all $i = 1, 2, \dots, 15$.			

1 6.2.2.3 Array sections

2 An **array section** is an array subobject optionally followed by a substring range.

3 In an *array-section* having a *section-subscript-list*, each *subscript-triplet* and *vector-subscript* in the
4 section subscript list indicates a sequence of subscripts, which may be empty. Each subscript in such a
5 sequence shall be within the bounds for its dimension unless the sequence is empty. The array section is
6 the set of elements from the array determined by all possible subscript lists obtainable from the single
7 subscripts or sequences of subscripts specified by each section subscript.

8 In an *array-section* with no *section-subscript-list*, the rank and shape of the array is the rank and shape
9 of the *part-ref* with nonzero rank; otherwise, the rank of the array section is the number of subscript
10 triplets and vector subscripts in the section subscript list. The shape is the rank-one array whose i th
11 element is the number of integer values in the sequence indicated by the i th subscript triplet or vector
12 subscript. If any of these sequences is empty, the array section has size zero. The subscript order of the
13 elements of an array section is that of the array data object that the array section represents.

14 6.2.2.3.1 Subscript triplet

15 A subscript triplet designates a regular sequence of subscripts consisting of zero or more subscript values.
16 The third expression in the subscript triplet is the increment between the subscript values and is called
17 the **stride**. The subscripts and stride of a subscript triplet are optional. An omitted first subscript in a
18 subscript triplet is equivalent to a subscript whose value is the lower bound for the array and an omitted
19 second subscript is equivalent to the upper bound. An omitted stride is equivalent to a stride of 1.

20 The stride shall not be zero.

21 When the stride is positive, the subscripts specified by a triplet form a regularly spaced sequence of
22 integers beginning with the first subscript and proceeding in increments of the stride to the largest such
23 integer not greater than the second subscript; the sequence is empty if the first subscript is greater than
24 the second.

NOTE 6.14

For example, suppose an array is declared as A (5, 4, 3). The section A (3 : 5, 2, 1 : 2) is the array of shape (3, 2):

NOTE 6.14 (cont.)

A (3, 2, 1)	A (3, 2, 2)
A (4, 2, 1)	A (4, 2, 2)
A (5, 2, 1)	A (5, 2, 2)

- 1 When the stride is negative, the sequence begins with the first subscript and proceeds in increments of
 2 the stride down to the smallest such integer equal to or greater than the second subscript; the sequence
 3 is empty if the second subscript is greater than the first.

NOTE 6.15

For example, if an array is declared B (10), the section B (9 : 1 : -2) is the array of shape (5) whose elements are B (9), B (7), B (5), B (3), and B (1), in that order.

NOTE 6.16

A subscript in a subscript triplet need not be within the declared bounds for that dimension if all values used in selecting the array elements are within the declared bounds.

For example, if an array is declared as B (10), the array section B (3 : 11 : 7) is the array of shape (2) consisting of the elements B (3) and B (10), in that order.

4 6.2.2.3.2 Vector subscript

- 5 A **vector subscript** designates a sequence of subscripts corresponding to the values of the elements
 6 of the expression. Each element of the expression shall be defined. A **many-one array section** is an
 7 array section with a vector subscript having two or more elements with the same value. A many-one
 8 array section shall appear neither on the left of the equals in an assignment statement nor as an input
 9 item in a READ statement.

- 10 An array section with a vector subscript shall not be argument associated with a dummy array that
 11 is defined or redefined. An array section with a vector subscript shall not be the target in a pointer
 12 assignment statement. An array section with a vector subscript shall not be an internal file.

NOTE 6.17

For example, suppose Z is a two-dimensional array of shape (5, 7) and U and V are one-dimensional arrays of shape (3) and (4), respectively. Assume the values of U and V are:

U = (/ 1, 3, 2 /)
 V = (/ 2, 1, 1, 3 /)

Then Z (3, V) consists of elements from the third row of Z in the order:

Z (3, 2) Z (3, 1) Z (3, 1) Z (3, 3)

and Z (U, 2) consists of the column elements:

Z (1, 2) Z (3, 2) Z (2, 2)

and Z (U, V) consists of the elements:

Z (1, 2) Z (1, 1) Z (1, 1) Z (1, 3)
 Z (3, 2) Z (3, 1) Z (3, 1) Z (3, 3)
 Z (2, 2) Z (2, 1) Z (2, 1) Z (2, 3)

NOTE 6.17 (cont.)

Because $Z(3, V)$ and $Z(U, V)$ contain duplicate elements from Z , the sections $Z(3, V)$ and $Z(U, V)$ shall not be redefined as sections.

1 **6.2.3 Image selectors**

2 An **image selector** specifies the image index for co-array data.

3 R624 *image-selector* **is** *lbracket co-subscript-list rbracket*
 4 R625 *co-subscript* **is** *scalar-int-expr*

5 The number of co-subscripts shall be equal to the co-rank of the object. The value of a co-subscript in
 6 an image selector shall be within the co-bounds for its co-dimension. Taking account of the co-bounds,
 7 the co-subscript list in an image selector determines the image index in the same way that a subscript
 8 list in an array element determines the subscript order value (6.2.2.2), taking account of the bounds. An
 9 image selector shall specify an image index value that is not greater than the number of images.

NOTE 6.18

For example, if there are 16 images and the co-array A is declared

```
REAL :: A(10)[5,*]
```

A(:)[1,4] is valid because it specifies image 16, but A(:)[2,4] is invalid because it specifies image 17.

10 **6.3 Dynamic association**

11 Dynamic control over the allocation, association, and deallocation of pointer targets is provided by
 12 the ALLOCATE, NULLIFY, and DEALLOCATE statements and pointer assignment. ALLOCATE
 13 (6.3.1) creates targets for pointers; pointer assignment (7.4.2) associates pointers with existing targets;
 14 NULLIFY (6.3.2) disassociates pointers from targets, and DEALLOCATE (6.3.3) deallocates targets.
 15 Dynamic association applies to scalars and arrays of any type.

16 The ALLOCATE and DEALLOCATE statements also are used to create and deallocate variables with
 17 the ALLOCATABLE attribute.

NOTE 6.19

Detailed remarks regarding pointers and dynamic association are in C.3.3.

18 **6.3.1 ALLOCATE statement**

19 The **ALLOCATE statement** dynamically creates pointer targets and allocatable variables.

20 R626 *allocate-stmt* **is** ALLOCATE ([*type-spec* ::] *allocation-list* ■
 21 ■ [*alloc-opt-list*])
 22 R627 *alloc-opt* **is** ERRMSG = *errmsg-variable*
 23 **or** MOLD = *source-expr*
 24 **or** SOURCE = *source-expr*
 25 **or** STAT = *stat-variable*
 26 R628 *stat-variable* **is** *scalar-int-variable*
 27 R629 *errmsg-variable* **is** *scalar-default-char-variable*
 28 R630 *source-expr* **is** *expr*
 29 R631 *allocation* **is** *allocate-object* [(*allocate-shape-spec-list*)] ■

- 1 ■ [lbracket allocate-co-array-spec rbracket]
- 2 R632 *allocate-object* **is** *variable-name*
- 3 **or** *structure-component*
- 4 R633 *allocate-shape-spec* **is** [*lower-bound-expr* :] *upper-bound-expr*
- 5 R634 *lower-bound-expr* **is** *scalar-int-expr*
- 6 R635 *upper-bound-expr* **is** *scalar-int-expr*
- 7 R636 *allocate-co-array-spec* **is** [*allocate-co-shape-spec-list* ,] [*lower-bound-expr* :] *
- 8 R637 *allocate-co-shape-spec* **is** [*lower-bound-expr* :] *upper-bound-expr*
- 9 C627 (R632) Each *allocate-object* shall be a nonprocedure pointer or an allocatable variable.
- 10 C628 (R626) If any *allocate-object* in the statement has a deferred type parameter, either *type-spec* or
- 11 *source-expr* shall appear.
- 12 C629 (R626) If a *type-spec* appears, it shall specify a type with which each *allocate-object* is type
- 13 compatible.
- 14 C630 (R626) If any *allocate-object* is unlimited polymorphic or is of abstract type, either *type-spec* or
- 15 *source-expr* shall appear.
- 16 C631 (R626) A *type-param-value* in a *type-spec* shall be an asterisk if and only if each *allocate-object*
- 17 is a dummy argument for which the corresponding type parameter is assumed.
- 18 C632 (R626) If a *type-spec* appears, the kind type parameter values of each *allocate-object* shall be
- 19 the same as the corresponding type parameter values of the *type-spec*.
- 20 C633 (R631) An *allocate-shape-spec-list* shall appear if and only if the *allocate-object* is an array. An
- 21 *allocate-co-array-spec* shall appear if and only if the *allocate-object* is a co-array.
- 22 C634 (R631) The number of *allocate-shape-specs* in an *allocate-shape-spec-list* shall be the same as the
- 23 rank of the *allocate-object*. The number of *allocate-co-shape-specs* in an *allocate-co-array-spec*
- 24 shall be one less than the co-rank of the *allocate-object*.
- 25 C635 (R627) No *alloc-opt* shall appear more than once in a given *alloc-opt-list*.
- 26 C636 (R626) At most one of *source-expr* and *type-spec* shall appear.
- 27 C637 (R626) Each *allocate-object* shall be type compatible (4.3.1.3) with *source-expr*. If SOURCE=
- 28 appears, *source-expr* shall be a scalar or have the same rank as each *allocate-object*.
- 29 C638 (R626) Corresponding kind type parameters of *allocate-object* and *source-expr* shall have the
- 30 same values.
- 31 C639 (R626) *type-spec* shall not specify a type that has a co-array ultimate component.
- 32 C640 (R630) The declared type of *source-expr* shall not have a co-array ultimate component.
- 33 C641 (R632) An *allocate-object* shall not be a co-indexed object.

NOTE 6.20

If a co-array is of a derived type that has an allocatable component, the component shall be allocated by its own image:

```

TYPE(SOMETHING), ALLOCATABLE :: T[:]
...
ALLOCATE(T[*])           ! Allowed - implies synchronization
ALLOCATE(T%AA(N))       ! Allowed - allocated by its own image

```

NOTE 6.20 (cont.)

```

ALLOCATE(T[Q]%AAC(N))    ! Not allowed, because it is not
                          ! necessarily executed on image Q.

```

- 1 An *allocate-object* or a bound or type parameter of an *allocate-object* shall not depend on the value of
2 *stat-variable*, the value of *errmsg-variable*, or on the value, bounds, length type parameters, allocation
3 status, or association status of any *allocate-object* in the same ALLOCATE statement.
- 4 Neither *stat-variable*, *source-expr*, nor *errmsg-variable* shall be allocated within the ALLOCATE state-
5 ment in which it appears; nor shall they depend on the value, bounds, length type parameters, allocation
6 status, or association status of any *allocate-object* in the same ALLOCATE statement.
- 7 If *type-spec* is specified, each *allocate-object* is allocated with the specified dynamic type and type pa-
8 rameter values; if *source-expr* is specified, each *allocate-object* is allocated with the dynamic type and
9 type parameter values of *source-expr*; otherwise, each *allocate-object* is allocated with its dynamic type
10 the same as its declared type.
- 11 If *type-spec* appears and the value of a type parameter it specifies differs from the value of the corre-
12 sponding nondeferred type parameter specified in the declaration of any *allocate-object*, an error condition
13 occurs. If the value of a nondeferred length type parameter of an *allocate-object* differs from the value
14 of the corresponding type parameter of *source-expr*, an error condition occurs.
- 15 If a *type-param-value* in a *type-spec* in an ALLOCATE statement is an asterisk, it denotes the current
16 value of that assumed type parameter. If it is an expression, subsequent redefinition or undefinition of
17 any entity in the expression does not affect the type parameter value.

NOTE 6.21

An example of an ALLOCATE statement is:

```

ALLOCATE (X (N), B (-3 : M, 0:9), STAT = IERR_ALLOC)

```

- 18 When an ALLOCATE statement is executed for an array, the values of the lower bound and upper
19 bound expressions determine the bounds of the array. Subsequent redefinition or undefinition of any
20 entities in the bound expressions do not affect the array bounds. If the lower bound is omitted, the
21 default value is 1. If the upper bound is less than the lower bound, the extent in that dimension is zero
22 and the array has zero size.
- 23 When an ALLOCATE statement is executed for a co-array, the values of the lower co-bound and upper
24 co-bound expressions determine the co-bounds of the co-array. Subsequent redefinition or undefinition
25 of any entities in the co-bound expressions do not affect the co-bounds. If the lower co-bound is omitted,
26 the default value is 1. The upper co-bound shall not be less than the lower co-bound.

NOTE 6.22

An *allocate-object* may be of type character with zero character length.

- 27 If SOURCE= appears, *source-expr* shall be conformable (2.4.5) with *allocation*. If the value of a non-
28 deferred length type parameter of *allocate-object* is different from the value of the corresponding type
29 parameter of *source-expr*, an error condition occurs. If the allocation is successful, the value of *allocate-*
30 *object* becomes that of *source-expr*.
- 31 If MOLD= appears and *source-expr* is a variable, its value need not be defined.

J3 internal note

Unresolved Technical Issue 081

Should we require allocatables to be allocated, or only when they have deferred type parameters? Furthermore, we should probably extend MOLD= (and SOURCE=) to be able to specify the shape, when no *array-spec* is provided, so the user can write

```
ALLOCATE(A,MOLD=B)
```

instead of

```
ALLOCATE(A(LBOUND(B,1):UBOUND(B,1),LBOUND(B,2):UBOUND(B,2)),MOLD=B)
```

NOTE 6.23

The *source-expr* can be an array or scalar independently of whether an *allocate-object* is an array or a scalar. For MOLD=*source-expr*, only the dynamic type and type parameter values of *source-expr* are relevant; its rank, shape, and value are not.

NOTE 6.24

An example of an ALLOCATE statement in which the value and dynamic type are determined by reference to another object is:

```
CLASS(*), ALLOCATABLE :: NEW
CLASS(*), POINTER :: OLD
! ...
ALLOCATE (NEW, SOURCE=OLD) ! Allocate NEW with the value and dynamic type of OLD
```

A more extensive example is given in C.3.2.

1 If the STAT= specifier appears, successful execution of the ALLOCATE statement causes the *stat-*
 2 *variable* to become defined with a value of zero. If an error condition occurs during the execution
 3 of the ALLOCATE statement, the *stat-variable* becomes defined with a processor-dependent positive
 4 integer value and each *allocate-object* will have a processor-dependent status; each *allocate-object* that
 5 was successfully allocated shall have an allocation status of allocated or a pointer association status of
 6 associated; each *allocate-object* that was not successfully allocated shall retain its previous allocation
 7 status or pointer association status.

8 If an error condition occurs during execution of an ALLOCATE statement that does not contain the
 9 STAT= specifier, execution of the program is terminated.

10 The ERRMSG= specifier is described in 6.3.1.3.

11 6.3.1.1 Allocation of allocatable variables

12 The allocation status of an allocatable entity is one of the following at any time.

- 13 (1) The status of an allocatable variable becomes **allocated** if it is allocated by an ALLOCATE
 14 statement, if it is allocated during assignment, or if it is given that status by the allocation
 15 transfer procedure (13.7.124). An allocatable variable with this status may be referenced,
 16 defined, or deallocated; allocating it causes an error condition in the ALLOCATE statement.
 17 The intrinsic function ALLOCATED (13.7.10) returns true for such a variable.
- 18 (2) An allocatable variable has a status of **unallocated** if it is not allocated. The status of
 19 an allocatable variable becomes unallocated if it is deallocated (6.3.3) or if it is given that
 20 status by the allocation transfer procedure. An allocatable variable with this status shall

1 not be referenced or defined. It shall not be supplied as an actual argument corresponding
2 to a nonallocatable dummy argument, except to certain intrinsic inquiry functions. It may
3 be allocated with the ALLOCATE statement. Deallocating it causes an error condition in
4 the DEALLOCATE statement. The intrinsic function ALLOCATED (13.7.10) returns false
5 for such a variable.

6 At the beginning of execution of a program, allocatable variables are unallocated.

7 A saved allocatable object has an initial status of unallocated. The status may change during the
8 execution of the program.

9 When the allocation status of an allocatable variable changes, the allocation status of any associated
10 allocatable variable changes accordingly. Allocation of an allocatable variable establishes values for the
11 deferred type parameters of all associated allocatable variables.

12 An unsaved allocatable local variable of a procedure has a status of unallocated at the beginning of
13 each invocation of the procedure; the status may change during execution of the procedure. An unsaved
14 allocatable local variable of a module or submodule, or a subobject thereof, has an initial status of
15 unallocated; the status may change during execution of the program. An unsaved local variable of a
16 construct has a status of unallocated at the beginning of each execution of the construct; the status may
17 change during execution of the construct.

18 When an object of derived type is created by an ALLOCATE statement, any allocatable ultimate
19 components have an allocation status of unallocated.

20 The value of each *lower-bound-expr* and each *upper-bound-expr* in an *allocate-co-array-spec* shall be the
21 same on each image.

22 If an *allocation* specifies a co-array, its dynamic type and the values of each length type parameter shall
23 be the same on each image.

24 There is implicit synchronization of all images in association with each ALLOCATE statement that
25 involves one or more co-arrays. On each image, execution of the segment (8.5.1) following the statement
26 is delayed until all other images have executed the same statement the same number of times.

NOTE 6.25

When an image executes an ALLOCATE statement, communication is not necessarily involved apart from any required for synchronization. The image allocates its co-array and records how the corresponding co-arrays on other images are to be addressed. The processor is not required to detect violations of the rule that the bounds are the same on all images, nor is it responsible for detecting or resolving deadlock problems (such as two images waiting on different ALLOCATE statements).

27 6.3.1.2 Allocation of pointer targets

28 Allocation of a pointer creates an object that implicitly has the TARGET attribute. Following successful
29 execution of an ALLOCATE statement for a pointer, the pointer is associated with the target and may
30 be used to reference or define the target. Additional pointers may become associated with the pointer
31 target or a part of the pointer target by pointer assignment. It is not an error to allocate a pointer
32 that is already associated with a target. In this case, a new pointer target is created as required by the
33 attributes of the pointer and any array bounds, type, and type parameters specified by the ALLOCATE
34 statement. The pointer is then associated with this new target. Any previous association of the pointer
35 with a target is broken. If the previous target had been created by allocation, it becomes inaccessible
36 unless other pointers are associated with it. The ASSOCIATED intrinsic function (13.7.15) may be used
37 to determine whether a pointer that does not have undefined association status is associated.

1 At the beginning of execution of a function whose result is a pointer, the association status of the result
 2 pointer is undefined. Before such a function returns, it shall either associate a target with this pointer
 3 or cause the association status of this pointer to become disassociated.

4 **6.3.1.3 ERRMSG= specifier**

5 If an error condition occurs during execution of an ALLOCATE or DEALLOCATE statement, the
 6 processor shall assign an explanatory message to *errmsg-variable*. If no such condition occurs, the
 7 processor shall not change the value of *errmsg-variable*.

8 **6.3.2 NULLIFY statement**

9 The **NULLIFY statement** causes pointers to be disassociated.

10 R638 *nullify-stmt* **is** NULLIFY (*pointer-object-list*)
 11 R639 *pointer-object* **is** *variable-name*
 12 **or** *structure-component*
 13 **or** *proc-pointer-name*

14 C642 (R639) Each *pointer-object* shall have the POINTER attribute.

15 A *pointer-object* shall not depend on the value, bounds, or association status of another *pointer-object*
 16 in the same NULLIFY statement.

NOTE 6.26

When a NULLIFY statement is applied to a polymorphic pointer (4.3.1.3), its dynamic type becomes the declared type.

17 **6.3.3 DEALLOCATE statement**

18 The **DEALLOCATE statement** causes allocatable variables to be deallocated; it causes pointer tar-
 19 gets to be deallocated and the pointers to be disassociated.

20 R640 *deallocate-stmt* **is** DEALLOCATE (*allocate-object-list* [, *dealloc-opt-list*])

21 C643 (R640) Each *allocate-object* shall be a nonprocedure pointer or an allocatable variable.

22 R641 *dealloc-opt* **is** STAT = *stat-variable*
 23 **or** ERRMSG = *errmsg-variable*

24 C644 (R641) No *dealloc-opt* shall appear more than once in a given *dealloc-opt-list*.

25 An *allocate-object* shall not depend on the value, bounds, allocation status, or association status of
 26 another *allocate-object* in the same DEALLOCATE statement; it also shall not depend on the value of
 27 the *stat-variable* or *errmsg-variable* in the same DEALLOCATE statement.

28 Neither *stat-variable* nor *errmsg-variable* shall be deallocated within the same DEALLOCATE state-
 29 ment; they also shall not depend on the value, bounds, allocation status, or association status of any
 30 *allocate-object* in the same DEALLOCATE statement.

31 If the STAT= specifier appears, successful execution of the DEALLOCATE statement causes the *stat-*
 32 *variable* to become defined with a value of zero. If an error condition occurs during the execution of
 33 the DEALLOCATE statement, the *stat-variable* becomes defined with a processor-dependent positive
 34 integer value and each *allocate-object* that was successfully deallocated shall have an allocation status of
 35 unallocated or a pointer association status of disassociated. Each *allocate-object* that was not successfully
 36 deallocated shall retain its previous allocation status or pointer association status.

NOTE 6.27

The status of objects that were not successfully deallocated can be individually checked with the ALLOCATED or ASSOCIATED intrinsic functions.

- 1 If an error condition occurs during execution of a DEALLOCATE statement that does not contain the
- 2 STAT= specifier, execution of the program is terminated.
- 3 The ERRMSG= specifier is described in 6.3.1.3.

NOTE 6.28

An example of a DEALLOCATE statement is:

```
DEALLOCATE (X, B)
```

4 **6.3.3.1 Deallocation of allocatable variables**

- 5 Deallocating an unallocated allocatable variable causes an error condition in the DEALLOCATE state-
- 6 ment. Deallocating an allocatable variable with the TARGET attribute causes the pointer association
- 7 status of any pointer associated with it to become undefined.
- 8 When the execution of a procedure is terminated by execution of a RETURN or END statement, an
- 9 unsaved allocatable local variable of the procedure retains its allocation and definition status if it is a
- 10 function result variable or a subobject thereof; otherwise, it is deallocated.
- 11 When a BLOCK construct terminates, an unsaved allocatable local variable of the construct is deallo-
- 12 cated.

NOTE 6.29

The ALLOCATED intrinsic function may be used to determine whether a variable is allocated or unallocated.

- 13 If an unsaved allocatable local variable of a module or submodule is allocated when execution of a
- 14 RETURN or END statement results in no active scoping unit referencing the module or submodule, it
- 15 is processor-dependent whether the object retains its allocation status or is deallocated.

NOTE 6.30

The following example illustrates the effects of SAVE on allocation status.

```
MODULE MOD1
  TYPE INITIALIZED_TYPE
    INTEGER :: I = 1 ! Default initialization
  END TYPE INITIALIZED_TYPE
  SAVE :: SAVED1, SAVED2
  INTEGER :: SAVED1, UNSAVED1
  TYPE(INITIALIZED_TYPE) :: SAVED2, UNSAVED2
  ALLOCATABLE :: SAVED1(:), SAVED2(:), UNSAVED1(:), UNSAVED2(:)
END MODULE MOD1
PROGRAM MAIN
  CALL SUB1 ! The values returned by the ALLOCATED intrinsic calls
            ! in the PRINT statement are:
            ! .FALSE., .FALSE., .FALSE., and .FALSE.
            ! Module MOD1 is used, and its variables are allocated.
            ! After return from the subroutine, whether the variables
```

NOTE 6.30 (cont.)

```

! which were not specified with the SAVE attribute
! retain their allocation status is processor dependent.
CALL SUB1 ! The values returned by the first two ALLOCATED intrinsic
! calls in the PRINT statement are:
! .TRUE., .TRUE.
! The values returned by the second two ALLOCATED
! intrinsic calls in the PRINT statement are
! processor dependent and each could be either
! .TRUE. or .FALSE.

CONTAINS
  SUBROUTINE SUB1
    USE MOD1 ! Brings in saved and unsaved variables.
    PRINT *, ALLOCATED(SAVED1), ALLOCATED(SAVED2), &
      ALLOCATED(UNSAVED1), ALLOCATED(UNSAVED2)
    IF (.NOT. ALLOCATED(SAVED1)) ALLOCATE(SAVED1(10))
    IF (.NOT. ALLOCATED(SAVED2)) ALLOCATE(SAVED2(10))
    IF (.NOT. ALLOCATED(UNSAVED1)) ALLOCATE(UNSAVED1(10))
    IF (.NOT. ALLOCATED(UNSAVED2)) ALLOCATE(UNSAVED2(10))
  END SUBROUTINE SUB1
END PROGRAM MAIN

```

- 1 If an executable construct references a function whose result is either allocatable or a structure with
- 2 a subobject that is allocatable, and the function reference is executed, an allocatable result and any
- 3 subobject that is an allocated allocatable entity in the result returned by the function is deallocated
- 4 after execution of the innermost executable construct containing the reference.

- 5 If a function whose result is either allocatable or a structure with an allocatable object is referenced in
- 6 the specification part of a scoping unit or BLOCK construct, and the function reference is executed, an
- 7 allocatable result and any subobject that is an allocated allocatable entity in the result returned by the
- 8 function is deallocated before execution of the executable constructs of the scoping unit or block.

- 9 When a procedure is invoked, any allocated allocatable object that is an actual argument associated with
- 10 an INTENT(OUT) allocatable dummy argument is deallocated; any allocated allocatable object that is
- 11 a subobject of an actual argument associated with an INTENT(OUT) dummy argument is deallocated.

- 12 When an intrinsic assignment statement (7.4.1.3) is executed, any allocated allocatable subobject of the
- 13 variable is deallocated before the assignment takes place.

- 14 When a variable of derived type is deallocated, any allocated allocatable subobject is deallocated.

- 15 If an allocatable component is a subobject of a finalizable object, that object is finalized before the
- 16 component is automatically deallocated.

- 17 The effect of automatic deallocation is the same as that of a DEALLOCATE statement without a
- 18 *dealloc-opt-list*.

NOTE 6.31

In the following example:

```

SUBROUTINE PROCESS
  REAL, ALLOCATABLE :: TEMP(:)
  REAL, ALLOCATABLE, SAVE :: X(:)
  ...

```

NOTE 6.31 (cont.)

END SUBROUTINE PROCESS

on return from subroutine PROCESS, the allocation status of X is preserved because X has the SAVE attribute. TEMP does not have the SAVE attribute, so it will be deallocated if it was allocated. On the next invocation of PROCESS, TEMP will have an allocation status of unallocated.

1 There is implicit synchronization of all images in association with each DEALLOCATE statement that
2 involves one or more co-arrays. On each image, execution of the segment (8.5.1) following the statement
3 is delayed until all other images have executed the same statement the same number of times.

4 There is also an implicit synchronization of all images in association with the deallocation of a co-array
5 or co-array subcomponent caused by the execution of a RETURN or END statement or the termination
6 of a BLOCK construct.

7 **6.3.3.2 Deallocation of pointer targets**

8 If a pointer appears in a DEALLOCATE statement, its association status shall be defined. Deallocating
9 a pointer that is disassociated or whose target was not created by an ALLOCATE statement causes an
10 error condition in the DEALLOCATE statement. If a pointer is associated with an allocatable entity,
11 the pointer shall not be deallocated.

12 If a pointer appears in a DEALLOCATE statement, it shall be associated with the whole of an object
13 that was created by allocation. Deallocating a pointer target causes the pointer association status of
14 any other pointer that is associated with the target or a portion of the target to become undefined.

1 7 Expressions and assignment

2 This clause describes the formation, interpretation, and evaluation rules for expressions, intrinsic and
3 defined assignment, pointer assignment, masked array assignment (WHERE), and FORALL.

4 7.1 Expressions

5 An **expression** represents either a data reference or a computation, and its value is either a scalar or
6 an array. An expression is formed from operands, operators, and parentheses.

7 An operand is either a scalar or an array. An operation is either intrinsic or defined (7.2). More
8 complicated expressions can be formed using operands which are themselves expressions.

9 Evaluation of an expression produces a value, which has a type, type parameters (if appropriate), and a
10 shape (7.1.4).

11 7.1.1 Form of an expression

12 An expression is defined in terms of several categories: primary, level-1 expression, level-2 expression,
13 level-3 expression, level-4 expression, and level-5 expression.

14 These categories are related to the different operator precedence levels and, in general, are defined in
15 terms of other categories. The simplest form of each expression category is a *primary*. The rules given
16 below specify the syntax of an expression. The semantics are specified in 7.2.

17 7.1.1.1 Primary

18	R701	<i>primary</i>	is	<i>constant</i>
19			or	<i>designator</i>
20			or	<i>array-constructor</i>
21			or	<i>structure-constructor</i>
22			or	<i>function-reference</i>
23			or	<i>type-param-inquiry</i>
24			or	<i>type-param-name</i>
25			or	<i>(expr)</i>

26 C701 (R701) The *type-param-name* shall be the name of a type parameter.

27 C702 (R701) The *designator* shall not be a whole assumed-size array.

NOTE 7.1

Examples of a *primary* are:

<u>Example</u>	<u>Syntactic class</u>
1.0	<i>constant</i>
'ABCDEFGHIJKLMNQRSTUWXYZ' (I:I)	<i>designator</i>
(/ 1.0, 2.0 /)	<i>array-constructor</i>
PERSON (12, 'Jones')	<i>structure-constructor</i>
F (X, Y)	<i>function-reference</i>
X%KIND	<i>type-param-inquiry</i>

NOTE 7.1 (cont.)

KIND (S + T)	<i>type-param-name</i> (<i>expr</i>)
-----------------	---

1 **7.1.1.2 Level-1 expressions**

2 Defined unary operators have the highest operator precedence (Table 7.10). Level-1 expressions are
3 primaries optionally operated on by defined unary operators:

4 R702 *level-1-expr* **is** [*defined-unary-op*] *primary*

5 R703 *defined-unary-op* **is** . *letter* [*letter*]

6 C703 (R703) A *defined-unary-op* shall not contain more than 63 letters and shall not be the same as
7 any *intrinsic-operator* or *logical-literal-constant*.

NOTE 7.2

Simple examples of a level-1 expression are:

<u>Example</u>	<u>Syntactic class</u>
A	<i>primary</i> (R701)
.INVERSE. B	<i>level-1-expr</i> (R702)

A more complicated example of a level-1 expression is:

.INVERSE. (A + B)

8 **7.1.1.3 Level-2 expressions**

9 Level-2 expressions are level-1 expressions optionally involving the numeric operators *power-op*, *mult-op*,
10 and *add-op*.

11 R704 *mult-operand* **is** *level-1-expr* [*power-op mult-operand*]

12 R705 *add-operand* **is** [*add-operand mult-op*] *mult-operand*

13 R706 *level-2-expr* **is** [[*level-2-expr*] *add-op*] *add-operand*

14 R707 *power-op* **is** **

15 R708 *mult-op* **is** *

16 **or** /

17 R709 *add-op* **is** +

18 **or** -

NOTE 7.3

Simple examples of a level-2 expression are:

<u>Example</u>	<u>Syntactic class</u>	<u>Remarks</u>
A	<i>level-1-expr</i>	A is a <i>primary</i> . (R702)
B ** C	<i>mult-operand</i>	B is a <i>level-1-expr</i> , ** is a <i>power-op</i> , and C is a <i>mult-operand</i> . (R704)
D * E	<i>add-operand</i>	D is an <i>add-operand</i> , * is a <i>mult-op</i> , and E is a <i>mult-operand</i> . (R705)
+1	<i>level-2-expr</i>	+ is an <i>add-op</i> and 1 is an <i>add-operand</i> . (R706)
F - I	<i>level-2-expr</i>	F is a <i>level-2-expr</i> , - is an <i>add-op</i> , and I is an <i>add-operand</i> . (R706)

NOTE 7.3 (cont.)

A more complicated example of a level-2 expression is:

```
- A + D * E + B ** C
```

1 **7.1.1.4 Level-3 expressions**

2 Level-3 expressions are level-2 expressions optionally involving the character operator and bits concatenate-
3 tion operator *concat-op*.

4 R710 *level-3-expr* is [*level-3-expr concat-op*] *level-2-expr*

5 R711 *concat-op* is //

NOTE 7.4

Simple examples of a level-3 expression are:

<u>Example</u>	<u>Syntactic class</u>
A	<i>level-2-expr</i> (R706)
B // C	<i>level-3-expr</i> (R710)

A more complicated example of a level-3 expression is:

```
X // Y // 'ABCD'
```

6 **7.1.1.5 Level-4 expressions**

7 Level-4 expressions are level-3 expressions optionally involving the relational operators *rel-op*.

8 R712 *level-4-expr* is [*level-3-expr rel-op*] *level-3-expr*

9 R713 *rel-op* is .EQ.

10 or .NE.

11 or .LT.

12 or .LE.

13 or .GT.

14 or .GE.

15 or ==

16 or /=

17 or <

18 or <=

19 or >

20 or >=

NOTE 7.5

Simple examples of a level-4 expression are:

<u>Example</u>	<u>Syntactic class</u>
A	<i>level-3-expr</i> (R710)
B == C	<i>level-4-expr</i> (R712)
D < E	<i>level-4-expr</i> (R712)

A more complicated example of a level-4 expression is:

```
(A + B) /= C
```

1 **7.1.1.6 Level-5 expressions**

2 Level-5 expressions are level-4 expressions optionally involving the logical and bits operators *not-op*,
3 *and-op*, *or-op*, and *equiv-op*.

4	R714	<i>and-operand</i>	is	[<i>not-op</i>] <i>level-4-expr</i>
5	R715	<i>or-operand</i>	is	[<i>or-operand and-op</i>] <i>and-operand</i>
6	R716	<i>equiv-operand</i>	is	[<i>equiv-operand or-op</i>] <i>or-operand</i>
7	R717	<i>level-5-expr</i>	is	[<i>level-5-expr equiv-op</i>] <i>equiv-operand</i>
8	R718	<i>not-op</i>	is	.NOT.
9	R719	<i>and-op</i>	is	.AND.
10	R720	<i>or-op</i>	is	.OR.
11	R721	<i>equiv-op</i>	is	.EQV.
12			or	.NEQV.
13			or	.XOR.

NOTE 7.6

Simple examples of a level-5 expression are:

<u>Example</u>	<u>Syntactic class</u>
A	<i>level-4-expr</i> (R712)
.NOT. B	<i>and-operand</i> (R714)
C .AND. D	<i>or-operand</i> (R715)
E .OR. F	<i>equiv-operand</i> (R716)
G .EQV. H	<i>level-5-expr</i> (R717)
S .NEQV. T	<i>level-5-expr</i> (R717)

A more complicated example of a level-5 expression is:

A .AND. B .EQV. .NOT. C

14 **7.1.1.7 General form of an expression**

15 Expressions are level-5 expressions optionally involving defined binary operators. Defined binary oper-
16 ators have the lowest operator precedence (Table 7.10).

17	R722	<i>expr</i>	is	[<i>expr defined-binary-op</i>] <i>level-5-expr</i>
18	R723	<i>defined-binary-op</i>	is	. <i>letter</i> [<i>letter</i>]

19 C704 (R723) A *defined-binary-op* shall not contain more than 63 letters and shall not be the same as
20 any *intrinsic-operator* or *logical-literal-constant*.

NOTE 7.7

Simple examples of an expression are:

<u>Example</u>	<u>Syntactic class</u>
A	<i>level-5-expr</i> (R717)
B.UNION.C	<i>expr</i> (R722)

More complicated examples of an expression are:

NOTE 7.7 (cont.)

```
(B .INTERSECT. C) .UNION. (X - Y)
A + B == C * D
.INVERSE. (A + B)
A + B .AND. C * D
E // G == H (1:10)
```

1 **7.1.2 Intrinsic operations**

2 An **intrinsic operation** is either an intrinsic unary operation or an intrinsic binary operation. An
 3 **intrinsic unary operation** is an operation of the form *intrinsic-operator* x_2 where x_2 is of an intrinsic
 4 type (4.4) listed in Table 7.1 for the unary intrinsic operator.

5 An **intrinsic binary operation** is an operation of the form x_1 *intrinsic-operator* x_2 where x_1 and
 6 x_2 are of the intrinsic types (4.4) listed in Table 7.1 for the binary intrinsic operator and are in shape
 7 conformance (7.1.5).

Table 7.1: Type of operands and results for intrinsic operators

Intrinsic operator <i>op</i>	Type of x_1	Type of x_2	Type of $[x_1] op x_2$
Unary +, -		I, R, Z	I, R, Z
Binary +, -, *, /, **	I	I, R, Z	I, R, Z
	R	I, R, Z	R, R, Z
	Z	I, R, Z	Z, Z, Z
//	C	C	C
	B	B	B
.EQ., .NE., ==, /=	I	I, R, Z, B	L, L, L, L
	R	I, R, Z, B	L, L, L, L
	Z	I, R, Z, B	L, L, L, L
	C	C	L
.GT., .GE., .LT., .LE. >, >=, <, <=	I	I, R	L, L
	R	I, R	L, L
	C	C	L
.NOT.		L, B	L, B
.AND., .OR., .EQV., .NEQV.	L	L	L
	B	B, I	B
	I	B	B
Note: The symbols I, R, Z, C, L, and B stand for the types integer, real, complex, character, logical, and bits, respectively. Where more than one type for x_2 is given, the type of the result of the operation is given in the same relative position in the next column. For the intrinsic operators with operands of type character, the kind type parameters of the operands shall be the same.			

8 A **numeric intrinsic operation** is an intrinsic operation for which the *intrinsic-operator* is a numeric
 9 operator (+, -, *, /, or **). A **numeric intrinsic operator** is the operator in a numeric intrinsic
 10 operation.

11 For numeric intrinsic binary operations, the two operands may be of different numeric types or different
 12 kind type parameters. Except for a value raised to an integer power, if the operands have different types
 13 or kind type parameters, the effect is as if each operand that differs in type or kind type parameter from
 14 those of the result is converted to the type and kind type parameter of the result before the operation
 15 is performed. When a value of type real or complex is raised to an integer power, the integer operand

1 need not be converted.

2 The **character intrinsic operation** is the intrinsic operation for which the *intrinsic-operator* is (//)
3 and both operands are of type character. The operands shall have the same kind type parameter. The
4 **character intrinsic operator** is the operator in a character intrinsic operation.

5 A **logical intrinsic operation** is an intrinsic operation for which the *intrinsic-operator* is .AND., .OR.,
6 .XOR., .NOT., .EQV., or .NEQV. and both operands are of type logical. A **logical intrinsic operator**
7 is the operator in a logical intrinsic operation.

8 A **bits intrinsic operation** is an intrinsic operation for which the *intrinsic-operator* is //, .AND., .OR.,
9 .XOR., .NOT., .EQV., or .NEQV. and at least one operand is of type bits. A **bits intrinsic operator**
10 is the operator in a bits intrinsic operation.

11 For bits intrinsic operations other than concatenation (//), the two operands may be of different types
12 or different kind type parameters. The effect is as if each operand that differs in type or kind type
13 parameter from those of the result is converted to the type and kind type parameter of the result before
14 the operation is performed.

15 A **relational intrinsic operator** is an *intrinsic-operator* that is .EQ., .NE., .GT., .GE., .LT., .LE.,
16 ==, /=, >, >=, <, or <=. A **relational intrinsic operation** is an intrinsic operation for which the
17 *intrinsic-operator* is a relational intrinsic operator. A **numeric relational intrinsic operation** is a
18 relational intrinsic operation for which both operands are of numeric type. A **character relational**
19 **intrinsic operation** is a relational intrinsic operation for which both operands are of type character.
20 The kind type parameters of the operands of a character relational intrinsic operation shall be the same.
21 A **bits relational intrinsic operation** is a relational intrinsic operation for which at least one of the
22 operands is of type bits.

23 If both operands of a bits relational operation do not have the same kind type parameter, the operand
24 with the smaller kind type parameter is converted to the same kind as the other operand. If one operand
25 of a bits relational operation is not of type bits, it is converted to type bits with the same kind type
26 parameter as the other operand. Any conversion takes place before the operation is evaluated.

27 7.1.3 Defined operations

28 A **defined operation** is either a defined unary operation or a defined binary operation. A **defined**
29 **unary operation** is an operation that has the form *defined-unary-op* x_2 or *intrinsic-operator* x_2 and
30 that is defined by a function and a generic interface (4.5.2, 12.4.3.3).

31 A function defines the unary operation *op* x_2 if

- 32 (1) the function is specified with a FUNCTION (12.6.2.1) or ENTRY (12.6.2.5) statement that
33 specifies one dummy argument d_2 ,
- 34 (2) either
 - 35 (a) a generic interface (12.4.3.2) provides the function with a *generic-spec* of OPERA-
36 TOR (*op*), or
 - 37 (b) there is a generic binding (4.5.2) in the declared type of x_2 with a *generic-spec* of
38 OPERATOR (*op*) and there is a corresponding binding to the function in the dynamic
39 type of x_2 ,
- 40 (3) the type of d_2 is compatible with the dynamic type of x_2 ,
- 41 (4) the type parameters, if any, of d_2 match the corresponding type parameters of x_2 , and
- 42 (5) either
 - 43 (a) the rank of x_2 matches that of d_2 or
 - 44 (b) the function is elemental and there is no other function that defines the operation.

1 If d_2 is an array, the shape of x_2 shall match the shape of d_2 .

2 A **defined binary operation** is an operation that has the form x_1 *defined-binary-op* x_2 or x_1 *intrinsic-*
3 *operator* x_2 and that is defined by a function and a generic interface.

4 A function defines the binary operation x_1 *op* x_2 if

- 5 (1) the function is specified with a FUNCTION (12.6.2.1) or ENTRY (12.6.2.5) statement that
6 specifies two dummy arguments, d_1 and d_2 ,
- 7 (2) either
 - 8 (a) a generic interface (12.4.3.2) provides the function with a *generic-spec* of OPERA-
9 TOR (*op*), or
 - 10 (b) there is a generic binding (4.5.2) in the declared type of x_1 or x_2 with a *generic-*
11 *spec* of OPERATOR (*op*) and there is a corresponding binding to the function in the
12 dynamic type of x_1 or x_2 , respectively,
- 13 (3) the types of d_1 and d_2 are compatible with the dynamic types of x_1 and x_2 , respectively,
- 14 (4) the type parameters, if any, of d_1 and d_2 match the corresponding type parameters of x_1
15 and x_2 , respectively, and
- 16 (5) either
 - 17 (a) the ranks of x_1 and x_2 match those of d_1 and d_2 or
 - 18 (b) the function is elemental, x_1 and x_2 are conformable, and there is no other function
19 that defines the operation.

20 If d_1 or d_2 is an array, the shapes of x_1 and x_2 shall match the shapes of d_1 and d_2 , respectively.

NOTE 7.8

An intrinsic operator may be used as the operator in a defined operation. In such a case, the generic properties of the operator are extended.

21 An **extension operation** is a defined operation in which the operator is of the form *defined-unary-op*
22 or *defined-binary-op*. Such an operator is called an **extension operator**. The operator used in an
23 extension operation may be such that a generic interface for the operator may specify more than one
24 function.

25 A **defined elemental operation** is a defined operation for which the function is elemental (12.8).

26 7.1.4 Type, type parameters, and shape of an expression

27 The type, type parameters, and shape of an expression depend on the operators and on the types, type
28 parameters, and shapes of the primaries used in the expression, and are determined recursively from
29 the syntactic form of the expression. The type of an expression is one of the intrinsic types (4.4) or a
30 derived type (4.5).

31 If an expression is a polymorphic primary or defined operation, the type parameters and the declared and
32 dynamic types of the expression are the same as those of the primary or defined operation. Otherwise
33 the type parameters and dynamic type of the expression are the same as its declared type and type
34 parameters; they are referred to simply as the type and type parameters of the expression.

35 R724 *logical-expr* **is** *expr*

36 C705 (R724) *logical-expr* shall be of type logical.

R725 *char-expr* **is** *expr*

1

2 C706 (R725) *char-expr* shall be of type character.3 R726 *default-char-expr* is *expr*4 C707 (R726) *default-char-expr* shall be of type default character.5 R727 *int-expr* is *expr*6 C708 (R727) *int-expr* shall be of type integer.7 R728 *numeric-expr* is *expr*8 C709 (R728) *numeric-expr* shall be of type integer, real, or complex.9 **7.1.4.1 Type, type parameters, and shape of a primary**

10 The type, type parameters, and shape of a primary are determined according to whether the primary is a
 11 constant, variable, array constructor, structure constructor, function reference, type parameter inquiry,
 12 type parameter name, or parenthesized expression. If a primary is a constant, its type, type parameters,
 13 and shape are those of the constant. If it is a structure constructor, it is scalar and its type and type
 14 parameters are as described in 4.5.10. If it is an array constructor, its type, type parameters, and shape
 15 are as described in 4.7. If it is a variable or function reference, its type, type parameters, and shape are
 16 those of the variable (5.2, 5.3) or the function reference (12.5.3), respectively. If the function reference
 17 is generic (12.4.3.2, 13.5) then its type, type parameters, and shape are those of the specific function
 18 referenced, which is determined by the types, type parameters, and ranks of its actual arguments as
 19 specified in 12.5.5.2. If it is a type parameter inquiry or type parameter name, it is a scalar integer with
 20 the kind of the type parameter.

21 If a primary is a parenthesized expression, its type, type parameters, and shape are those of the expres-
 22 sion.

23 The associated target object is referenced if a pointer appears as

- 24 (1) a primary in an intrinsic or defined operation,
- 25 (2) the *expr* of a parenthesized primary, or
- 26 (3) the only primary on the right-hand side of an intrinsic assignment statement.

27 The type, type parameters, and shape of the primary are those of the current target. If the pointer is
 28 not associated with a target, it may appear as a primary only as an actual argument in a reference to
 29 a procedure whose corresponding dummy argument is declared to be a pointer, or as the target in a
 30 pointer assignment statement.

31 A disassociated array pointer or an unallocated allocatable array has no shape but does have rank.
 32 The type, type parameters, and rank of the result of the NULL intrinsic function depend on context
 33 (13.7.131).

34 **7.1.4.2 Type, type parameters, and shape of the result of an operation**

35 The type of the result of an intrinsic operation $[x_1] \text{ op } x_2$ is specified by Table 7.1. The shape of the
 36 result of an intrinsic operation is the shape of x_2 if *op* is unary or if x_1 is scalar, and is the shape of x_1
 37 otherwise.

38 The type, type parameters, and shape of the result of a defined operation $[x_1] \text{ op } x_2$ are specified by the
 39 function defining the operation (7.2).

40 An expression of an intrinsic type has a kind type parameter. An expression of type character also has

1 a character length parameter.

2 The type parameters of the result of an intrinsic operation are as follows.

- 3 (1) For an expression $x_1 // x_2$ where $//$ is the character intrinsic operator and x_1 and x_2 are
4 of type character, the character length parameter is the sum of the lengths of the operands
5 and the kind type parameter is the kind type parameter of x_1 , which shall be the same as
6 the kind type parameter of x_2 .
- 7 (2) For an expression $op x_2$ where op is an intrinsic unary operator and x_2 is of type integer,
8 real, complex, logical, or bits, the kind type parameter of the expression is that of the
9 operand.
- 10 (3) For an expression $x_1 op x_2$ where op is a numeric intrinsic binary operator with one operand
11 of type integer and the other of type real or complex, the kind type parameter of the
12 expression is that of the real or complex operand.
- 13 (4) For an expression $x_1 op x_2$ where op is a numeric intrinsic binary operator with both
14 operands of the same type and kind type parameters, or with one real and one complex
15 with the same kind type parameters, the kind type parameter of the expression is identical
16 to that of each operand. In the case where both operands are integer with different kind type
17 parameters, the kind type parameter of the expression is that of the operand with the greater
18 decimal exponent range if the decimal exponent ranges are different; if the decimal exponent
19 ranges are the same, the kind type parameter of the expression is processor dependent, but
20 it is the same as that of one of the operands. In the case where both operands are any
21 of type real or complex with different kind type parameters, the kind type parameter of
22 the expression is that of the operand with the greater decimal precision if the decimal
23 precisions are different; if the decimal precisions are the same, the kind type parameter of
24 the expression is processor dependent, but it is the same as that of one of the operands.
- 25 (5) For an expression $x_1 op x_2$ where op is a logical intrinsic binary operator with both operands
26 of the same kind type parameter, the kind type parameter of the expression is identical to
27 that of each operand. In the case where both operands are of type logical with different
28 kind type parameters, the kind type parameter of the expression is processor dependent,
29 but it is the same as that of one of the operands.
- 30 (6) For an expression $x_1 // x_2$ where both operands are of type bits, the kind type parameter
31 of the result is the sum of the kind type parameters of the operands.
- 32 (7) For an expression $x_1 op x_2$ where op is a bits intrinsic binary operator other than $//$, the
33 type of the expression is bits and the kind type parameter of the expression is the maximum
34 of the kind type parameters of x_1 and x_2 .
- 35 (8) For an expression $x_1 op x_2$ where op is a relational intrinsic operator, the expression has
36 the default logical kind type parameter.

37 C710 The kind type parameter for the result of a bits concatenation operation expression shall be a
38 bits kind type parameter value supported by the processor.

39 7.1.5 Conformability rules for elemental operations

40 An **elemental operation** is an intrinsic operation or a defined elemental operation. Two entities are
41 in **shape conformance** if both are arrays of the same shape, or one or both are scalars.

42 For all elemental binary operations, the two operands shall be in shape conformance. In the case where
43 one is a scalar and the other an array, the scalar is treated as if it were an array of the same shape as
44 the array operand with every element, if any, of the array equal to the value of the scalar.

45 7.1.6 Specification expression

1 A **specification expression** is an expression with limitations that make it suitable for use in speci-
 2 fications such as length type parameters (C404) and array bounds (R513, R514). A *specification-expr*
 3 shall be an initialization expression unless it is in an interface body (12.4.3.2), the specification part of
 4 a subprogram, or the *declaration-type-spec* of a FUNCTION statement (12.6.2.1).

5 R729 *specification-expr* is *scalar-int-expr*

6 C711 (R729) The *scalar-int-expr* shall be a restricted expression.

7 A **restricted expression** is an expression in which each operation is intrinsic and each primary is

- 8 (1) a constant or subobject of a constant,
- 9 (2) an object designator with a base object that is a dummy argument that has neither the
 10 OPTIONAL nor the INTENT (OUT) attribute,
- 11 (3) an object designator with a base object that is in a common block,
- 12 (4) an object designator with a base object that is made accessible by use association or host
 13 association,
- 14 (5) an object designator with a base object that is a local variable of the procedure containing
 15 the BLOCK construct in which the restricted expression appears,
- 16 (6) an object designator with a base object that is a local variable of an outer BLOCK construct
 17 containing the BLOCK construct in which the restricted expression appears,
- 18 (7) an array constructor where each element and each *scalar-int-expr* of each *ac-implied-do-*
 19 *control* is a restricted expression,
- 20 (8) a structure constructor where each component is a restricted expression,
- 21 (9) a specification inquiry where each designator or function argument is
 - 22 (a) a restricted expression or
 - 23 (b) a variable whose properties inquired about are not
 - 24 (i) dependent on the upper bound of the last dimension of an assumed-size array,
 - 25 (ii) deferred, or
 - 26 (iii) defined by an expression that is not a restricted expression,
- 27 (10) a reference to any other standard intrinsic function where each argument is a restricted
 28 expression,
- 29 (11) a reference to a specification function where each argument is a restricted expression,
- 30 (12) a type parameter of the derived type being defined,
- 31 (13) an *ac-do-variable* within an array constructor where each *scalar-int-expr* of the correspond-
 32 ing *ac-implied-do-control* is a restricted expression, or
- 33 (14) a restricted expression enclosed in parentheses,

34 where each subscript, section subscript, substring starting point, substring ending point, and type pa-
 35 rameter value is a restricted expression, and where any final subroutine that is invoked is pure.

36 A **specification inquiry** is a reference to

- 37 (1) an array inquiry function (13.5.7),
- 38 (2) the bit inquiry function BIT_SIZE,
- 39 (3) the character inquiry function LEN,
- 40 (4) the kind inquiry function KIND,
- 41 (5) the bits kind inquiry function BITS_KIND,
- 42 (6) the character inquiry function NEW_LINE,
- 43 (7) a numeric inquiry function (13.5.6),
- 44 (8) a type parameter inquiry (6.1.4),

- 1 (9) an IEEE inquiry function (14.9.1),
 2 (10) the function C_SIZEOF from the intrinsic module ISO_C_BINDING (15.2.3.6) , or
 3 (11) the COMPILER_VERSION or COMPILER_OPTIONS inquiry functions from the intrinsic
 4 module ISO_FORTRAN_ENV (13.8.3.3, 13.8.3.4).

5 A function is a **specification function** if it is a pure function, is not a standard intrinsic function, is
 6 not an internal function, is not a statement function, and does not have a dummy procedure argument.

7 Evaluation of a specification expression shall not directly or indirectly cause a procedure defined by the
 8 subprogram in which it appears to be invoked.

NOTE 7.9

Specification functions are nonintrinsic functions that may be used in specification expressions to determine the attributes of data objects. The requirement that they be pure ensures that they cannot have side effects that could affect other objects being declared in the same *specification-part*. The requirement that they not be internal ensures that they cannot inquire, via host association, about other objects being declared in the same *specification-part*. The prohibition against recursion avoids the creation of a new instance of a procedure while construction of one is in progress.

9 A variable in a specification expression shall have its type and type parameters, if any, specified by a
 10 previous declaration in the same scoping unit, by the implicit typing rules in effect for the scoping unit,
 11 or by host or use association. If a variable in a specification expression is typed by the implicit typing
 12 rules, its appearance in any subsequent type declaration statement shall confirm the implied type and
 13 type parameters.

14 If a specification expression includes a specification inquiry that depends on a type parameter or an
 15 array bound of an entity specified in the same *specification-part*, the type parameter or array bound
 16 shall be specified in a prior specification of the *specification-part*. The prior specification may be to the
 17 left of the specification inquiry in the same statement, but shall not be within the same *entity-decl*. If a
 18 specification expression includes a reference to the value of an element of an array specified in the same
 19 *specification-part*, the array shall be completely specified in prior declarations.

20 If a specification expression in a module or submodule includes a reference to a generic entity, that
 21 generic entity shall have no specific procedures defined in the module or submodule subsequent to the
 22 specification expression.

NOTE 7.10

The following are examples of specification expressions:

```
LBOUND (B, 1) + 5 ! B is an assumed-shape dummy array
M + LEN (C)      ! M and C are dummy arguments
2 * PRECISION (A) ! A is a real variable made accessible
                  ! by a USE statement
```

23 7.1.7 Initialization expression

24 An **initialization expression** is an expression with limitations that make it suitable for use as a kind
 25 type parameter, initializer, or named constant. It is an expression in which each operation is intrinsic,
 26 and each primary is

- 27 (1) a constant or subobject of a constant,
 28 (2) an array constructor where each element and each *scalar-int-expr* of each *ac-implied-do-control*
 29 is an initialization expression,

- 1 (3) a structure constructor where each *component-spec* corresponding to
 2 (a) an allocatable component is a reference to the intrinsic function NULL,
 3 (b) a pointer component is a reference to the intrinsic function NULL or an initialization
 4 target, and
 5 (c) any other component is an initialization expression,
 6 (4) a reference to an elemental standard intrinsic function, where each argument is an initial-
 7 ization expression,
 8 (5) a reference to a transformational standard intrinsic function other than NULL, where each
 9 argument is an initialization expression,
 10 (6) A reference to the transformational intrinsic function NULL that does not have an argu-
 11 ment with a type parameter that is assumed or is defined by an expression that is not an
 12 initialization expression,
 13 (7) a reference to the transformational function IEEE_SELECTED_REAL_KIND from the in-
 14 trinsic module IEEE_ARITHMETIC (14), where each argument is an initialization expres-
 15 sion.
 16 (8) a specification inquiry where each designator or function argument is
 17 (a) an initialization expression or
 18 (b) a variable whose properties inquired about are not
 19 (i) assumed,
 20 (ii) deferred, or
 21 (iii) defined by an expression that is not an initialization expression,
 22 (9) a kind type parameter of the derived type being defined,
 23 (10) a *data-i-do-variable* within a *data-implied-do*,
 24 (11) an *ac-do-variable* within an array constructor where each *scalar-int-expr* of the correspond-
 25 ing *ac-implied-do-control* is an initialization expression, or
 26 (12) an initialization expression enclosed in parentheses,

27 and where each subscript, section subscript, substring starting point, substring ending point, and type
 28 parameter value is an initialization expression.

29 R730 *initialization-expr* **is** *expr*

30 C712 (R730) *initialization-expr* shall be an initialization expression.

31 R731 *char-initialization-expr* **is** *char-expr*

32 C713 (R731) *char-initialization-expr* shall be an initialization expression.

33 R732 *int-initialization-expr* **is** *int-expr*

34 C714 (R732) *int-initialization-expr* shall be an initialization expression.

35 R733 *logical-initialization-expr* **is** *logical-expr*

36 C715 (R733) *logical-initialization-expr* shall be an initialization expression.

37 If an initialization expression includes a specification inquiry that depends on a type parameter or an
 38 array bound of an entity specified in the same *specification-part*, the type parameter or array bound
 39 shall be specified in a prior specification of the *specification-part*. The prior specification may be to the
 40 left of the specification inquiry in the same statement, but shall not be within the same *entity-decl*.

41 If an initialization expression in a module or submodule includes a reference to a generic entity, that
 42 generic entity shall have no specific procedures defined in the module or submodule subsequent to the

1 initialization expression.

NOTE 7.11

The following are examples of initialization expressions:

```

3
-3 + 4
'AB'
'AB' // 'CD'
('AB' // 'CD') // 'EF'
SIZE (A)
DIGITS (X) + 4
4.0 * atan(1.0)
ceiling(number_of_decimal_digits / log10(radix(0.0)))

```

where A is an explicit-shaped array with constant bounds and X is of type default real.

2 7.1.8 Evaluation

3 7.1.8.1 Evaluation of operations

4 An intrinsic operation requires the values of its operands.

5 The evaluation of a function reference shall neither affect nor be affected by the evaluation of any other
6 entity within the statement. If a function reference causes definition or undefinition of an actual argument
7 of the function, that argument or any associated entities shall not appear elsewhere in the same statement.
8 However, execution of a function reference in the logical expression in an IF statement (8.1.8.4), the mask
9 expression in a WHERE statement (7.4.3.1), or the subscripts and strides in a FORALL statement (7.4.4)
10 is permitted to define variables in the statement that is conditionally executed.

NOTE 7.12

For example, the statements

```

A (I) = F (I)
Y = G (X) + X

```

are prohibited if the reference to F defines or undefines I or the reference to G defines or undefines X.

However, in the statements

```

IF (F (X)) A = X
WHERE (G (X)) B = X

```

F or G may define X.

11 The declared type of an expression in which a function reference appears does not affect, and is not
12 affected by, the evaluation of the actual arguments of the function.

13 Execution of an array element reference requires the evaluation of its subscripts. The type of an expres-
14 sion in which the array element reference appears does not affect, and is not affected by, the evaluation
15 of its subscripts. Execution of an array section reference requires the evaluation of its section subscripts.
16 The type of an expression in which an array section appears does not affect, and is not affected by, the
17 evaluation of the array section subscripts. Execution of a substring reference requires the evaluation of

- 1 its substring expressions. The type of an expression in which a substring appears does not affect, and
 2 is not affected by, the evaluation of the substring expressions. It is not necessary for a processor to
 3 evaluate any subscript expressions or substring expressions for an array of zero size or a character entity
 4 of zero length.
- 5 The appearance of an array constructor requires the evaluation of each *scalar-int-expr* of the *ac-implicit-*
 6 *do-control* in any *ac-implicit-do* it may contain. The type of an expression in which an array constructor
 7 appears does not affect, and is not affected by, the evaluation of such bounds and stride expressions.
- 8 When an elemental binary operation is applied to a scalar and an array or to two arrays of the same
 9 shape, the operation is performed element-by-element on corresponding array elements of the array
 10 operands.

NOTE 7.13

For example, the array expression

$$A + B$$

produces an array of the same shape as A and B. The individual array elements of the result have the values of the first element of A added to the first element of B, the second element of A added to the second element of B, etc.

- 11 When an elemental unary operator operates on an array operand, the operation is performed element-
 12 by-element, and the result is the same shape as the operand.

NOTE 7.14

If an elemental operation is intrinsically pure or is implemented by a pure elemental function (12.8), the element operations may be performed simultaneously or in any order.

13 **7.1.8.2 Evaluation of operands**

- 14 It is not necessary for a processor to evaluate all of the operands of an expression, or to evaluate entirely
 15 each operand, if the value of the expression can be determined otherwise.

NOTE 7.15

This principle is most often applicable to logical expressions, zero-sized arrays, and zero-length strings, but it applies to all expressions.

For example, in evaluating the expression

$$X > Y \text{ .OR. } L(Z)$$

where X, Y, and Z are real and L is a function of type logical, the function reference L (Z) need not be evaluated if X is greater than Y. Similarly, in the array expression

$$W(Z) + A$$

where A is of size zero and W is a function, the function reference W (Z) need not be evaluated.

- 16 If a statement contains a function reference in a part of an expression that need not be evaluated, all
 17 entities that would have become defined in the execution of that reference become undefined at the
 18 completion of evaluation of the expression containing the function reference.

NOTE 7.16

In the examples in Note 7.15, if L or W defines its argument, evaluation of the expressions under the specified conditions causes Z to become undefined, no matter whether or not L(Z) or W(Z) is evaluated.

- 1 If a statement contains a function reference in a part of an expression that need not be evaluated, no
 2 invocation of that function in that part of the expression shall execute an image control statement other
 3 than CRITICAL or END CRITICAL.

NOTE 7.17

This restriction is intended to avoid inadvertant deadlock caused by optimization.

4 **7.1.8.3 Integrity of parentheses**

- 5 Subclauses 7.1.8.3 to 7.1.8.8 state certain conditions under which a processor may evaluate an expression
 6 that is different from the one specified by applying the rules given in 7.1.1 and 7.2. However, any
 7 expression in parentheses shall be treated as a data entity.

NOTE 7.18

For example, in evaluating the expression $A + (B - C)$ where A, B, and C are of numeric types, the difference of B and C shall be evaluated before the addition operation is performed; the processor shall not evaluate the mathematically equivalent expression $(A + B) - C$.

8 **7.1.8.4 Evaluation of numeric intrinsic operations**

- 9 The rules given in 7.2.2 specify the interpretation of a numeric intrinsic operation. Once the interpreta-
 10 tion has been established in accordance with those rules, the processor may evaluate any mathematically
 11 equivalent expression, provided that the integrity of parentheses is not violated.
- 12 Two expressions of a numeric type are mathematically equivalent if, for all possible values of their
 13 primaries, their mathematical values are equal. However, mathematically equivalent expressions of
 14 numeric type may produce different computational results.

NOTE 7.19

Any difference between the values of the expressions $(1./3.)*3.$ and 1. is a computational difference, not a mathematical difference. The difference between the values of the expressions $5/2$ and $5./2.$ is a mathematical difference, not a computational difference.

The mathematical definition of integer division is given in 7.2.2.1.

NOTE 7.20

The following are examples of expressions with allowable alternative forms that may be used by the processor in the evaluation of those expressions. A, B, and C represent arbitrary real or complex operands; I and J represent arbitrary integer operands; and X, Y, and Z represent arbitrary operands of numeric type.

<u>Expression</u>	<u>Allowable alternative form</u>
$X + Y$	$Y + X$
$X * Y$	$Y * X$
$-X + Y$	$Y - X$
$X + Y + Z$	$X + (Y + Z)$
$X - Y + Z$	$X - (Y - Z)$

NOTE 7.20 (cont.)

$X * A / Z$	$X * (A / Z)$
$X * Y - X * Z$	$X * (Y - Z)$
$A / B / C$	$A / (B * C)$
$A / 5.0$	$0.2 * A$

The following are examples of expressions with forbidden alternative forms that shall not be used by a processor in the evaluation of those expressions.

<u>Expression</u>	<u>Forbidden alternative form</u>
$I / 2$	$0.5 * I$
$X * I / J$	$X * (I / J)$
$I / J / A$	$I / (J * A)$
$(X + Y) + Z$	$X + (Y + Z)$
$(X * Y) - (X * Z)$	$X * (Y - Z)$
$X * (Y - Z)$	$X * Y - X * Z$

- 1 The execution of any numeric operation whose result is not defined by the arithmetic used by the
- 2 processor is prohibited. Raising a negative-valued primary of type real to a real power is prohibited.
- 3 In addition to the parentheses required to establish the desired interpretation, parentheses may be
- 4 included to restrict the alternative forms that may be used by the processor in the actual evaluation
- 5 of the expression. This is useful for controlling the magnitude and accuracy of intermediate values
- 6 developed during the evaluation of an expression.

NOTE 7.21

For example, in the expression

$$A + (B - C)$$

the parenthesized expression $(B - C)$ shall be evaluated and then added to A .

The inclusion of parentheses may change the mathematical value of an expression. For example, the two expressions

$$A * I / J$$

$$A * (I / J)$$

may have different mathematical values if I and J are of type integer.

- 7 Each operand in a numeric intrinsic operation has a type that may depend on the order of evaluation
- 8 used by the processor.

NOTE 7.22

For example, in the evaluation of the expression

$$Z + R + I$$

where Z , R , and I represent data objects of complex, real, and integer type, respectively, the type of the operand that is added to I may be either complex or real, depending on which pair of operands (Z and R , R and I , or Z and I) is added first.

1 7.1.8.5 Evaluation of the character intrinsic operation

2 The rules given in 7.2.3 specify the interpretation of the character intrinsic operation. A processor is
 3 only required to evaluate as much of the character intrinsic operation as is required by the context in
 4 which the expression appears.

NOTE 7.23

For example, the statements

```
CHARACTER (LEN = 2) C1, C2, C3, CF
C1 = C2 // CF (C3)
```

do not require the function CF to be evaluated, because only the value of C2 is needed to determine the value of C1 because C1 and C2 both have a length of 2.

5 7.1.8.6 Evaluation of relational intrinsic operations

6 The rules given in 7.2.4 specify the interpretation of relational intrinsic operations. Once the interpre-
 7 tation of an expression has been established in accordance with those rules, the processor may evaluate
 8 any other expression that is relationally equivalent, provided that the integrity of parentheses in any
 9 expression is not violated.

NOTE 7.24

For example, the processor may choose to evaluate the expression

```
I > J
```

where I and J are integer variables, as

```
J - I < 0
```

10 Two relational intrinsic operations are relationally equivalent if their logical values are equal for all
 11 possible values of their primaries.

12 7.1.8.7 Evaluation of logical intrinsic operations

13 The rules given in 7.2.5 specify the interpretation of logical intrinsic operations. Once the interpretation
 14 of an expression has been established in accordance with those rules, the processor may evaluate any
 15 other expression that is logically equivalent, provided that the integrity of parentheses in any expression
 16 is not violated.

NOTE 7.25

For example, for the variables L1, L2, and L3 of type logical, the processor may choose to evaluate the expression

```
L1 .AND. L2 .AND. L3
```

as

```
L1 .AND. (L2 .AND. L3)
```

17 Two expressions of type logical are logically equivalent if their values are equal for all possible values of
 18 their primaries.

1 7.1.8.8 Evaluation of bits intrinsic operations

2 The rules given in 7.2.6 specify the interpretation of bits intrinsic operations. Once the interpretation
 3 of an expression has been established in accordance with those rules, the processor may evaluate any
 4 other expression that is computationally equivalent, provided that the integrity of parentheses in any
 5 expression is not violated.

NOTE 7.26

For example, for the variables B1, B2, and B3 of type bits, the processor may choose to evaluate the expression

$$B1 \text{ .XOR. } B2 \text{ .XOR. } B3$$

as

$$B1 \text{ .XOR. } (B2 \text{ .XOR. } B3)$$

6 Two expressions of type bits are computationally equivalent if their values are equal for all possible
 7 values of their primaries.

8 7.1.8.9 Evaluation of a defined operation

9 The rules given in 7.2 specify the interpretation of a defined operation. Once the interpretation of an
 10 expression has been established in accordance with those rules, the processor may evaluate any other
 11 expression that is equivalent, provided that the integrity of parentheses is not violated.

12 Two expressions of derived type are equivalent if their values are equal for all possible values of their
 13 primaries.

14 7.2 Interpretation of operations

15 7.2.1 General

16 The intrinsic operations are those defined in 7.1.2. These operations are divided into the following
 17 categories: numeric, character, relational, logical, and bits. The interpretations defined in subclause 7.2
 18 apply to both scalars and arrays; the interpretation for arrays is obtained by applying the interpretation
 19 for scalars element by element.

20 The interpretation of a defined operation is provided by the function that defines the operation. The type,
 21 type parameters and interpretation of an expression that consists of an intrinsic or defined operation are
 22 independent of the type and type parameters of the context or any larger expression in which it appears.

NOTE 7.27

For example, if X is of type real, J is of type integer, and INT is the real-to-integer intrinsic conversion function, the expression INT (X + J) is an integer expression and X + J is a real expression.

23 The operators <, <=, >, >=, ==, and /= always have the same interpretations as the operators .LT.,
 24 .LE., .GT., .GE., .EQ., and .NE., respectively.

25 7.2.2 Numeric intrinsic operations

1 A numeric operation is used to express a numeric computation. Evaluation of a numeric operation
 2 produces a numeric value. The permitted data types for operands of the numeric intrinsic operations
 3 are specified in 7.1.2.

4 The numeric operators and their interpretation in an expression are given in Table 7.3, where x_1 denotes
 5 the operand to the left of the operator and x_2 denotes the operand to the right of the operator.

Table 7.3: **Interpretation of the numeric intrinsic operators**

Operator	Representing	Use of operator	Interpretation
**	Exponentiation	$x_1 ** x_2$	Raise x_1 to the power x_2
/	Division	x_1 / x_2	Divide x_1 by x_2
*	Multiplication	$x_1 * x_2$	Multiply x_1 by x_2
-	Subtraction	$x_1 - x_2$	Subtract x_2 from x_1
-	Negation	$- x_2$	Negate x_2
+	Addition	$x_1 + x_2$	Add x_1 and x_2
+	Identity	$+ x_2$	Same as x_2

6 The interpretation of a division operation depends on the types of the operands (7.2.2.1).

7 If x_1 and x_2 are of type integer and x_2 has a negative value, the interpretation of $x_1 ** x_2$ is the same
 8 as the interpretation of $1/(x_1 ** \text{ABS}(x_2))$, which is subject to the rules of integer division (7.2.2.1).

NOTE 7.28

For example, $2 ** (-3)$ has the value of $1/(2 ** 3)$, which is zero.

9 **7.2.2.1 Integer division**

10 One operand of type integer may be divided by another operand of type integer. Although the math-
 11 ematical quotient of two integers is not necessarily an integer, Table 7.1 specifies that an expression
 12 involving the division operator with two operands of type integer is interpreted as an expression of type
 13 integer. The result of such an operation is the integer closest to the mathematical quotient and between
 14 zero and the mathematical quotient inclusively.

NOTE 7.29

For example, the expression $(-8) / 3$ has the value (-2) .

15 **7.2.2.2 Complex exponentiation**

16 In the case of a complex value raised to a complex power, the value of the operation $x_1 ** x_2$ is the
 17 principal value of $x_1^{x_2}$.

18 **7.2.3 Character intrinsic operation**

19 The character intrinsic operator `//` is used to concatenate two operands of type character with the same
 20 kind type parameter. Evaluation of the character intrinsic operation produces a result of type character.

21 The interpretation of the character intrinsic operator `//` when used to form an expression is given in
 22 Table 7.4, where x_1 denotes the operand to the left of the operator and x_2 denotes the operand to the
 23 right of the operator.

Table 7.4: **Interpretation of the character intrinsic operator `//`**

Operator	Representing	Use of operator	Interpretation
//	Concatenation	$x_1 // x_2$	Concatenate x_1 with x_2

- 1 The result of the character intrinsic operation `//` is a character string whose value is the value of x_1
 2 concatenated on the right with the value of x_2 and whose length is the sum of the lengths of x_1 and x_2 .
 3 Parentheses used to specify the order of evaluation have no effect on the value of a character expression.

NOTE 7.30

For example, the value of `('AB' // 'CDE') // 'F'` is the string 'ABCDEF'. Also, the value of `'AB' // ('CDE' // 'F')` is the string 'ABCDEF'.

4 **7.2.4 Relational intrinsic operations**

- 5 A relational intrinsic operation is used to compare values of two operands using the relational intrinsic
 6 operators `.LT.`, `.LE.`, `.GT.`, `.GE.`, `.EQ.`, `.NE.`, `<`, `<=`, `>`, `>=`, `==`, and `/=`. The permitted types for
 7 operands of the relational intrinsic operators are specified in 7.1.2.

NOTE 7.31

As shown in Table 7.1, a relational intrinsic operator cannot be used to compare the value of an expression of a numeric type with one of type character or logical. Also, two operands of type logical cannot be compared, a complex operand may be compared with another numeric operand only when the operator is `.EQ.`, `.NE.`, `==`, or `/=`, and two character operands cannot be compared unless they have the same kind type parameter value.

- 8 Evaluation of a relational intrinsic operation produces a result of type default logical.
 9 The interpretation of the relational intrinsic operators is given in Table 7.5, where x_1 denotes the operand
 10 to the left of the operator and x_2 denotes the operand to the right of the operator.

Table 7.5: **Interpretation of the relational intrinsic operators**

Operator	Representing	Use of operator	Interpretation
<code>.LT.</code>	Less than	x_1 <code>.LT.</code> x_2	x_1 less than x_2
<code><</code>	Less than	$x_1 < x_2$	x_1 less than x_2
<code>.LE.</code>	Less than or equal to	x_1 <code>.LE.</code> x_2	x_1 less than or equal to x_2
<code><=</code>	Less than or equal to	$x_1 <= x_2$	x_1 less than or equal to x_2
<code>.GT.</code>	Greater than	x_1 <code>.GT.</code> x_2	x_1 greater than x_2
<code>></code>	Greater than	$x_1 > x_2$	x_1 greater than x_2
<code>.GE.</code>	Greater than or equal to	x_1 <code>.GE.</code> x_2	x_1 greater than or equal to x_2
<code>>=</code>	Greater than or equal to	$x_1 >= x_2$	x_1 greater than or equal to x_2
<code>.EQ.</code>	Equal to	x_1 <code>.EQ.</code> x_2	x_1 equal to x_2
<code>==</code>	Equal to	$x_1 == x_2$	x_1 equal to x_2
<code>.NE.</code>	Not equal to	x_1 <code>.NE.</code> x_2	x_1 not equal to x_2
<code>/=</code>	Not equal to	$x_1 /= x_2$	x_1 not equal to x_2

- 11 A numeric relational intrinsic operation is interpreted as having the logical value true if and only if the
 12 values of the operands satisfy the relation specified by the operator.

- 13 In the numeric relational operation

$$14 \quad x_1 \text{ rel-op } x_2$$

- 15 if the types or kind type parameters of x_1 and x_2 differ, their values are converted to the type and kind
 16 type parameter of the expression $x_1 + x_2$ before evaluation.

- 17 A character relational intrinsic operation is interpreted as having the logical value true if and only if the
 18 values of the operands satisfy the relation specified by the operator.

1 For a character relational intrinsic operation, the operands are compared one character at a time in
 2 order, beginning with the first character of each character operand. If the operands are of unequal
 3 length, the shorter operand is treated as if it were extended on the right with blanks to the length of
 4 the longer operand. If both x_1 and x_2 are of zero length, x_1 is equal to x_2 ; if every character of x_1 is
 5 the same as the character in the corresponding position in x_2 , x_1 is equal to x_2 . Otherwise, at the first
 6 position where the character operands differ, the character operand x_1 is considered to be less than x_2
 7 if the character value of x_1 at this position precedes the value of x_2 in the collating sequence (4.4.5.4);
 8 x_1 is greater than x_2 if the character value of x_1 at this position follows the value of x_2 in the collating
 9 sequence.

NOTE 7.32

The collating sequence depends partially on the processor; however, the result of the use of the operators .EQ., .NE., ==, and /= does not depend on the collating sequence.

For nondefault character types, the blank padding character is processor dependent.

10 A bits relational intrinsic operation is interpreted as having the logical value true if and only if the values
 11 of the operands satisfy the relation specified by the operator.

12 For a bits relational intrinsic operation, x_1 and x_2 are equal if and only if each corresponding bit has
 13 the same value. If x_1 and x_2 are not equal, and the leftmost unequal corresponding bit of x_1 is 1 and
 14 x_2 is 0 then x_1 is greater than x_2 ; otherwise x_1 is less than x_2 .

15 **7.2.5 Logical intrinsic operations**

16 A logical operation is used to express a logical computation. Evaluation of a logical operation produces
 17 a result of type logical. The permitted types for operands of the logical intrinsic operations are specified
 18 in 7.1.2.

19 The logical operators and their interpretation when used to form an expression are given in Table 7.6,
 20 where x_1 denotes the operand to the left of the operator and x_2 denotes the operand to the right of the
 21 operator.

Table 7.6: Interpretation of the logical intrinsic operators

Operator	Representing	Use of operator	Interpretation
.NOT.	Logical negation	.NOT. x_2	True if x_2 is false
.AND.	Logical conjunction	x_1 .AND. x_2	True if x_1 and x_2 are both true
.OR.	Logical inclusive disjunction	x_1 .OR. x_2	True if x_1 and/or x_2 is true
.EQV.	Logical equivalence	x_1 .EQV. x_2	True if both x_1 and x_2 are true or both are false
.NEQV.	Logical nonequivalence	x_1 .NEQV. x_2	True if either x_1 or x_2 is true, but not both
.XOR.	Logical nonequivalence	x_1 .XOR. x_2	True if either x_1 or x_2 is true, but not both

22 The values of the logical intrinsic operations are shown in Table 7.7.

Table 7.7: The values of operations involving logical intrinsic operators

x_1	x_2	.NOT. x_2	x_1 .AND. x_2	x_1 .OR. x_2	x_1 .EQV. x_2	x_1 .NEQV. x_2	x_1 .XOR. x_2
true	true	false	true	true	true	false	false
true	false	true	false	true	false	true	true
false	true	false	false	true	false	true	true
false	false	true	false	false	true	false	false

1 **7.2.6 Bits intrinsic operations**

2 Bit operations are used to express bitwise operations on sequences of bits, or to concatenate such
 3 sequences. Evaluation of a bits operation produces a result of type bits. The permitted types of
 4 operands of the bits intrinsic operations are specified in 7.1.2.

5 The bits operators and their interpretation when used to form an expression are given in Table 7.8,
 6 where x_1 denotes the operand of type bits to the left of the operator and x_2 denotes the operand of type
 7 bits to the right of the operator.

Table 7.8: Interpretation of the bits intrinsic operators

Operator	Representing	Use of operator	Interpretation
<code>(//)</code>	Concatenation	$x_1 // x_2$	Concatenation of x_1 and x_2
<code>.NOT.</code>	Bitwise NOT	<code>.NOT.</code> x_2	Bitwise NOT of x_2
<code>.AND.</code>	Bitwise AND	x_1 <code>.AND.</code> x_2	Bitwise AND of x_1 and x_2
<code>.OR.</code>	Bitwise inclusive OR	x_1 <code>.OR.</code> x_2	Bitwise OR of x_1 and x_2
<code>.EQV.</code>	Bitwise equivalence	x_1 <code>.EQV.</code> x_2	Bitwise equivalence of x_1 and x_2
<code>.NEQV.</code>	Bitwise nonequivalence	x_1 <code>.NEQV.</code> x_2	Bitwise nonequivalence of x_1 and x_2
<code>.XOR.</code>	Bitwise exclusive OR	x_1 <code>.XOR.</code> x_2	Bitwise exclusive OR of x_1 and x_2

8 The leftmost $KIND(x_1)$ bits of the result of the bits concatenation operation are the value of x_1 and the
 9 rightmost $KIND(x_2)$ bits of the result are the value of x_2 .

10 For a bits intrinsic operation other than `//`, the result value is computed separately for each pair of bits
 11 at corresponding positions in each operand. The value of each bit operation, for bits denoted b_1 and b_2
 12 are given in Table 7.9.

Table 7.9: The values of bits intrinsic operations other than `//`

x_1	x_2	<code>.NOT.</code> x_2	x_1 <code>.AND.</code> x_2	x_1 <code>.OR.</code> x_2	x_1 <code>.EQV.</code> x_2	x_1 <code>.NEQV.</code> x_2	x_1 <code>.XOR.</code> x_2
1	1	0	1	1	1	0	0
1	0	1	0	1	0	1	1
0	1	0	0	1	0	1	1
0	0	1	0	0	1	0	0

13 **7.3 Precedence of operators**

14 There is a precedence among the intrinsic and extension operations corresponding to the form of expres-
 15 sions specified in 7.1.1, which determines the order in which the operands are combined unless the order
 16 is changed by the use of parentheses. This precedence order is summarized in Table 7.10.

Table 7.10: Categories of operations and relative precedence

Category of operation	Operators	Precedence
Extension	<i>defined-unary-op</i>	Highest
Numeric	<code>**</code>	.
Numeric	<code>*</code> , <code>/</code>	.
Numeric	unary <code>+</code> , <code>-</code>	.
Numeric	binary <code>+</code> , <code>-</code>	.
Character	<code>//</code>	.
Relational	<code>.EQ.</code> , <code>.NE.</code> , <code>.LT.</code> , <code>.LE.</code> , <code>.GT.</code> , <code>.GE.</code> , <code>==</code> , <code>/=</code> , <code><</code> , <code><=</code> , <code>></code> , <code>>=</code>	.
Logical, Bits	<code>.NOT.</code>	.
Logical, Bits	<code>.AND.</code>	.

Categories of operations and relative precedence (cont.)

Category of operation	Operators	Precedence
Logical, Bits	.OR.	.
Logical, Bits	.EQV., .NEQV., .XOR.	.
Extension	<i>defined-binary-op</i>	Lowest

- 1 The precedence of a defined operation is that of its operator.

NOTE 7.33

For example, in the expression

```
-A ** 2
```

the exponentiation operator (**) has precedence over the negation operator (-); therefore, the operands of the exponentiation operator are combined to form an expression that is used as the operand of the negation operator. The interpretation of the above expression is the same as the interpretation of the expression

```
- (A ** 2)
```

- 2 The general form of an expression (7.1.1) also establishes a precedence among operators in the same
 3 syntactic class. This precedence determines the order in which the operands are to be combined in
 4 determining the interpretation of the expression unless the order is changed by the use of parentheses.

NOTE 7.34

In interpreting a *level-2-expr* containing two or more binary operators + or -, each operand (*add-operand*) is combined from left to right. Similarly, the same left-to-right interpretation for a *mult-operand* in *add-operand*, as well as for other kinds of expressions, is a consequence of the general form. However, for interpreting a *mult-operand* expression when two or more exponentiation operators ** combine *level-1-expr* operands, each *level-1-expr* is combined from right to left.

For example, the expressions

```
2.1 + 3.4 + 4.9
2.1 * 3.4 * 4.9
2.1 / 3.4 / 4.9
2 ** 3 ** 4
'AB' // 'CD' // 'EF'
```

have the same interpretations as the expressions

```
(2.1 + 3.4) + 4.9
(2.1 * 3.4) * 4.9
(2.1 / 3.4) / 4.9
2 ** (3 ** 4)
('AB' // 'CD') // 'EF'
```

As a consequence of the general form (7.1.1), only the first *add-operand* of a *level-2-expr* may be preceded by the identity (+) or negation (-) operator. These formation rules do not permit expressions containing two consecutive numeric operators, such as A ** -B or A + -B. However, expressions such as A ** (-B) and A + (-B) are permitted. The rules do allow a binary operator or an intrinsic unary operator to be followed by a defined unary operator, such as:

NOTE 7.34 (cont.)

```
A * .INVERSE. B
- .INVERSE. (B)
```

As another example, in the expression

```
A .OR. B .AND. C
```

the general form implies a higher precedence for the .AND. operator than for the .OR. operator; therefore, the interpretation of the above expression is the same as the interpretation of the expression

```
A .OR. (B .AND. C)
```

NOTE 7.35

An expression may contain more than one category of operator. The logical expression

```
L .OR. A + B >= C
```

where A, B, and C are of type real, and L is of type logical, contains a numeric operator, a relational operator, and a logical operator. This expression would be interpreted the same as the expression

```
L .OR. ((A + B) >= C)
```

NOTE 7.36

If

- (1) The operator ** is extended to type logical,
- (2) The operator .STARSTAR. is defined to duplicate the function of ** on type real,
- (3) .MINUS. is defined to duplicate the unary operator −, and
- (4) L1 and L2 are type logical and X and Y are type real,

then in precedence: L1 ** L2 is higher than X * Y; X * Y is higher than X .STARSTAR. Y; and .MINUS. X is higher than −X.

1 7.4 Assignment

2 Execution of an assignment statement causes a variable to become defined or redefined. Execution of a
 3 pointer assignment statement causes a pointer to become associated with a target or causes its pointer
 4 association status to become disassociated or undefined. Execution of a WHERE statement or WHERE
 5 construct masks the evaluation of expressions and assignment of values in array assignment statements
 6 according to the value of a logical array expression. Execution of a FORALL statement or FORALL
 7 construct controls the assignment to elements of arrays by using a set of index variables and a mask
 8 expression.

9 7.4.1 Assignment statement

10 A variable may be defined or redefined by execution of an assignment statement.

11 7.4.1.1 General form

- 1 R734 *assignment-stmt* **is** *variable = expr*
 2 C716 (R734) The *variable* shall not be a whole assumed-size array.

NOTE 7.37

Examples of an assignment statement are:

```
A = 3.5 + X * Y
I = INT (A)
```

- 3 An *assignment-stmt* shall meet the requirements of either a defined assignment statement or an intrinsic
 4 assignment statement.

7.4.1.2 Intrinsic assignment statement

- 6 An **intrinsic assignment statement** is an assignment statement that is not a defined assignment
 7 statement (7.4.1.4). In an intrinsic assignment statement,

- 8 (1) if the variable is polymorphic it shall be allocatable,
 9 (2) if *variable* is a co-indexed object, it shall not be of a type that has an allocatable ultimate
 10 component,

J3 internal note**Unresolved Technical Issue 023**

06-174r3 had no edit about this at all; when it came up in discussion with JKR et al they suggested making a runtime requirement that the types and shapes of the allocatable components must match; however, that would be very unsafe. Since the only safe way of programming such an assignment would be to avoid the intrinsic assignment altogether, we should not allow the intrinsic assignment in this case. Furthermore, giving different semantics to intrinsic assignment when a component is a co-indexed object is a bad idea - it will confuse the users - better to disallow the discrepant inconsistency.

The fact that this inherently unsafe idea affects cross-image variables with all the fun of race conditions and other goodies to make debugging impossible simply underlines that this is a bad idea.”

- 11 (3) if *expr* is an array then the variable shall also be an array,
 12 (4) the shapes of the variable and *expr* shall conform unless the variable is an allocatable array
 13 that has the same rank as *expr* and is neither a co-array nor a co-indexed object,
 14 (5) if the variable is an allocatable co-array or co-indexed object, it shall not be polymorphic,
 15 (6) if the variable is polymorphic it shall be type compatible with *expr* and have the same rank;
 16 otherwise the declared types of the variable and *expr* shall conform as specified in Table
 17 7.11,
 18 (7) if the variable is of derived type each kind type parameter of the variable shall have the
 19 same value as the corresponding type parameter of *expr*, and
 20 (8) if the variable is of derived type each length type parameter of the variable shall have the
 21 same value as the corresponding type parameter of *expr* unless the variable is allocatable
 22 and its corresponding type parameter is deferred.

23 Table 7.11: **Type conformance for the intrinsic assignment statement**

Type of the variable	Type of <i>expr</i>
integer	integer, real, complex, bits
real	integer, real, complex, bits
complex	integer, real, complex, bits

Type conformance for the intrinsic assignment statement

(cont.)

Type of the variable	Type of <i>expr</i>
ISO 10646, ASCII, or default character	ISO 10646, ASCII, or default character
other character	character of the same kind type parameter as the variable
logical	logical, bits
bits	integer, real, complex, bits
derived type	same derived type as the variable

1 A **numeric intrinsic assignment statement** is an intrinsic assignment statement for which the vari-
 2 able and *expr* are of numeric type. A **character intrinsic assignment statement** is an intrinsic
 3 assignment statement for which the variable and *expr* are of type character. A **logical intrinsic as-**
 4 **signment statement** is an intrinsic assignment statement for which the variable and *expr* are of type
 5 logical. A **bits intrinsic assignment statement** is an intrinsic assignment statement for which either
 6 the variable or *expr* is of type bits. A **derived-type intrinsic assignment statement** is an intrinsic
 7 assignment statement for which the variable and *expr* are of derived type.

8 An **array intrinsic assignment statement** is an intrinsic assignment statement for which the variable
 9 is an array. The the variable shall not be a many-one array section (6.2.2.3.2).

10 If the variable is a pointer, it shall be associated with a definable target such that the type, type
 11 parameters, and shape of the target and *expr* conform.

12 **7.4.1.3 Interpretation of intrinsic assignments**

13 Execution of an intrinsic assignment causes, in effect, the evaluation of the expression *expr* and all
 14 expressions within *variable* (7.1.8), the possible conversion of *expr* to the type and type parameters of
 15 the variable (Table 7.12), and the definition of the variable with the resulting value. The execution of
 16 the assignment shall have the same effect as if the evaluation of *expr* and the evaluation of all expressions
 17 in *variable* occurred before any portion of the variable is defined by the assignment. The evaluation of
 18 expressions within *variable* shall neither affect nor be affected by the evaluation of *expr*. No value is
 19 assigned to the variable if it is of type character and zero length, or is an array of size zero.

20 If the variable is a pointer, the value of *expr* is assigned to the target of the variable.

21 If the variable is an allocated allocatable variable, it is deallocated if *expr* is an array of different shape,
 22 any of the corresponding length type parameter values of the variable and *expr* differ, or the variable
 23 is polymorphic and the dynamic type of the variable and *expr* differ. If the variable is or becomes an
 24 unallocated allocatable variable, then it is allocated with each deferred type parameter equal to the
 25 corresponding type parameter of *expr*, with the shape of *expr*, with each lower bound equal to the
 26 corresponding element of LBOUND(*expr*), and with the same dynamic type as *expr*.

NOTE 7.38

For example, given the declaration

```
CHARACTER(:), ALLOCATABLE :: NAME
```

then after the assignment statement

```
NAME = 'Dr. '//FIRST_NAME//' '//SURNAME
```

NAME will have the length LEN(FIRST_NAME)+LEN(SURNAME)+5, even if it had previously been unallocated, or allocated with a different length. However, for the assignment statement

```
NAME(:) = 'Dr. '//FIRST_NAME//' '//SURNAME
```


NOTE 7.38 (cont.)

NAME must already be allocated at the time of the assignment; the assigned value is truncated or blank padded to the previously allocated length of NAME.

- 1 Both *variable* and *expr* may contain references to any portion of the variable.

NOTE 7.39

For example, in the character intrinsic assignment statement:

```
STRING (2:5) = STRING (1:4)
```

the assignment of the first character of STRING to the second character does not affect the evaluation of STRING (1:4). If the value of STRING prior to the assignment was 'ABCDEF', the value following the assignment is 'AABCDF'.

- 2 If *expr* is a scalar and the variable is an array, the *expr* is treated as if it were an array of the same
3 shape as the variable with every element of the array equal to the scalar value of *expr*.
- 4 If the variable is an array, the assignment is performed element-by-element on corresponding array
5 elements of the variable and *expr*.

NOTE 7.40

For example, if A and B are arrays of the same shape, the array intrinsic assignment

```
A = B
```

assigns the corresponding elements of B to those of A; that is, the first element of B is assigned to the first element of A, the second element of B is assigned to the second element of A, etc.

If C is an allocatable array of rank 1, then

```
C = PACK (ARRAY, ARRAY > 0)
```

will cause C to contain all the positive elements of ARRAY in array element order; if C is not allocated or is allocated with the wrong size, it will be re-allocated to be of the correct size to hold the result of PACK.

- 6 The processor may perform the element-by-element assignment in any order.

NOTE 7.41

For example, the following program segment results in the values of the elements of array X being reversed:

```
REAL X (10)
...
X (1:10) = X (10:1:-1)
```

- 7 For a numeric intrinsic assignment statement, the variable and *expr* may have different numeric types
8 or different kind type parameters, in which case the value of *expr* is converted to the type and kind type
9 parameter of the variable according to the rules of Table 7.12.

Table 7.12: **Numeric conversion and the assignment statement**

Type of the variable	Value Assigned
integer	INT (<i>expr</i> , KIND = KIND (<i>variable</i>))
real	REAL (<i>expr</i> , KIND = KIND (<i>variable</i>))
complex	CMPLX (<i>expr</i> , KIND = KIND (<i>variable</i>))
Note: The functions INT, REAL, CMPLX, and KIND are the generic functions defined in 13.7.	

- 1 For a logical intrinsic assignment statement, the variable and *expr* may have different kind type param-
 2 eters, in which case the value of *expr* is converted to the kind type parameter of the variable.
- 3 For a character intrinsic assignment statement, the variable and *expr* may have different character length
 4 parameters in which case the conversion of *expr* to the length of the variable is as follows.
- 5 (1) If the length of the variable is less than that of *expr*, the value of *expr* is truncated from
 6 the right until it is the same length as the variable.
- 7 (2) If the length of the variable is greater than that of *expr*, the value of *expr* is extended on
 8 the right with blanks until it is the same length as the variable.
- 9 If the variable and *expr* have different kind type parameters, each character *c* in *expr* is converted to
 10 the kind type parameter of the variable by ACHAR(IACHAR(*c*),KIND(*variable*)).

NOTE 7.42

For nondefault character types, the blank padding character is processor dependent. When assigning a character expression to a variable of a different kind, each character of the expression that is not representable in the kind of the variable is replaced by a processor-dependent character.

- 11 For a bits intrinsic assignment statement, the variable and *expr* may have different types or different
 12 kind type parameters, in which case the value of *expr* is converted to the type and kind type parameter
 13 of the variable according to the rules of Table 7.13.

Table 7.13: **Bits conversion and the assignment statement**

Type of the variable	Value Assigned
integer	INT (<i>expr</i> , KIND = KIND (<i>variable</i>))
real	REAL (<i>expr</i> , KIND = KIND (<i>variable</i>))
complex	CMPLX (<i>expr</i> , KIND = KIND (<i>variable</i>))
logical	LOGICAL (<i>expr</i> , KIND = KIND (<i>variable</i>))
bits	BITS (<i>expr</i> , KIND = KIND (<i>variable</i>))
Note: The functions BITS, INT, REAL, CMPLX, and KIND are the generic functions defined in 13.7.	

NOTE 7.43

Bits assignment is not always the same as the result of the TRANSFER intrinsic, because:

- bits assignment operates elementally, whereas TRANSFER does not preserve array element boundaries;
- for scalars, if the source is larger TRANSFER uses those bits which occur first in memory whereas bits assignment always uses the “rightmost” bits (according to the model for bits values), independent of the endianness of the processor’s memory addressing;
- if the source is smaller, TRANSFER uses it for the part of the result which occurs first in

NOTE 7.43 (cont.)

memory address order and leaves the rest of the result processor-dependent, whereas bits assignment copies the source to the rightmost bits and makes the remaining bits all zero.

1 A derived-type intrinsic assignment is performed as if each component of the variable were assigned
 2 from the corresponding component of *expr* using pointer assignment (7.4.2) for each pointer component,
 3 defined assignment for each nonpointer nonallocatable component of a type that has a type-bound defined
 4 assignment consistent with the component, intrinsic assignment for each other nonpointer nonallocatable
 5 component, and intrinsic assignment for each allocated co-array component. For unallocated co-array
 6 components, the corresponding component of the variable shall be unallocated. For a non-co-array
 7 allocatable component the following sequence of operations is applied.

- 8 (1) If the component of the variable is allocated, it is deallocated.
- 9 (2) If the component of the value of *expr* is allocated, the corresponding component of the
 10 variable is allocated with the same dynamic type and type parameters as the component
 11 of the value of *expr*. If it is an array, it is allocated with the same bounds. The value of
 12 the component of the value of *expr* is then assigned to the corresponding component of the
 13 variable using defined assignment if the declared type of the component has a type-bound
 14 defined assignment consistent with the component, and intrinsic assignment for the dynamic
 15 type of that component otherwise.

16 The processor may perform the component-by-component assignment in any order or by any means that
 17 has the same effect.

NOTE 7.44

For an example of a derived-type intrinsic assignment statement, if C and D are of the same derived type with a pointer component P and nonpointer components S, T, U, and V of type integer, logical, character, and another derived type, respectively, the intrinsic

$$C = D$$

pointer assigns D%P to C%P. It assigns D%S to C%S, D%T to C%T, and D%U to C%U using intrinsic assignment. It assigns D%V to C%V using defined assignment if objects of that type have a compatible type-bound defined assignment, and intrinsic assignment otherwise.

NOTE 7.45

If an allocatable component of *expr* is unallocated, the corresponding component of the variable has an allocation status of unallocated after execution of the assignment.

18 **7.4.1.4 Defined assignment statement**

19 A **defined assignment statement** is an assignment statement that is defined by a subroutine and a
 20 generic interface (4.5.2, 12.4.3.3.2) that specifies ASSIGNMENT (=). A **defined elemental assign-**
 21 **ment statement** is a defined assignment statement for which the subroutine is elemental (12.8).

22 A subroutine defines the defined assignment $x_1 = x_2$ if

- 23 (1) the subroutine is specified with a SUBROUTINE (12.6.2.2) or ENTRY (12.6.2.5) statement
 24 that specifies two dummy arguments, d_1 and d_2 ,
- 25 (2) either
 - 26 (a) a generic interface (12.4.3.2) provides the subroutine with a *generic-spec* of ASSIGN-
 27 MENT (=), or

- 1 (b) there is a generic binding (4.5.2) in the declared type of x_1 or x_2 with a *generic-spec*
 2 of ASSIGNMENT (=) and there is a corresponding binding to the subroutine in the
 3 dynamic type of x_1 or x_2 , respectively,
- 4 (3) the types of d_1 and d_2 are compatible with the dynamic types of x_1 and x_2 , respectively,
 5 (4) the type parameters, if any, of d_1 and d_2 match the corresponding type parameters of x_1
 6 and x_2 , respectively, and
 7 (5) either
 8 (a) the ranks of x_1 and x_2 match those of d_1 and d_2 or
 9 (b) the subroutine is elemental, x_1 and x_2 are conformable, and there is no other subrou-
 10 tine that defines the operation.

11 If d_1 or d_2 is an array, the shapes of x_1 and x_2 shall match the shapes of d_1 and d_2 , respectively.

12 7.4.1.5 Interpretation of defined assignment statements

13 The interpretation of a defined assignment is provided by the subroutine that defines it.

14 If the defined assignment is an elemental assignment and the the variable in the assignment is an array,
 15 the defined assignment is performed element-by-element, on corresponding elements of the variable and
 16 *expr*. If *expr* is a scalar, it is treated as if it were an array of the same shape as the variable with every
 17 element of the array equal to the scalar value of *expr*.

NOTE 7.46

The rules of defined assignment (12.4.3.3.2), procedure references (12.5), subroutine references (12.5.4), and elemental subroutine arguments (12.8.3) ensure that the defined assignment has the same effect as if the evaluation of all operations in x_2 and x_1 occurs before any portion of x_1 is defined. If an elemental assignment is defined by a pure elemental subroutine, the element assignments may be performed simultaneously or in any order.

18 7.4.2 Pointer assignment

19 7.4.2.1 General

20 Pointer assignment causes a pointer to become associated with a target or causes its pointer association
 21 status to become disassociated or undefined. Any previous association between the pointer and a target
 22 is broken.

23 Pointer assignment for a pointer component of a structure may also take place by execution of a derived-
 24 type intrinsic assignment statement (7.4.1.3).

25 A pointer may also become associated with a target by allocation of the pointer.

26 7.4.2.2 Syntax

27 R735 *pointer-assignment-stmt* **is** *data-pointer-object* [(*bounds-spec-list*)] => *data-target*
 28 **or** *data-pointer-object* (*bounds-remapping-list*) => *data-target*
 29 **or** *proc-pointer-object* => *proc-target*
 30 R736 *data-pointer-object* **is** *variable-name*
 31 **or** *scalar-variable* % *data-pointer-component-name*

32 C717 (R735) If *data-target* is not unlimited polymorphic, *data-pointer-object* shall be type compatible
 33 (4.3.1.3) with it and the corresponding kind type parameters shall be equal.

- 1 C718 (R735) If *data-target* is unlimited polymorphic, *data-pointer-object* shall be unlimited polymorphic, of a sequence derived type, or of a type with the BIND attribute.
2
- 3 C719 (R735) If *bounds-spec-list* is specified, the number of *bounds-specs* shall equal the rank of *data-pointer-object*.
4
- 5 C720 (R735) If *bounds-remapping-list* is specified, the number of *bounds-remappings* shall equal the rank of *data-pointer-object*.
6
- 7 C721 (R735) If *bounds-remapping-list* is not specified, the ranks of *data-pointer-object* and *data-target* shall be the same.
8
- 9 C722 (R736) A *variable-name* shall have the POINTER attribute.
- 10 C723 (R736) A *scalar-variable* shall be a *data-ref*.
- 11 C724 (R736) A *data-pointer-component-name* shall be the name of a component of *scalar-variable* that is a data pointer.
12
- 13 C725 (R736) A *data-pointer-object* shall not be a co-indexed object.
- 14 R737 *bounds-spec* **is** *lower-bound-expr* :
15 R738 *bounds-remapping* **is** *lower-bound-expr* : *upper-bound-expr*
16 R739 *data-target* **is** *variable*
17 **or** *expr*
- 18 C726 (R739) A *variable* shall have either the TARGET or POINTER attribute, and shall not be an array section with a vector subscript.
19
- 20 C727 (R739) A *data-target* shall not be a co-indexed object.

NOTE 7.47

A data pointer and its target are always on the same image. A co-array may be of a derived type with pointer or allocatable subcomponents. For example, if PTR is a pointer component, Z[P]%PTR is a reference to the target of component PTR of Z on image P. This target is on image P and its association with Z[P]%PTR must have been established by the execution of an ALLOCATE statement or a pointer assignment on image P.

- 21 C728 (R739) An *expr* shall be a reference to a function whose result is a data pointer.
- 22 R740 *proc-pointer-object* **is** *proc-pointer-name*
23 **or** *proc-component-ref*
24 R741 *proc-component-ref* **is** *scalar-variable* % *procedure-component-name*
- 25 C729 (R741) The *scalar-variable* shall be a *data-ref*.
- 26 C730 (R741) The *procedure-component-name* shall be the name of a procedure pointer component of the declared type of *scalar-variable*.
27
- 28 R742 *proc-target* **is** *expr*
29 **or** *procedure-name*
30 **or** *proc-component-ref*
- 31 C731 (R742) An *expr* shall be a reference to a function whose result is a procedure pointer.
- 32 C732 (R742) A *procedure-name* shall be the name of an external, internal, module, or dummy procedure, a procedure pointer, or a specific intrinsic function listed in 13.6 and not marked with a
33

1 bullet (●).

2 C733 (R742) The *proc-target* shall not be a nonintrinsic elemental procedure.

3 **7.4.2.3 Data pointer assignment**

4 If *data-pointer-object* is not polymorphic and *data-target* is polymorphic with dynamic type that differs
5 from its declared type, the assignment target is the ancestor component of *data-target* that has the type
6 of *data-pointer-object*. Otherwise, the assignment target is *data-target*.

7 If *data-target* is not a pointer, *data-pointer-object* becomes pointer associated with the assignment target.
8 Otherwise, the pointer association status of *data-pointer-object* becomes that of *data-target*; if *data-target*
9 is associated with an object, *data-pointer-object* becomes associated with the assignment target. If *data-*
10 *target* is allocatable, it shall be allocated.

11 If *data-pointer-object* is polymorphic (4.3.1.3), it assumes the dynamic type of *data-target*. If *data-*
12 *pointer-object* is of sequence derived type or a type with the BIND attribute, the dynamic type of
13 *data-target* shall be that derived type.

14 If *data-target* is a disassociated pointer, all nondeferred type parameters of the declared type of *data-*
15 *pointer-object* that correspond to nondeferred type parameters of *data-target* shall have the same values
16 as the corresponding type parameters of *data-target*.

17 Otherwise, all nondeferred type parameters of the declared type of *data-pointer-object* shall have the
18 same values as the corresponding type parameters of *data-target*.

19 If *data-pointer-object* has nondeferred type parameters that correspond to deferred type parameters of
20 *data-target*, *data-target* shall not be a pointer with undefined association status.

21 If *data-pointer-object* has the CONTIGUOUS attribute, *data-target* shall be contiguous.

22 If *bounds-remapping-list* is specified, *data-target* shall be contiguous (5.3.6) or of rank one. It shall not
23 be a disassociated or undefined pointer, and the size of *data-target* shall not be less than the size of
24 *data-pointer-object*. The elements of the target of *data-pointer-object*, in array element order (6.2.2.2),
25 are the first SIZE(*data-pointer-object*) elements of *data-target*.

26 If no *bounds-remapping-list* is specified, the extent of a dimension of *data-pointer-object* is the extent of
27 the corresponding dimension of *data-target*. If *bounds-spec-list* appears, it specifies the lower bounds;
28 otherwise, the lower bound of each dimension is the result of the intrinsic function LBOUND (13.7.97)
29 applied to the corresponding dimension of *data-target*. The upper bound of each dimension is one less
30 than the sum of the lower bound and the extent.

31 **7.4.2.4 Procedure pointer assignment**

32 If the *proc-target* is not a pointer, *proc-pointer-object* becomes pointer associated with *proc-target*. Other-
33 wise, the pointer association status of *proc-pointer-object* becomes that of *proc-target*; if *proc-target* is
34 associated with a procedure, *proc-pointer-object* becomes associated with the same procedure.

35 If *proc-target* is the name of an internal procedure the **host instance** of *proc-pointer-object* becomes
36 the innermost currently executing instance of the host procedure. Otherwise if *proc-target* has a host
37 instance the host instance of *proc-pointer-object* becomes that instance. Otherwise *proc-pointer-object*
38 has no host instance.

39 If *proc-pointer-object* has an explicit interface, its characteristics shall be the same as *proc-target* except
40 that *proc-target* may be pure even if *proc-pointer-object* is not pure and *proc-target* may be an elemental
41 intrinsic procedure even if *proc-pointer-object* is not elemental.

- 1 If the characteristics of *proc-pointer-object* or *proc-target* are such that an explicit interface is required,
- 2 both *proc-pointer-object* and *proc-target* shall have an explicit interface.
- 3 If *proc-pointer-object* has an implicit interface and is explicitly typed or referenced as a function, *proc-*
- 4 *target* shall be a function. If *proc-pointer-object* has an implicit interface and is referenced as a subroutine,
- 5 *proc-target* shall be a subroutine.
- 6 If *proc-target* and *proc-pointer-object* are functions, they shall have the same type; corresponding type
- 7 parameters shall either both be deferred or both have the same value.
- 8 If *procedure-name* is a specific procedure name that is also a generic name, only the specific procedure
- 9 is associated with pointer-object.

10 7.4.2.5 Examples

NOTE 7.48

The following are examples of pointer assignment statements. (See Note 12.16 for declarations of P and BESSEL.)

```

NEW_NODE % LEFT => CURRENT_NODE
SIMPLE_NAME => TARGET_STRUCTURE % SUBSTRUCT % COMPONENT
PTR => NULL ( )
ROW => MAT2D (N, :)
WINDOW => MAT2D (I-1:I+1, J-1:J+1)
POINTER_OBJECT => POINTER_FUNCTION (ARG_1, ARG_2)
EVERY_OTHER => VECTOR (1:N:2)
WINDOW2 (0:, 0:) => MAT2D (ML:MU, NL:NU)
! P is a procedure pointer and BESSEL is a procedure with a
! compatible interface.
P => BESSEL

! Likewise for a structure component.
STRUCT % COMPONENT => BESSEL

```

NOTE 7.49

It is possible to obtain different-rank views of parts of an object by specifying upper bounds in pointer assignment statements. This requires that the object be either rank one or contiguous. Consider the following example, in which a matrix is under consideration. The matrix is stored as a rank-one object in MYDATA because its diagonal is needed for some reason – the diagonal cannot be gotten as a single object from a rank-two representation. The matrix is represented as a rank-two view of MYDATA.

```

real, target :: MYDATA ( NR*NC )      ! An automatic array
real, pointer :: MATRIX ( :, : )     ! A rank-two view of MYDATA
real, pointer :: VIEW_DIAG ( : )
MATRIX( 1:NR, 1:NC ) => MYDATA       ! The MATRIX view of the data
VIEW_DIAG => MYDATA( 1::NR+1 )       ! The diagonal of MATRIX

```

Rows, columns, or blocks of the matrix can be accessed as sections of MATRIX.

11 7.4.3 Masked array assignment – WHERE

1 The masked array assignment is used to mask the evaluation of expressions and assignment of values in
 2 array assignment statements, according to the value of a logical array expression.

3 7.4.3.1 General form of the masked array assignment

4 A **masked array assignment** is either a WHERE statement or a WHERE construct.

5	R743	<i>where-stmt</i>	is	WHERE (<i>mask-expr</i>) <i>where-assignment-stmt</i>
6	R744	<i>where-construct</i>	is	<i>where-construct-stmt</i>
7				[<i>where-body-construct</i>] ...
8				[<i>masked-elsewhere-stmt</i>
9				[<i>where-body-construct</i>] ...] ...
10				[<i>elsewhere-stmt</i>
11				[<i>where-body-construct</i>] ...]
12				<i>end-where-stmt</i>
13	R745	<i>where-construct-stmt</i>	is	[<i>where-construct-name</i> :] WHERE (<i>mask-expr</i>)
14	R746	<i>where-body-construct</i>	is	<i>where-assignment-stmt</i>
15			or	<i>where-stmt</i>
16			or	<i>where-construct</i>
17	R747	<i>where-assignment-stmt</i>	is	<i>assignment-stmt</i>
18	R748	<i>mask-expr</i>	is	<i>logical-expr</i>
19	R749	<i>masked-elsewhere-stmt</i>	is	ELSEWHERE (<i>mask-expr</i>) [<i>where-construct-name</i>]
20	R750	<i>elsewhere-stmt</i>	is	ELSEWHERE [<i>where-construct-name</i>]
21	R751	<i>end-where-stmt</i>	is	END WHERE [<i>where-construct-name</i>]

22 C734 (R747) A *where-assignment-stmt* that is a defined assignment shall be elemental.

23 C735 (R744) If the *where-construct-stmt* is identified by a *where-construct-name*, the corresponding
 24 *end-where-stmt* shall specify the same *where-construct-name*. If the *where-construct-stmt* is
 25 not identified by a *where-construct-name*, the corresponding *end-where-stmt* shall not specify
 26 a *where-construct-name*. If an *elsewhere-stmt* or a *masked-elsewhere-stmt* is identified by a
 27 *where-construct-name*, the corresponding *where-construct-stmt* shall specify the same *where-*
 28 *construct-name*.

29 C736 (R746) A statement that is part of a *where-body-construct* shall not be a branch target statement.

30 If a *where-construct* contains a *where-stmt*, a *masked-elsewhere-stmt*, or another *where-construct* then
 31 each *mask-expr* within the *where-construct* shall have the same shape. In each *where-assignment-stmt*,
 32 the *mask-expr* and the variable being defined shall be arrays of the same shape.

NOTE 7.50

Examples of a masked array assignment are:

```
WHERE (TEMP > 100.0) TEMP = TEMP - REDUCE_TEMP
WHERE (PRESSURE <= 1.0)
  PRESSURE = PRESSURE + INC_PRESSURE
  TEMP = TEMP - 5.0
ELSEWHERE
  RAINING = .TRUE.
END WHERE
```

33 7.4.3.2 Interpretation of masked array assignments

34 When a WHERE statement or a *where-construct-stmt* is executed, a control mask is established. In
 35 addition, when a WHERE construct statement is executed, a pending control mask is established. If

1 the statement does not appear as part of a *where-body-construct*, the *mask-expr* of the statement is
 2 evaluated, and the control mask is established to be the value of *mask-expr*. The pending control mask
 3 is established to have the value *.NOT. mask-expr* upon execution of a WHERE construct statement that
 4 does not appear as part of a *where-body-construct*. The *mask-expr* is evaluated only once.

5 Each statement in a WHERE construct is executed in sequence.

6 Upon execution of a *masked-elsewhere-stmt*, the following actions take place in sequence.

- 7 (1) The control mask m_c is established to have the value of the pending control mask.
- 8 (2) The pending control mask is established to have the value m_c *.AND.* (*.NOT. mask-expr*).
- 9 (3) The control mask m_c is established to have the value m_c *.AND.* *mask-expr*.

10 The *mask-expr* is evaluated at most once.

11 Upon execution of an ELSEWHERE statement, the control mask is established to have the value of the
 12 pending control mask. No new pending control mask value is established.

13 Upon execution of an ENDWHERE statement, the control mask and pending control mask are es-
 14 tablished to have the values they had prior to the execution of the corresponding WHERE construct
 15 statement. Following the execution of a WHERE statement that appears as a *where-body-construct*, the
 16 control mask is established to have the value it had prior to the execution of the WHERE statement.

NOTE 7.51

The establishment of control masks and the pending control mask is illustrated with the following example:

```

WHERE(cond1)      ! Statement 1
. . .
ELSEWHERE(cond2) ! Statement 2
. . .
ELSEWHERE        ! Statement 3
. . .
END WHERE

```

Following execution of statement 1, the control mask has the value *cond1* and the pending control mask has the value *.NOT. cond1*. Following execution of statement 2, the control mask has the value (*.NOT. cond1*) *.AND.* *cond2* and the pending control mask has the value (*.NOT. cond1*) *.AND.* (*.NOT. cond2*). Following execution of statement 3, the control mask has the value (*.NOT. cond1*) *.AND.* (*.NOT. cond2*). The false condition values are propagated through the execution of the masked ELSEWHERE statement.

17 Upon execution of a WHERE construct statement that is part of a *where-body-construct*, the pending
 18 control mask is established to have the value m_c *.AND.* (*.NOT. mask-expr*). The control mask is then
 19 established to have the value m_c *.AND.* *mask-expr*. The *mask-expr* is evaluated at most once.

20 Upon execution of a WHERE statement that is part of a *where-body-construct*, the control mask is
 21 established to have the value m_c *.AND.* *mask-expr*. The pending mask is not altered.

22 If a nonelemental function reference occurs in the *expr* or *variable* of a *where-assignment-stmt* or in a
 23 *mask-expr*, the function is evaluated without any masked control; that is, all of its argument expressions
 24 are fully evaluated and the function is fully evaluated. If the result is an array and the reference is not
 25 within the argument list of a nonelemental function, elements corresponding to true values in the control
 26 mask are selected for use in evaluating the *expr*, *variable* or *mask-expr*.

27 If an elemental operation or function reference occurs in the *expr* or *variable* of a *where-assignment-stmt*

- 1 or in a *mask-expr*, and is not within the argument list of a nonelemental function reference, the operation
 2 is performed or the function is evaluated only for the elements corresponding to true values of the control
 3 mask.
- 4 If an array constructor appears in a *where-assignment-stmt* or in a *mask-expr*, the array constructor is
 5 evaluated without any masked control and then the *where-assignment-stmt* is executed or the *mask-expr*
 6 is evaluated.
- 7 When a *where-assignment-stmt* is executed, the values of *expr* that correspond to true values of the
 8 control mask are assigned to the corresponding elements of the variable.
- 9 The value of the control mask is established by the execution of a WHERE statement, a WHERE con-
 10 struct statement, an ELSEWHERE statement, a masked ELSEWHERE statement, or an ENDWHERE
 11 statement. Subsequent changes to the value of entities in a *mask-expr* have no effect on the value of the
 12 control mask. The execution of a function reference in the mask expression of a WHERE statement is
 13 permitted to affect entities in the assignment statement.

NOTE 7.52

Examples of function references in masked array assignments are:

```

WHERE (A > 0.0)
A = LOG (A)           ! LOG is invoked only for positive elements.
A = A / SUM (LOG (A)) ! LOG is invoked for all elements
                    ! because SUM is transformational.
END WHERE

```

14 **7.4.4 FORALL**

- 15 FORALL constructs and statements are used to control the execution of assignment and pointer assign-
 16 ment statements with selection by sets of index values and an optional mask expression.

17 **7.4.4.1 The FORALL Construct**

- 18 The FORALL construct allows multiple assignments, masked array (WHERE) assignments, and nested
 19 FORALL constructs and statements to be controlled by a single *forall-triplet-spec-list* and *scalar-mask-*
 20 *expr*.

- | | | | | |
|----|------|-------------------------------|-----------|---|
| 21 | R752 | <i>forall-construct</i> | is | <i>forall-construct-stmt</i>
[<i>forall-body-construct</i>] ...
<i>end-forall-stmt</i> |
| 24 | R753 | <i>forall-construct-stmt</i> | is | [<i>forall-construct-name</i> :] FORALL <i>forall-header</i> |
| 25 | R754 | <i>forall-header</i> | is | (<i>forall-triplet-spec-list</i> [, <i>scalar-mask-expr</i>]) |
| 26 | R755 | <i>forall-triplet-spec</i> | is | <i>index-name</i> = <i>subscript</i> : <i>subscript</i> [: <i>stride</i>] |
| 27 | R619 | <i>subscript</i> | is | <i>scalar-int-expr</i> |
| 28 | R622 | <i>stride</i> | is | <i>scalar-int-expr</i> |
| 29 | R756 | <i>forall-body-construct</i> | is | <i>forall-assignment-stmt</i>
or <i>where-stmt</i>
or <i>where-construct</i>
or <i>forall-construct</i>
or <i>forall-stmt</i> |
| 34 | R757 | <i>forall-assignment-stmt</i> | is | <i>assignment-stmt</i>
or <i>pointer-assignment-stmt</i> |
| 36 | R758 | <i>end-forall-stmt</i> | is | END FORALL [<i>forall-construct-name</i>] |
| 37 | C737 | (R758) | | If the <i>forall-construct-stmt</i> has a <i>forall-construct-name</i> , the <i>end-forall-stmt</i> shall have |

- 1 the same *forall-construct-name*. If the *end-forall-stmt* has a *forall-construct-name*, the *forall-*
 2 *construct-stmt* shall have the same *forall-construct-name*.
- 3 C738 (R754) The *scalar-mask-expr* shall be scalar and of type logical.
- 4 C739 (R754) Any procedure referenced in the *scalar-mask-expr*, including one referenced by a defined
 5 operation, shall be a pure procedure (12.7).
- 6 C740 (R755) The *index-name* shall be a named scalar variable of type integer.
- 7 C741 (R755) A *subscript* or *stride* in a *forall-triplet-spec* shall not contain a reference to any *index-*
 8 *name* in the *forall-triplet-spec-list* in which it appears.
- 9 C742 (R756) A statement in a *forall-body-construct* shall not define an *index-name* of the *forall-*
 10 *construct*.
- 11 C743 (R756) Any procedure referenced in a *forall-body-construct*, including one referenced by a defined
 12 operation, assignment, or finalization, shall be a pure procedure.
- 13 C744 (R756) A *forall-body-construct* shall not be a branch target.
- 14 C745 (R757) The *variable* in an *assignment-stmt* that is a *forall-assignment-stmt* shall be a *designator*.

NOTE 7.53

An example of a FORALL construct is:

```
REAL :: A(10, 10), B(10, 10) = 1.0
. . .
FORALL (I = 1:10, J = 1:10, B(I, J) /= 0.0)
  A(I, J) = REAL (I + J - 2)
  B(I, J) = A(I, J) + B(I, J) * REAL (I * J)
END FORALL
```

NOTE 7.54

An assignment statement that is a FORALL body construct may be a scalar or array assignment statement, or a defined assignment statement. The variable being defined will normally use each index name in the *forall-triplet-spec-list*. For example

```
FORALL (I = 1:N, J = 1:N)
  A(:, I, :, J) = 1.0 / REAL(I + J - 1)
END FORALL
```

broadcasts scalar values to rank-two subarrays of A.

NOTE 7.55

An example of a FORALL construct containing a pointer assignment statement is:

```
TYPE ELEMENT
  REAL ELEMENT_WT
  CHARACTER (32), POINTER :: NAME
END TYPE ELEMENT
TYPE(ELEMENT) CHART(200)
REAL WEIGHTS (1000)
CHARACTER (32), TARGET :: NAMES (1000)
```

NOTE 7.55 (cont.)

```

. . .
FORALL (I = 1:200, WEIGHTS (I + N - 1) > .5)
  CHART(I) % ELEMENT_WT = WEIGHTS (I + N - 1)
  CHART(I) % NAME => NAMES (I + N - 1)
END FORALL

```

The results of this FORALL construct cannot be achieved with a WHERE construct because a pointer assignment statement is not permitted in a WHERE construct.

- 1 An *index-name* in a *forall-construct* has a scope of the construct (16.4). It is a scalar variable that has
- 2 the type and type parameters that it would have if it were the name of a variable in the scoping unit
- 3 that includes the FORALL, and this type shall be integer type; it has no other attributes.

NOTE 7.56

The use of *index-name* variables in a FORALL construct does not affect variables of the same name, for example:

```

INTEGER :: X = -1
REAL A(5, 4)
J = 100
. . .
FORALL (X = 1:5, J = 1:4)
  A (X, J) = J
END FORALL

```

After execution of the FORALL, the variables X and J have the values -1 and 100 and A has the value

```

1 2 3 4
1 2 3 4
1 2 3 4
1 2 3 4
1 2 3 4

```

4 7.4.4.2 Execution of the FORALL construct

5 There are three stages in the execution of a FORALL construct:

- 6 (1) determination of the values for *index-name* variables,
- 7 (2) evaluation of the *scalar-mask-expr*, and
- 8 (3) execution of the FORALL body constructs.

9 7.4.4.2.1 Determination of the values for index variables

10 The subscript and stride expressions in the *forall-triplet-spec-list* are evaluated. These expressions may
 11 be evaluated in any order. The set of values that a particular *index-name* variable assumes is determined
 12 as follows.

- 13 (1) The lower bound m_1 , the upper bound m_2 , and the stride m_3 are of type integer with the
 14 same kind type parameter as the *index-name*. Their values are established by evaluating
 15 the first subscript, the second subscript, and the stride expressions, respectively, including,
 16 if necessary, conversion to the kind type parameter of the *index-name* according to the rules
 17 for numeric conversion (Table 7.12). If a stride does not appear, m_3 has the value 1. The

- 1 value m_3 shall not be zero.
 2 (2) Let the value of max be $(m_2 - m_1 + m_3)/m_3$. If $max \leq 0$ for some *index-name*, the execution
 3 of the construct is complete. Otherwise, the set of values for the *index-name* is
 4 $m_1 + (k - 1) \times m_3$ where $k = 1, 2, \dots, max$.

5 The set of combinations of *index-name* values is the Cartesian product of the sets defined by each triplet
 6 specification. An *index-name* becomes defined when this set is evaluated.

NOTE 7.57

The *stride* may be positive or negative; the FORALL body constructs are executed as long as $max > 0$. For the *forall-triplet-spec*

```
I = 10:1:-1
```

max has the value 10

7 **7.4.4.2.2 Evaluation of the mask expression**

8 The *scalar-mask-expr*, if any, is evaluated for each combination of *index-name* values. If the *scalar-*
 9 *mask-expr* is not present, it is as if it were present with the value true. The *index-name* variables may
 10 be primaries in the *scalar-mask-expr*.

11 The **active combination of *index-name* values** is defined to be the subset of all possible combinations
 12 (7.4.4.2.1) for which the *scalar-mask-expr* has the value true.

NOTE 7.58

The *index-name* variables may appear in the mask, for example

```
FORALL (I=1:10, J=1:10, A(I) > 0.0 .AND. B(J) < 1.0)
. . .
```

13 **7.4.4.2.3 Execution of the FORALL body constructs**

14 The *forall-body-constructs* are executed in the order in which they appear. Each construct is executed
 15 for all active combinations of the *index-name* values with the following interpretation:

16 Execution of a *forall-assignment-stmt* that is an *assignment-stmt* causes the evaluation of *expr* and all
 17 expressions within *variable* for all active combinations of *index-name* values. These evaluations may be
 18 done in any order. After all these evaluations have been performed, each *expr* value is assigned to the
 19 corresponding *variable*. The assignments may occur in any order.

20 Execution of a *forall-assignment-stmt* that is a *pointer-assignment-stmt* causes the evaluation of all
 21 expressions within *data-target* and *data-pointer-object* or *proc-target* and *proc-pointer-object*, the de-
 22 termination of any pointers within *data-pointer-object* or *proc-pointer-object*, and the determination of
 23 the target for all active combinations of *index-name* values. These evaluations may be done in any
 24 order. After all these evaluations have been performed, each *data-pointer-object* or *proc-pointer-object*
 25 is associated with the corresponding target. These associations may occur in any order.

26 In a *forall-assignment-stmt*, a defined assignment subroutine shall not reference any *variable* that be-
 27 comes defined by the statement.

NOTE 7.59

The following FORALL construct contains two assignment statements. The assignment to array B uses the values of array A computed in the previous statement, not the values A had prior to

NOTE 7.59 (cont.)

execution of the FORALL.

```
FORALL ( I = 2:N-1, J = 2:N-1 )
  A ( I, J ) = A(I, J-1) + A(I, J+1) + A(I-1, J) + A(I+1, J)
  B ( I, J ) = 1.0 / A(I, J)
END FORALL
```

Computations that would otherwise cause error conditions can be avoided by using an appropriate *scalar-mask-expr* that limits the active combinations of the *index-name* values. For example:

```
FORALL ( I = 1:N, Y(I) /= 0.0 )
  X(I) = 1.0 / Y(I)
END FORALL
```

- 1 Each statement in a *where-construct* (7.4.3) within a *forall-construct* is executed in sequence. When
- 2 a *where-stmt*, *where-construct-stmt* or *masked-elsewhere-stmt* is executed, the statement's *mask-expr* is
- 3 evaluated for all active combinations of *index-name* values as determined by the outer *forall-constructs*,
- 4 masked by any control mask corresponding to outer *where-constructs*. Any *where-assignment-stmt* is
- 5 executed for all active combinations of *index-name* values, masked by the control mask in effect for the
- 6 *where-assignment-stmt*.

NOTE 7.60

This FORALL construct contains a WHERE statement and an assignment statement.

```
INTEGER A(5,4), B(5,4)
FORALL ( I = 1:5 )
  WHERE ( A(I,:) == 0 ) A(I,:) = I
  B ( I, :) = I / A(I, :)
END FORALL
```

When executed with the input array

```

      0  0  0  0
      1  1  1  0
A =   2  2  0  2
      1  0  2  3
      0  0  0  0
```

the results will be

```

      1  1  1  1
      1  1  1  2
A =   2  2  3  2
      1  4  2  3
      5  5  5  5

      1  1  1  1
      2  2  2  1
B =   1  1  1  1
      4  1  2  1
      1  1  1  1
```

For an example of a FORALL construct containing a WHERE construct with an ELSEWHERE statement, see C.4.5.

- 7 Execution of a *forall-stmt* or *forall-construct* causes the evaluation of the *subscript* and *stride* expressions
- 8 in the *forall-triplet-spec-list* for all active combinations of the *index-name* values of the outer FORALL

- 1 construct. The set of combinations of *index-name* values for the inner FORALL is the union of the sets
 2 defined by these bounds and strides for each active combination of the outer *index-name* values; it also
 3 includes the outer *index-name* values. The *scalar-mask-expr* is then evaluated for all combinations of the
 4 *index-name* values of the inner construct to produce a set of active combinations for the inner construct.
 5 If there is no *scalar-mask-expr*, it is as if it were present with the value .TRUE.. Each statement in the
 6 inner FORALL is then executed for each active combination of the *index-name* values.

NOTE 7.61

This FORALL construct contains a nested FORALL construct. It assigns the transpose of the strict lower triangle of array A (the section below the main diagonal) to the strict upper triangle of A.

```

INTEGER A (3, 3)
FORALL (I = 1:N-1 )
  FORALL ( J=I+1:N )
    A(I,J) = A(J,I)
  END FORALL
END FORALL

```

If prior to execution N = 3 and

```

      0  3  6
A =   1  4  7
      2  5  8

```

then after execution

```

      0  1  2
A =   1  4  5
      2  5  8

```

7 7.4.4.3 The FORALL statement

- 8 The FORALL statement allows a single assignment statement or pointer assignment to be controlled by
 9 a set of index values and an optional mask expression.

10 R759 *forall-stmt* **is** FORALL *forall-header forall-assignment-stmt*

- 11 A FORALL statement is equivalent to a FORALL construct containing a single *forall-body-construct*
 12 that is a *forall-assignment-stmt*.

- 13 The scope of an *index-name* in a *forall-stmt* is the statement itself (16.4).

NOTE 7.62

Examples of FORALL statements are:

```
FORALL (I=1:N) A(I,I) = X(I)
```

This statement assigns the elements of vector X to the elements of the main diagonal of matrix A.

```
FORALL (I = 1:N, J = 1:N) X(I,J) = 1.0 / REAL (I+J-1)
```

Array element X(I,J) is assigned the value (1.0 / REAL (I+J-1)) for values of I and J between 1

NOTE 7.62 (cont.)

and N, inclusive.

```
FORALL (I=1:N, J=1:N, Y(I,J) /= 0 .AND. I /= J) X(I,J) = 1.0 / Y(I,J)
```

This statement takes the reciprocal of each nonzero off-diagonal element of array Y(1:N, 1:N) and assigns it to the corresponding element of array X. Elements of Y that are zero or on the diagonal do not participate, and no assignments are made to the corresponding elements of X. The results from the execution of the example in Note 7.61 could be obtained with a single FORALL statement:

```
FORALL ( I = 1:N-1, J=1:N, J > I ) A(I,J) = A(J,I)
```

For more examples of FORALL statements, see C.4.6.

1 7.4.4.4 Restrictions on FORALL constructs and statements

- 2 A many-to-one assignment is more than one assignment to the same object, or association of more
 3 than one target with the same pointer, whether the object is referenced directly or indirectly through a
 4 pointer. A many-to-one assignment shall not occur within a single statement in a FORALL construct or
 5 statement. It is possible to assign or pointer assign to the same object in different assignment statements
 6 in a FORALL construct.

NOTE 7.63

The appearance of each *index-name* in the identification of the left-hand side of an assignment statement is helpful in eliminating many-to-one assignments, but it is not sufficient to guarantee there will be none. For example, the following is allowed

```
FORALL (I = 1:10)
  A (INDEX (I)) = B(I)
END FORALL
```

if and only if INDEX(1:10) contains no repeated values.

- 7 Within the scope of a FORALL construct, a nested FORALL statement or FORALL construct shall
 8 not have the same *index-name*. The *forall-header* expressions within a nested FORALL may depend on
 9 the values of outer *index-name* variables.

1 8 Execution control

2 8.1 Executable constructs containing blocks

3 8.1.1 General

4 The following are executable constructs that contain blocks:

- 5 (1) ASSOCIATE construct;
- 6 (2) BLOCK construct;
- 7 (3) CASE construct;
- 8 (4) CRITICAL construct;
- 9 (5) DO construct;
- 10 (6) IF construct;
- 11 (7) SELECT TYPE construct.

12 There is also a nonblock form of the DO construct.

13 A **block** is a sequence of executable constructs that is treated as a unit.

14 R801 *block* **is** [*execution-part-construct*] ...

15 Executable constructs may be used to control which blocks of a program are executed or how many times
 16 a block is executed. Blocks are always bounded by statements that are particular to the construct in
 17 which they are embedded; however, in some forms of the DO construct, a sequence of executable constructs without
 18 a terminating boundary statement shall obey all other rules governing blocks (8.1.2).

NOTE 8.1

A block need not contain any executable constructs. Execution of such a block has no effect.

NOTE 8.2

An example of a construct containing a block is:

```
IF (A > 0.0) THEN
  B = SQRT (A)  ! These two statements
  C = LOG (A)  ! form a block.
END IF
```

19 8.1.2 Rules governing blocks

20 8.1.2.1 Control flow in blocks

21 Transfer of control to the interior of a block from outside the block is prohibited. Transfers within a
 22 block and transfers from the interior of a block to outside the block may occur.

23 Subroutine and function references (12.5.3, 12.5.4) may appear in a block.

24 8.1.2.2 Execution of a block

1 Execution of a block begins with the execution of the first executable construct in the block. Execution
 2 of the block is completed when the last executable construct in the sequence is executed, when a branch
 3 (8.2) within the block that has a branch target outside the block occurs, when a RETURN statement
 4 within the block is executed, or when an EXIT or CYCLE statement that belongs to a construct that
 5 contains the block is executed.

NOTE 8.3

The action that takes place at the terminal boundary depends on the particular construct and on the block within that construct. It is usually a transfer of control.

6 8.1.3 ASSOCIATE construct

7 The **ASSOCIATE construct** associates named entities with expressions or variables during the exe-
 8 cution of its block. These named construct entities (16.4) are associating entities (16.5.1.6). The names
 9 are **associate names**.

10 8.1.3.1 Form of the ASSOCIATE construct

11	R802	<i>associate-construct</i>	is	<i>associate-stmt</i>
12				<i>block</i>
13				<i>end-associate-stmt</i>
14	R803	<i>associate-stmt</i>	is	[<i>associate-construct-name</i> :] ASSOCIATE ■
15				■ (<i>association-list</i>)
16	R804	<i>association</i>	is	<i>associate-name</i> => <i>selector</i>
17	R805	<i>selector</i>	is	<i>expr</i>
18			or	<i>variable</i>

19 C801 (R804) If *selector* is not a *variable* or is a *variable* that has a vector subscript, *associate-name*
 20 shall not appear in a variable definition context (16.6.7).

21 C802 (R804) An *associate-name* shall not be the same as another *associate-name* in the same *associate-*
 22 *stmt*.

23 C803 (R805) *expr* shall not be a reference to a function that has a pointer result.

24	R806	<i>end-associate-stmt</i>	is	END ASSOCIATE [<i>associate-construct-name</i>]
----	------	---------------------------	-----------	---

25 C804 (R806) If the *associate-stmt* of an *associate-construct* specifies an *associate-construct-name*,
 26 the corresponding *end-associate-stmt* shall specify the same *associate-construct-name*. If the
 27 *associate-stmt* of an *associate-construct* does not specify an *associate-construct-name*, the cor-
 28 responding *end-associate-stmt* shall not specify an *associate-construct-name*.

29 8.1.3.2 Execution of the ASSOCIATE construct

30 Execution of an ASSOCIATE construct causes execution of its *associate-stmt* followed by execution
 31 of its block. During execution of that block each associate name identifies an entity, which is associ-
 32 ated (16.5.1.6) with the corresponding selector. The associating entity assumes the declared type and
 33 type parameters of the selector. If and only if the selector is polymorphic, the associating entity is
 34 polymorphic.

35 The other attributes of the associating entity are described in 8.1.3.3.

36 It is permissible to branch to an *end-associate-stmt* only from within its ASSOCIATE construct.

37 8.1.3.3 Attributes of associate names

- 1 Within a SELECT TYPE or ASSOCIATE construct, each associating entity has the same rank as its
 2 associated selector. The lower bound of each dimension is the result of the intrinsic function LBOUND
 3 (13.7.97) applied to the corresponding dimension of *selector*. The upper bound of each dimension is one
 4 less than the sum of the lower bound and the extent. The associating entity has the ASYNCHRONOUS
 5 or VOLATILE attribute if and only if the selector is a variable and has the attribute. The associating
 6 entity has the TARGET attribute if and only if the selector is a variable and has either the TARGET
 7 or POINTER attribute. If the associating entity is polymorphic, it assumes the dynamic type and type
 8 parameter values of the selector. If the selector has the OPTIONAL attribute, it shall be present. The
 9 associating entity is contiguous if and only if the selector is contiguous.
- 10 If the selector (8.1.3.1) is not permitted to appear in a variable definition context (16.6.7), the associate
 11 name shall not appear in a variable definition context.

12 8.1.3.4 Examples of the ASSOCIATE construct

NOTE 8.4

The following example illustrates an association with an expression.

```
ASSOCIATE ( Z => EXP(-(X**2+Y**2)) * COS(THETA) )
  PRINT *, A+Z, A-Z
END ASSOCIATE
```

The following example illustrates an association with a derived-type variable.

```
ASSOCIATE ( XC => AX%B(I,J)%C )
  XC%DV = XC%DV + PRODUCT(XC%EV(1:N))
END ASSOCIATE
```

The following example illustrates association with an array section.

```
ASSOCIATE ( ARRAY => AX%B(I,:)%C )
  ARRAY(N)%EV = ARRAY(N-1)%EV
END ASSOCIATE
```

The following example illustrates multiple associations.

```
ASSOCIATE ( W => RESULT(I,J)%W, ZX => AX%B(I,J)%D, ZY => AY%B(I,J)%D )
  W = ZX*X + ZY*Y
END ASSOCIATE
```

13 8.1.4 BLOCK construct

- 14 The BLOCK construct is an executable construct which may contain declarations.

```
15 R807  block-construct          is  block-stmt
16                                     [ specification-part ]
17                                     block
18                                     end-block-stmt
```

```
19 R808  block-stmt              is  [ block-construct-name : ] BLOCK
```

```
20 R809  end-block-stmt         is  END BLOCK [ block-construct-name ]
```

- 21 C805 (R807) The *specification-part* of a BLOCK construct shall not contain a COMMON statement, IMPLICIT statement, INTENT statement, OPTIONAL statement, or USE statement.

J3 internal note

Unresolved Technical Issue 084

It would be better to allow the USE statement. Some work is required to handle modules going out of scope on exiting a BLOCK construct.

1 C806 (R807) If the *block-stmt* of a *block-construct* specifies a *block-construct-name*, the corresponding
 2 *end-block-stmt* shall specify the same *block-construct-name*. If the *block-stmt* does not specify
 3 a *block-construct-name*, the corresponding *end-block-stmt* shall not specify a *block-construct-*
 4 *name*.

5 Except for the ASYNCHRONOUS and VOLATILE statements, specifications in a BLOCK construct
 6 declare construct entities whose scope is that of the BLOCK construct (16.4).

7 Execution of a BLOCK construct causes evaluation of the specification expressions within its *specifica-*
 8 *tion-part* in a processor-dependent order, followed by execution of its block.

9 8.1.5 CASE construct**10 8.1.5.1 Form of the CASE construct**

11 The **CASE construct** selects for execution at most one of its constituent blocks. The selection is based
 12 on the value of an expression.

13 R810	<i>case-construct</i>	is	<i>select-case-stmt</i> [<i>case-stmt</i> <i>block</i>] ... <i>end-select-stmt</i>
17 R811	<i>select-case-stmt</i>	is	[<i>case-construct-name</i> :] SELECT CASE (<i>case-expr</i>)
18 R812	<i>case-stmt</i>	is	CASE <i>case-selector</i> [<i>case-construct-name</i>]
19 R813	<i>end-select-stmt</i>	is	END SELECT [<i>case-construct-name</i>]

20 C807 (R810) If the *select-case-stmt* of a *case-construct* specifies a *case-construct-name*, the corre-
 21 sponding *end-select-stmt* shall specify the same *case-construct-name*. If the *select-case-stmt*
 22 of a *case-construct* does not specify a *case-construct-name*, the corresponding *end-select-stmt*
 23 shall not specify a *case-construct-name*. If a *case-stmt* specifies a *case-construct-name*, the
 24 corresponding *select-case-stmt* shall specify the same *case-construct-name*.

25 R814	<i>case-expr</i>	is	<i>scalar-int-expr</i> or <i>scalar-char-expr</i> or <i>scalar-logical-expr</i>
28 R815	<i>case-selector</i>	is	(<i>case-value-range-list</i>) or DEFAULT

30 C808 (R810) No more than one of the selectors of one of the CASE statements shall be DEFAULT.

31 R816	<i>case-value-range</i>	is	<i>case-value</i> or <i>case-value</i> : or : <i>case-value</i> or <i>case-value</i> : <i>case-value</i>
35 R817	<i>case-value</i>	is	<i>scalar-int-initialization-expr</i> or <i>scalar-char-initialization-expr</i> or <i>scalar-logical-initialization-expr</i>

38 C809 (R810) For a given *case-construct*, each *case-value* shall be of the same type as *case-expr*. For
 39 character type, the kind type parameters shall be the same; character length differences are

1 allowed.

2 C810 (R810) A *case-value-range* using a colon shall not be used if *case-expr* is of type logical.

3 C811 (R810) For a given *case-construct*, the *case-value-ranges* shall not overlap; that is, there shall
4 be no possible value of the *case-expr* that matches more than one *case-value-range*.

5 8.1.5.2 Execution of a CASE construct

6 The execution of the SELECT CASE statement causes the case expression to be evaluated. The resulting
7 value is called the **case index**. For a case value range list, a match occurs if the case index matches any
8 of the case value ranges in the list. For a case index with a value of *c*, a match is determined as follows.

9 (1) If the case value range contains a single value *v* without a colon, a match occurs for type
10 logical if the expression *c* .EQV. *v* is true, and a match occurs for type integer or character
11 if the expression *c* == *v* is true.

12 (2) If the case value range is of the form *low* : *high*, a match occurs if the expression *low* <= *c*
13 .AND. *c* <= *high* is true.

14 (3) If the case value range is of the form *low* :, a match occurs if the expression *low* <= *c* is true.

15 (4) If the case value range is of the form : *high*, a match occurs if the expression *c* <= *high* is
16 true.

17 (5) If no other selector matches and a DEFAULT selector appears, it matches the case index.

18 (6) If no other selector matches and the DEFAULT selector does not appear, there is no match.

19 The block following the CASE statement containing the matching selector, if any, is executed. This
20 completes execution of the construct.

21 It is permissible to branch to an *end-select-stmt* only from within its CASE construct.

22 8.1.5.3 Examples of CASE constructs

NOTE 8.5

An integer signum function:

```

INTEGER FUNCTION SIGNUM (N)
SELECT CASE (N)
CASE (:-1)
    SIGNUM = -1
CASE (0)
    SIGNUM = 0
CASE (1:)
    SIGNUM = 1
END SELECT
END

```

NOTE 8.6

A code fragment to check for balanced parentheses:

```

CHARACTER (80) :: LINE
...
LEVEL = 0
SCAN_LINE: DO I = 1, 80
    CHECK_PARENS: SELECT CASE (LINE (I:I))
    CASE ('(')

```

NOTE 8.6 (cont.)

```

    LEVEL = LEVEL + 1
CASE (')')
    LEVEL = LEVEL - 1
    IF (LEVEL < 0) THEN
        PRINT *, 'UNEXPECTED RIGHT PARENTHESIS'
        EXIT SCAN_LINE
    END IF
CASE DEFAULT
    ! Ignore all other characters
END SELECT CHECK_PARENS
END DO SCAN_LINE
IF (LEVEL > 0) THEN
    PRINT *, 'MISSING RIGHT PARENTHESIS'
END IF

```

NOTE 8.7

The following three fragments are equivalent:

```

IF (SILLY == 1) THEN
    CALL THIS
ELSE
    CALL THAT
END IF
SELECT CASE (SILLY == 1)
CASE (.TRUE.)
    CALL THIS
CASE (.FALSE.)
    CALL THAT
END SELECT
SELECT CASE (SILLY)
CASE DEFAULT
    CALL THAT
CASE (1)
    CALL THIS
END SELECT

```

NOTE 8.8

A code fragment showing several selections of one block:

```

SELECT CASE (N)
CASE (1, 3:5, 8) ! Selects 1, 3, 4, 5, 8
    CALL SUB
CASE DEFAULT
    CALL OTHER
END SELECT

```

1 8.1.6 CRITICAL construct

- 2 A CRITICAL construct limits execution of a block to one image at a time.

- 1 R818 *critical-construct* is *critical-stmt*
 2 *block*
 3 *end-critical-stmt*
- 4 R819 *critical-stmt* is [*critical-construct-name* :] CRITICAL
- 5 R820 *end-critical-stmt* is END CRITICAL [*critical-construct-name*]
- 6 C812 (R818) If the *critical-stmt* of a *critical-construct* specifies a *critical-construct-name*, the corre-
 7 sponding *end-critical-stmt* shall specify the same *critical-construct-name*. If the *critical-stmt* of a
 8 *critical-construct* does not specify a *critical-construct-name*, the corresponding *end-critical-stmt*
 9 shall not specify a *critical-construct-name*.
- 10 C813 (R818) The *block* of a *critical-construct* shall not contain an image control statement.
- 11 Execution of the CRITICAL construct is completed when execution of its block is completed.
- 12 The processor shall ensure that once an image has commenced executing *block*, no other image shall
 13 commence executing it until the image has completed executing it. If image T is the next to execute the
 14 construct after image M, the segments (8.5.1) that executed before the construct on image M precede
 15 the segments that execute after the construct on image T. No image control statement shall be executed
 16 during the execution of a critical construct by the image executing the CRITICAL construct.

NOTE 8.9

If more than one image executes the block of a CRITICAL construct, its execution by one image always either precedes or succeeds its execution by another image. Typically no other statement ordering is needed. Consider the following example:

```
CRITICAL
  GLOBAL_COUNTER[1] = GLOBAL_COUNTER[1] + 1
END CRITICAL
```

The definition of GLOBAL_COUNTER[1] by a particular image will always precede the reference to the same variable by the next image to execute the block.

NOTE 8.10

The following example permits a large number of jobs to be shared among the images:

```
INTEGER :: NUM_JOBS[*], JOB

IF (THIS_IMAGE() == 1) READ(*,*) NUM_JOBS
SYNC ALL
DO
  CRITICAL
    JOB = NUM_JOBS[1]
    NUM_JOBS[1] = JOB - 1
  END CRITICAL
  IF (JOB > 0) THEN
    ! Work on JOB
  ELSE
    EXIT
  END IF
END DO
SYNC ALL
```

1 8.1.7 DO construct

2 8.1.7.1 General

3 The **DO construct** specifies the repeated execution of a sequence of executable constructs. Such a
4 repeated sequence is called a **loop**.

5 The number of iterations of a loop may be determined at the beginning of execution of the DO construct,
6 or may be left indefinite (“DO forever” or DO WHILE). Except in the case of a DO CONCURRENT
7 construct, the loop can be terminated immediately (8.1.7.5.4). The current iteration of the loop may be
8 curtailed by executing a CYCLE statement (8.1.7.5.3).

9 There are three phases in the execution of a DO construct: initiation of the loop, execution of the loop
10 range, and termination of the loop.

11 The DO CONCURRENT construct is a DO construct whose DO statement contains the CONCURRENT
12 keyword.

13 8.1.7.2 Forms of the DO construct

14 The DO construct may be written in either a block form or a nonblock form.

15 R821 *do-construct* **is** *block-do-construct*
16 **or** *nonblock-do-construct*

17 8.1.7.2.1 Form of the block DO construct

18 R822 *block-do-construct* **is** *do-stmt*
19 *do-block*
20 *end-do*
21 R823 *do-stmt* **is** *label-do-stmt*
22 **or** *nonlabel-do-stmt*
23 R824 *label-do-stmt* **is** [*do-construct-name* :] DO *label* [*loop-control*]
24 R825 *nonlabel-do-stmt* **is** [*do-construct-name* :] DO [*loop-control*]
25 R826 *loop-control* **is** [,] *do-variable* = *scalar-int-expr*, *scalar-int-expr* ■
26 ■ [, *scalar-int-expr*]
27 **or** [,] WHILE (*scalar-logical-expr*)
28 **or** [,] CONCURRENT *forall-header*
29 R827 *do-variable* **is** *scalar-int-variable-name*

30 C814 (R827) The *do-variable* shall be a variable of type integer.

31 R828 *do-block* **is** *block*
32 R829 *end-do* **is** *end-do-stmt*
33 **or** *continue-stmt*
34 R830 *end-do-stmt* **is** END DO [*do-construct-name*]

35 C815 (R822) If the *do-stmt* of a *block-do-construct* specifies a *do-construct-name*, the corresponding
36 *end-do* shall be an *end-do-stmt* specifying the same *do-construct-name*. If the *do-stmt* of a
37 *block-do-construct* does not specify a *do-construct-name*, the corresponding *end-do* shall not
38 specify a *do-construct-name*.

39 C816 (R822) If the *do-stmt* is a *nonlabel-do-stmt*, the corresponding *end-do* shall be an *end-do-stmt*.

40 C817 (R822) If the *do-stmt* is a *label-do-stmt*, the corresponding *end-do* shall be identified with the
41 same *label*.

1 8.1.7.2 Form of the nonblock DO construct

2	R831	<i>nonblock-do-construct</i>	is	<i>action-term-do-construct</i>
3			or	<i>outer-shared-do-construct</i>
4	R832	<i>action-term-do-construct</i>	is	<i>label-do-stmt</i>
5				<i>do-body</i>
6				<i>do-term-action-stmt</i>
7	R833	<i>do-body</i>	is	[<i>execution-part-construct</i>] ...
8	R834	<i>do-term-action-stmt</i>	is	<i>action-stmt</i>

9 C818 (R834) A *do-term-action-stmt* shall not be a *continue-stmt*, a *goto-stmt*, a *return-stmt*, a *stop-stmt*, an *exit-stmt*,
10 a *cycle-stmt*, an *end-function-stmt*, an *end-subroutine-stmt*, an *end-program-stmt*, or an *arithmetic-if-stmt*.

11 C819 (R831) The *do-term-action-stmt* shall be identified with a label and the corresponding *label-do-stmt* shall refer
12 to the same label.

13	R835	<i>outer-shared-do-construct</i>	is	<i>label-do-stmt</i>
14				<i>do-body</i>
15				<i>shared-term-do-construct</i>
16	R836	<i>shared-term-do-construct</i>	is	<i>outer-shared-do-construct</i>
17			or	<i>inner-shared-do-construct</i>
18	R837	<i>inner-shared-do-construct</i>	is	<i>label-do-stmt</i>
19				<i>do-body</i>
20				<i>do-term-shared-stmt</i>
21	R838	<i>do-term-shared-stmt</i>	is	<i>action-stmt</i>

22 C820 (R838) A *do-term-shared-stmt* shall not be a *goto-stmt*, a *return-stmt*, a *stop-stmt*, an *exit-stmt*, a *cycle-stmt*,
23 an *end-function-stmt*, an *end-subroutine-stmt*, an *end-program-stmt*, or an *arithmetic-if-stmt*.

24 C821 (R836) The *do-term-shared-stmt* shall be identified with a label and all of the *label-do-stmts* of the *inner-shared-*
25 *do-construct* and *outer-shared-do-construct* shall refer to the same label.

26 The *do-term-action-stmt*, *do-term-shared-stmt*, or *shared-term-do-construct* following the *do-body* of a nonblock DO con-
27 struct is called the **DO termination** of that construct.

28 Within a scoping unit, all DO constructs whose DO statements refer to the same label are nonblock DO constructs, and
29 share the statement identified by that label.

30 8.1.7.3 Range of the DO construct

31 The **range** of a block DO construct is the *do-block*, which shall satisfy the rules for blocks (8.1.2). In
32 particular, transfer of control to the interior of such a block from outside the block is prohibited. It
33 is permitted to branch to the *end-do* of a block DO construct only from within the range of that DO
34 construct.

35 The range of a nonblock DO construct consists of the *do-body* and the following DO termination. The end of such a
36 range is not bounded by a particular statement as for the other executable constructs (e.g., END IF); nevertheless, the
37 range satisfies the rules for blocks (8.1.2). Transfer of control into the *do-body* or to the DO termination from outside the
38 range is prohibited; in particular, it is permitted to branch to a *do-term-shared-stmt* only from within the range of the
39 corresponding *inner-shared-do-construct*.

40 8.1.7.4 Active and inactive DO constructs

41 A DO construct is either **active** or **inactive**. Initially inactive, a DO construct becomes active only
42 when its DO statement is executed.

43 Once active, the DO construct becomes inactive only when it terminates (8.1.7.5.4).

44 8.1.7.5 Execution of a DO construct

45 8.1.7.5.1 Loop initiation

When the DO statement is executed, the DO construct becomes active. If *loop-control* is

1 [,] *do-variable* = *scalar-int-expr*₁ , *scalar-int-expr*₂ [, *scalar-int-expr*₃]

2 the following steps are performed in sequence.

- 3 (1) The initial parameter m_1 , the terminal parameter m_2 , and the incrementation parameter
4 m_3 are of type integer with the same kind type parameter as the *do-variable*. Their values
5 are established by evaluating *scalar-int-expr*₁, *scalar-int-expr*₂, and *scalar-int-expr*₃, re-
6 spectively, including, if necessary, conversion to the kind type parameter of the *do-variable*
7 according to the rules for numeric conversion (Table 7.12). If *scalar-int-expr*₃ does not
8 appear, m_3 has the value 1. The value of m_3 shall not be zero.
- 9 (2) The DO variable becomes defined with the value of the initial parameter m_1 .
- 10 (3) The **iteration count** is established and is the value of the expression $(m_2 - m_1 + m_3)/m_3$,
11 unless that value is negative, in which case the iteration count is 0.

NOTE 8.11

The iteration count is zero whenever:

$m_1 > m_2$ and $m_3 > 0$, or
 $m_1 < m_2$ and $m_3 < 0$.

12 If *loop-control* is omitted, no iteration count is calculated. The effect is as if a large positive iteration
13 count, impossible to decrement to zero, were established. If *loop-control* is [,] WHILE (*scalar-logical-*
14 *expr*), the effect is as if *loop-control* were omitted and the following statement inserted as the first
15 statement of the *do-block*:

16 IF (.NOT. (*scalar-logical-expr*)) EXIT

17 For a DO CONCURRENT construct, the values of the index variables for the iterations of the construct
18 are determined by the rules for the index variables of the FORALL construct (7.4.4.2.1 and 7.4.4.2.2).

19 An *index-name* in a DO CONCURRENT construct has a scope of the construct (16.4). It is a scalar
20 variable that has the type and type parameters that it would have if it were the name of a variable in the
21 scoping unit that includes the construct, and this type shall be integer type; it has no other attributes.

22 At the completion of the execution of the DO statement, the execution cycle begins.

23 8.1.7.5.2 The execution cycle

24 The **execution cycle** of a DO construct that is not a DO CONCURRENT construct consists of the
25 following steps performed in sequence repeatedly until termination.

- 26 (1) The iteration count, if any, is tested. If it is zero, the loop terminates and the DO construct
27 becomes inactive. If *loop-control* is [,] WHILE (*scalar-logical-expr*), the *scalar-logical-expr*
28 is evaluated; if the value of this expression is false, the loop terminates and the DO construct
29 becomes inactive. If, as a result, all of the DO constructs sharing the *do-term-shared-stmt* are inactive,
30 the execution of all of these constructs is complete. However, if some of the DO constructs sharing the
31 *do-term-shared-stmt* are active, execution continues with step (3) of the execution cycle of the active DO
32 construct whose DO statement was most recently executed.
- 33 (2) The range of the loop is executed.
- 34 (3) The iteration count, if any, is decremented by one. The DO variable, if any, is incremented
35 by the value of the incrementation parameter m_3 .

36 Except for the incrementation of the DO variable that occurs in step (3), the DO variable shall neither
37 be redefined nor become undefined while the DO construct is active.

38 The range of a DO CONCURRENT construct is executed for all of the active combinations of the
index-name values. Each execution of the range is an iteration. The executions may occur in any order.

1 8.1.7.5.3 CYCLE statement

2 Execution of the range of the loop may be curtailed by executing a CYCLE statement from within the
3 range of the loop.

4 R839 *cycle-stmt* is CYCLE [*do-construct-name*]

5 C822 (R839) If a *do-construct-name* appears, the CYCLE statement shall be within the range of that
6 *do-construct*; otherwise, it shall be within the range of at least one *do-construct*.

7 C823 (R839) A *cycle-stmt* shall not appear within the range of a DO CONCURRENT construct if it
8 belongs to an outer construct.

9 A CYCLE statement **belongs** to a particular DO construct. If the CYCLE statement contains a DO
10 construct name, it belongs to that DO construct; otherwise, it belongs to the innermost DO construct
11 in which it appears.

12 Execution of a CYCLE statement that belongs to a DO construct that is not a DO CONCURRENT
13 construct causes immediate progression to step (3) of the current execution cycle of the DO construct
14 to which it belongs. If this construct is a nonblock DO construct, the *do-term-action-stmt* or *do-term-shared-stmt* is
15 not executed.

16 Execution of a CYCLE statement that belongs to a DO CONCURRENT construct completes execution
17 of that iteration of the construct.

18 In a block DO construct, a transfer of control to the *end-do* has the same effect as execution of a CYCLE
19 statement belonging to that construct. In a nonblock DO construct, transfer of control to the *do-term-action-stmt*
20 or *do-term-shared-stmt* causes that statement to be executed. Unless a further transfer of control results, step (3) of the
21 current execution cycle of the DO construct is then executed.

22 8.1.7.5.4 Loop termination

23 For a DO construct that is not a DO CONCURRENT construct, the loop terminates, and the DO
24 construct becomes inactive, when any of the following occurs.

- 25 (1) The iteration count is determined to be zero or the *scalar-logical-expr* is false, when tested
26 during step (1) of the above execution cycle.
- 27 (2) An EXIT statement belonging to the DO construct is executed.
- 28 (3) An EXIT or CYCLE statement that belongs to an outer construct and is within the range
29 of the DO construct is executed.
- 30 (4) Control is transferred from a statement within the range of a DO construct to a statement
31 that is neither the *end-do* nor within the range of the same DO construct.
- 32 (5) A RETURN statement within the range of the DO construct is executed.

33 For a DO CONCURRENT construct, the loop terminates, and the DO construct becomes inactive when
34 all of the iterations have completed execution.

35 When a DO construct becomes inactive, the DO variable, if any, of the DO construct retains its last
36 defined value.

37 8.1.7.6 Restrictions on DO CONCURRENT constructs

38 C824 A RETURN statement shall not appear within a DO CONCURRENT construct.

39 C825 A branch (8.2) within a DO CONCURRENT construct shall not have a branch target that is
40 outside the construct.

- 1 C826 A reference to a nonpure procedure shall not appear within a DO CONCURRENT construct.
- 2 C827 A reference to the procedure IEEE.GET_FLAG, IEEE.SET_HALTING_MODE, or IEEE.GET_-
3 HALTING_MODE from the intrinsic module IEEE_EXCEPTIONS, shall not be appear within
4 a DO CONCURRENT construct.

5 The following additional restrictions apply to DO CONCURRENT constructs.

- 6 • A variable that is referenced in an iteration shall either be previously defined during that iteration,
7 or shall be defined or become undefined during any other iteration of the current execution of
8 the construct. A variable that is defined or becomes undefined by more than one iteration of the
9 current execution of the construct becomes undefined when the current execution of the construct
10 terminates.
- 11 • A pointer that is referenced in an iteration either shall be previously pointer associated during that
12 iteration, or shall not have its pointer association changed during any iteration. A pointer that has
13 its pointer association changed in more than one iteration has a processor dependent association
14 status when the construct terminates.
- 15 • An allocatable object that is allocated in more than one iteration shall be subsequently deallocated
16 during the same iteration in which it was allocated. An object that is allocated or deallocated in
17 only one iteration shall not be deallocated, allocated, referenced, defined, or become undefined in
18 a different iteration.
- 19 • An input/output statement shall not write data to a file record or position in one iteration and
20 read from the same record or position in a different iteration of the same execution of the construct.
- 21 • Records written by output statements in the loop range to a sequential access file appear in the
22 file in an indeterminate order.

NOTE 8.12

The restrictions on referencing variables defined in an iteration of a DO CONCURRENT construct apply to any procedure invoked within the loop.

NOTE 8.13

The restrictions on the statements in the loop range of a DO CONCURRENT construct are designed to ensure there are no data dependencies between iterations of the loop. This permits code optimizations that might otherwise be difficult or impossible because they would depend on characteristics of the program not visible to the compiler.

23 8.1.7.7 Examples of DO constructs

NOTE 8.14

The following program fragment computes a tensor product of two arrays:

```
DO I = 1, M
  DO J = 1, N
    C (I, J) = SUM (A (I, J, :) * B (:, I, J))
  END DO
END DO
```

NOTE 8.15

The following program fragment contains a DO construct that uses the WHILE form of *loop-control*. The loop will continue to execute until an end-of-file or input/output error is encountered, at which point the DO statement terminates the loop. When a negative value of X is read, the program skips immediately to the next READ statement, bypassing most of the range of the loop.

```

READ (IUN, '(1X, G14.7)', IOSTAT = IOS) X
DO WHILE (IOS == 0)
  IF (X >= 0.) THEN
    CALL SUBA (X)
    CALL SUBB (X)
    . . .
    CALL SUBZ (X)
  ENDIF
  READ (IUN, '(1X, G14.7)', IOSTAT = IOS) X
END DO

```

NOTE 8.16

The following example behaves exactly the same as the one in Note 8.15. However, the READ statement has been moved to the interior of the range, so that only one READ statement is needed. Also, a CYCLE statement has been used to avoid an extra level of IF nesting.

```

DO      ! A "DO WHILE + 1/2" loop
  READ (IUN, '(1X, G14.7)', IOSTAT = IOS) X
  IF (IOS /= 0) EXIT
  IF (X < 0.) CYCLE
  CALL SUBA (X)
  CALL SUBB (X)
  . . .
  CALL SUBZ (X)
END DO

```

NOTE 8.17

The following example represents a case in which the user knows that the elements of the array IND form a permutation of the integers 1, ..., N. The DO CONCURRENT construct will allow the compiler to generate vector gather/scatter code, unroll the loop, or parallelize the code for this loop, significantly improving performance.

```

INTEGER :: A(N,N),IND(N)

DO CONCURRENT (I=1:N, J=1:N)
  A(IND(I),IND(J)) = A(IND(I),IND(J)) + 1
END DO

```

NOTE 8.18

Additional examples of DO constructs are in C.5.3.

1 8.1.8 IF construct and statement

1 8.1.8.1 Form of the IF construct

2 The **IF construct** selects for execution at most one of its constituent blocks. The selection is based on
3 a sequence of logical expressions.

```

4 R840  if-construct           is  if-then-stmt
5                                     block
6                                     [ else-if-stmt
7                                     block ] ...
8                                     [ else-stmt
9                                     block ]
10                                    end-if-stmt
11 R841  if-then-stmt          is  [ if-construct-name : ] IF ( scalar-logical-expr ) THEN
12 R842  else-if-stmt          is  ELSE IF ( scalar-logical-expr ) THEN [ if-construct-name ]
13 R843  else-stmt             is  ELSE [ if-construct-name ]
14 R844  end-if-stmt           is  END IF [ if-construct-name ]

```

15 C828 (R840) If the *if-then-stmt* of an *if-construct* specifies an *if-construct-name*, the corresponding
16 *end-if-stmt* shall specify the same *if-construct-name*. If the *if-then-stmt* of an *if-construct* does
17 not specify an *if-construct-name*, the corresponding *end-if-stmt* shall not specify an *if-construct-*
18 *name*. If an *else-if-stmt* or *else-stmt* specifies an *if-construct-name*, the corresponding *if-then-*
19 *stmt* shall specify the same *if-construct-name*.

20 8.1.8.2 Execution of an IF construct

21 At most one of the blocks in the IF construct is executed. If there is an ELSE statement in the construct,
22 exactly one of the blocks in the construct is executed. The scalar logical expressions are evaluated in
23 the order of their appearance in the construct until a true value is found or an ELSE statement or END
24 IF statement is encountered. If a true value or an ELSE statement is found, the block immediately
25 following is executed and this completes the execution of the construct. The scalar logical expressions
26 in any remaining ELSE IF statements of the IF construct are not evaluated. If none of the evaluated
27 expressions is true and there is no ELSE statement, the execution of the construct is completed without
28 the execution of any block within the construct.

29 It is permissible to branch to an END IF statement only from within its IF construct. Execution of an
30 END IF statement has no effect.

31 8.1.8.3 Examples of IF constructs

NOTE 8.19

```

IF (CVAR == 'RESET') THEN
  I = 0; J = 0; K = 0
END IF
PROOF_DONE: IF (PROP) THEN
  WRITE (3, '( 'QED' )')
  STOP
ELSE
  PROP = NEXTPROP
END IF PROOF_DONE
IF (A > 0) THEN
  B = C/A
  IF (B > 0) THEN
    D = 1.0
  END IF

```

NOTE 8.19 (cont.)

```

ELSE IF ( C > 0 ) THEN
  B = A/C
  D = -1.0
ELSE
  B = ABS ( MAX ( A, C ) )
  D = 0
END IF

```

1 8.1.8.4 IF statement

2 The **IF statement** controls the execution of a single action statement based on a single logical expression.

3 R845 *if-stmt* **is** IF (*scalar-logical-expr*) *action-stmt*

4 C829 (R845) The *action-stmt* in the *if-stmt* shall not be an *if-stmt*, *end-program-stmt*, *end-function-*
5 *stmt*, or *end-subroutine-stmt*.

6 Execution of an IF statement causes evaluation of the scalar logical expression. If the value of the
7 expression is true, the action statement is executed. If the value is false, the action statement is not
8 executed and execution continues.

9 The execution of a function reference in the scalar logical expression may affect entities in the action
10 statement.

NOTE 8.20

An example of an IF statement is:

```

IF ( A > 0.0 ) A = LOG ( A )

```

11 8.1.9 SELECT TYPE construct**12 8.1.9.1 Form of the SELECT TYPE construct**

13 The **SELECT TYPE** construct selects for execution at most one of its constituent blocks. The selection
14 is based on the dynamic type of an expression. A name is associated with the expression or variable
15 (16.4, 16.5.1.6), in the same way as for the ASSOCIATE construct.

16 R846 *select-type-construct* **is** *select-type-stmt*
17 [*type-guard-stmt*
18 *block*] ...
19 *end-select-type-stmt*

20 R847 *select-type-stmt* **is** [*select-construct-name* :] SELECT TYPE ■
21 ■ ([*associate-name* =>] *selector*)

22 C830 (R847) If *selector* is not a named *variable*, *associate-name* => shall appear.

23 C831 (R847) If *selector* is not a *variable* or is a *variable* that has a vector subscript, *associate-name*
24 shall not appear in a variable definition context (16.6.7).

25 C832 (R847) The *selector* in a *select-type-stmt* shall be polymorphic.

26 R848 *type-guard-stmt* **is** TYPE IS (*type-spec*) [*select-construct-name*]
27 **or** CLASS IS (*derived-type-spec*) [*select-construct-name*]
or CLASS DEFAULT [*select-construct-name*]

- 1 C833 (R848) The *type-spec* or *derived-type-spec* shall specify that each length type parameter is assumed.
2
- 3 C834 (R848) The *type-spec* or *derived-type-spec* shall not specify a sequence derived type or a type with the BIND attribute.
4
- 5 C835 (R848) If *selector* is not unlimited polymorphic, the *type-spec* or *derived-type-spec* shall specify an extension of the declared type of *selector*.
6
- 7 C836 (R848) For a given *select-type-construct*, the same type and kind type parameter values shall not be specified in more than one TYPE IS *type-guard-stmt* and shall not be specified in more than one CLASS IS *type-guard-stmt*.
8
9
- 10 C837 (R848) For a given *select-type-construct*, there shall be at most one CLASS DEFAULT *type-guard-stmt*.
11
- 12 R849 *end-select-type-stmt* **is** END SELECT [*select-construct-name*]
- 13 C838 (R846) If the *select-type-stmt* of a *select-type-construct* specifies a *select-construct-name*, the corresponding *end-select-type-stmt* shall specify the same *select-construct-name*. If the *select-type-stmt* of a *select-type-construct* does not specify a *select-construct-name*, the corresponding *end-select-type-stmt* shall not specify a *select-construct-name*. If a *type-guard-stmt* specifies a *select-construct-name*, the corresponding *select-type-stmt* shall specify the same *select-construct-name*.
14
15
16
17
18
- 19 The associate name of a SELECT TYPE construct is the *associate-name* if specified; otherwise it is the
20 *name* that constitutes the *selector*.

21 8.1.9.2 Execution of the SELECT TYPE construct

- 22 Execution of a SELECT TYPE construct whose selector is not a *variable* causes the selector expression
23 to be evaluated.
- 24 A SELECT TYPE construct selects at most one block to be executed. During execution of that block,
25 the associate name identifies an entity, which is associated (16.5.1.6) with the selector.
- 26 A TYPE IS type guard statement matches the selector if the dynamic type and kind type parameter
27 values of the selector are the same as those specified by the statement. A CLASS IS type guard
28 statement matches the selector if the dynamic type of the selector is an extension of the type specified
29 by the statement and the kind type parameter values specified by the statement are the same as the
30 corresponding type parameter values of the dynamic type of the selector.
- 31 The block to be executed is selected as follows.
- 32 (1) If a TYPE IS type guard statement matches the selector, the block following that statement
33 is executed.
 - 34 (2) Otherwise, if exactly one CLASS IS type guard statement matches the selector, the block
35 following that statement is executed.
 - 36 (3) Otherwise, if several CLASS IS type guard statements match the selector, one of these
37 statements must specify a type that is an extension of all the types specified in the others;
38 the block following that statement is executed.
 - 39 (4) Otherwise, if there is a CLASS DEFAULT type guard statement, the block following that
40 statement is executed.

NOTE 8.21

This algorithm does not examine the type guard statements in source text order when it looks for

NOTE 8.21 (cont.)

a match; it selects the most particular type guard when there are several potential matches.

- 1 Within the block following a TYPE IS type guard statement, the associating entity (16.5.5) is not
- 2 polymorphic (4.3.1.3), has the type named in the type guard statement, and has the type parameter
- 3 values of the selector.
- 4 Within the block following a CLASS IS type guard statement, the associating entity is polymorphic and
- 5 has the declared type named in the type guard statement. The type parameter values of the associating
- 6 entity are the corresponding type parameter values of the selector.
- 7 Within the block following a CLASS DEFAULT type guard statement, the associating entity is poly-
- 8 morphic and has the same declared type as the selector. The type parameter values of the associating
- 9 entity are those of the declared type of the selector.

NOTE 8.22

If the declared type of the *selector* is T, specifying CLASS DEFAULT has the same effect as specifying CLASS IS (T).

- 10 The other attributes of the associating entity are described in 8.1.3.3.
- 11 It is permissible to branch to an *end-select-type-stmt* only from within its SELECT TYPE construct.
- 12 **8.1.9.3 Examples of the SELECT TYPE construct**

NOTE 8.23

```

TYPE POINT
  REAL :: X, Y
END TYPE POINT
TYPE, EXTENDS(POINT) :: POINT_3D
  REAL :: Z
END TYPE POINT_3D
TYPE, EXTENDS(POINT) :: COLOR_POINT
  INTEGER :: COLOR
END TYPE COLOR_POINT

TYPE(POINT), TARGET :: P
TYPE(POINT_3D), TARGET :: P3
TYPE(COLOR_POINT), TARGET :: C
CLASS(POINT), POINTER :: P_OR_C
P_OR_C => C
SELECT TYPE ( A => P_OR_C )
CLASS IS ( POINT )
  ! "CLASS ( POINT ) :: A" implied here
  PRINT *, A%X, A%Y ! This block gets executed
TYPE IS ( POINT_3D )
  ! "TYPE ( POINT_3D ) :: A" implied here
  PRINT *, A%X, A%Y, A%Z
END SELECT

```

NOTE 8.24

The following example illustrates the omission of *associate-name*. It uses the declarations from Note 8.23.

```
P_OR_C => P3
SELECT TYPE ( P_OR_C )
CLASS IS ( POINT )
  ! "CLASS ( POINT ) :: P_OR_C" implied here
  PRINT *, P_OR_C%X, P_OR_C%Y
TYPE IS ( POINT_3D )
  ! "TYPE ( POINT_3D ) :: P_OR_C" implied here
  PRINT *, P_OR_C%X, P_OR_C%Y, P_OR_C%Z ! This block gets executed
END SELECT
```

1 **8.1.10 EXIT statement**

2 The EXIT statement provides one way of terminating a construct.

3 R850 *exit-stmt* **is** EXIT [*construct-name*]

4 C839 If a *construct-name* appears, the EXIT statement shall be within that construct; otherwise, it
5 shall be within the range of at least one *do-construct*.

6 An EXIT statement **belongs** to a particular construct. If the EXIT statement contains a construct
7 name, it belongs to that construct; otherwise, it belongs to the innermost DO construct in which it
8 appears.

9 C840 An *exit-stmt* shall not belong to a DO CONCURRENT construct, nor shall it appear within
10 the range of a DO CONCURRENT construct if it belongs to a construct that contains that DO
11 CONCURRENT construct.

12 When an EXIT statement that belongs to a DO construct is executed, it terminates the loop (8.1.7.5.4)
13 and any active loops contained within the terminated loop. When an EXIT statement that belongs to
14 a non-DO construct is executed, it terminates any active loops contained within that construct, and
15 completes execution of that construct.

16 **8.2 Branching**

17 **Branching** is used to alter the normal execution sequence. A branch causes a transfer of control from
18 one statement in a scoping unit to a labeled branch target statement in the same scoping unit. Branching
19 may be caused by a GOTO statement, a computed GOTO statement, an arithmetic IF statement, a
20 CALL statement that has an *alt-return-spec*, or an input/output statement that has an END= or ERR=
21 specifier. Although procedure references and control constructs can cause transfer of control, they are
22 not branches. A **branch target statement** is an *action-stmt*, an *associate-stmt*, an *end-associate-*
23 *stmt*, an *if-then-stmt*, an *end-if-stmt*, a *select-case-stmt*, an *end-select-stmt*, a *select-type-stmt*, an *end-*
24 *select-type-stmt*, a *do-stmt*, an *end-do-stmt*, *block-stmt*, *end-block-stmt*, *critical-stmt*, *end-critical-stmt*,
25 a *forall-construct-stmt*, a *do-term-action-stmt*, a *do-term-shared-stmt*, or a *where-construct-stmt*.

26 **8.2.1 GO TO statement**

27 5

28 R851 *goto-stmt* **is** GO TO *label*

1 C841 (R851) The *label* shall be the statement label of a branch target statement that appears in the
2 same scoping unit as the *goto-stmt*.

3 Execution of a GO TO statement causes a transfer of control so that the branch target statement
4 identified by the label is executed next.

5 8.2.2 Computed GO TO statement

6 R852 *computed-goto-stmt* is GO TO (*label-list*) [,] *scalar-int-expr*

7 C842 (R852) Each *label* in *label-list* shall be the statement label of a branch target statement that appears in the same
8 scoping unit as the *computed-goto-stmt*.

NOTE 8.25

The same statement label may appear more than once in a label list.

9 Execution of a computed GO TO statement causes evaluation of the scalar integer expression. If this value is i such that
10 $1 \leq i \leq n$ where n is the number of labels in *label-list*, a transfer of control occurs so that the next statement executed is
11 the one identified by the i th label in the list of labels. If i is less than 1 or greater than n , the execution sequence continues
12 as though a CONTINUE statement were executed.

13 8.2.3 Arithmetic IF statement

14 R853 *arithmetic-if-stmt* is IF (*scalar-numeric-expr*) *label* , *label* , *label*

15 C843 (R853) Each *label* shall be the label of a branch target statement that appears in the same scoping unit as the
16 *arithmetic-if-stmt*.

17 C844 (R853) The *scalar-numeric-expr* shall not be of type complex.

NOTE 8.26

The same label may appear more than once in one arithmetic IF statement.

18 Execution of an arithmetic IF statement causes evaluation of the numeric expression followed by a transfer of control. The
19 branch target statement identified by the first label, the second label, or the third label is executed next depending on
20 whether the value of the numeric expression is less than zero, equal to zero, or greater than zero, respectively.

21 8.3 CONTINUE statement

22 Execution of a CONTINUE statement has no effect.

23 R854 *continue-stmt* is CONTINUE

24 8.4 STOP statement

25 R855 *stop-stmt* is STOP [*stop-code*]

26 R856 *stop-code* is *scalar-char-initialization-expr*

27 or *scalar-int-initialization-expr*

28 C845 (R856) The *scalar-char-initialization-expr* shall be of default kind.

29 C846 (R856) The *scalar-int-initialization-expr* shall be of default kind.

30 Execution of a STOP statement causes normal termination of execution of that image. If each image
31 of a set of images executes a STOP statement immediately following the execution of a construct that
32 performs a synchronization of the images in the set, normal termination of execution occurs for all of

- 1 the images in the set; the executions of all other images are terminated. If termination of execution of
 2 an image occurs for some other reason, termination of execution occurs on all other images.

J3 internal note

Unresolved Technical Issue 005

1. Context-dependent meaning (immediately after SYNC ALL) is a bad idea, especially since “immediately following” is not well-defined.
2. Mixing “normal termination” and “error termination” is also probably a bad idea. In the past, STOP has always been “normal”, and various other terminations have been “error” (DEALLOCATE of an unallocated allocatable). We have never felt the need to have a special statement for causing error termination, and having “STOP,ALL:” do “normal” termination on some images and “error” termination on other images is baroque.
3. We do not need to have semantics like “what happens when the user presses Control-C”; we have never had anything in the Fortran standard that did that before either.
4. Notation (editorial): we should probably define the term “error termination” and use it in the multiple places we mean that, instead of merely omitting “normal” before “termination”.

J3 internal note

Unresolved Technical Issue 006

There is no explanation, let alone an adequate one, of when this (termination of execution of remote images) occurs. That means that there is no reason for the processor to terminate any other image, it can just let them run.

Surely we want to at least terminate a remote image at the next attempt to SYNC with the image that executes “STOP,ALL:”.

NOTE 8.27

If all images execute a SYNC ALL statement and immediately afterwards execute a STOP statement, normal termination of execution occurs on all images.

If only a subset of the images have their executions terminate normally, how soon termination takes place on the other images is processor dependent.

- 3 When an image is terminated by a STOP statement, its stop code, if any, is made available in a
 4 processor-dependent manner. If any exception (14) is signaling on that image, the processor shall issue
 5 a warning indicating which exceptions are signaling; this warning shall be on the unit identified by the
 6 named constant ERROR_UNIT (13.8.3.5). It is recommended that the stop code is made available by
 7 formatted output to the same unit.

NOTE 8.28

When normal termination occurs on more than one image, it is expected that a processor-dependent summary of the stop codes and signaling exceptions will be made available.

NOTE 8.29

If the *stop-code* is an integer, it is recommended that the value also be used as the process exit status, if the operating system supports that concept. If the integer *stop-code* is used as the process exit status, the operating system might be able to interpret only values within a limited range, or only a limited portion of the integer value (for example, only the least-significant 8 bits).

J3 internal note

Unresolved Technical Issue 033

This recommendation for the process exit status does not provide any help for STOP without a code, or for END PROGRAM. Perhaps “STOP” should be equivalent to “STOP 42” whereas “END PROGRAM” should be equivalent to “STOP 5”?

Not giving an exit code recommendation for “STOP charstring” is not very good either.

1 8.5 Image execution control

2 8.5.1 Image control statements

3 The execution sequence on each image is as specified in 2.3.6.

4 An **image control** statement affects the execution ordering between images. Each of the following is
5 an image control statement:

- 6 • SYNC ALL statement;
- 7 • SYNC TEAM statement;
- 8 • SYNC IMAGES statement;
- 9 • SYNC MEMORY statement;
- 10 • NOTIFY statement;
- 11 • QUERY statement;
- 12 • ALLOCATE or DEALLOCATE statement involving a co-array;
- 13 • CRITICAL or END CRITICAL statement (8.1.6);
- 14 • OPEN statement with a TEAM= specifier;
- 15 • CLOSE statement for a file that is open with a TEAM= specifier;
- 16 • END, END BLOCK, or RETURN statement that involves an implicit deallocation of a co-array;
- 17 • END PROGRAM statement;
- 18 • CALL statement for a collective subroutine (13.1).

19 During an execution of a statement that contains more than one reference to a procedure, image control
20 statements other than CRITICAL or END CRITICAL shall be executed in at most one of the procedures
21 invoked.

22 On each image, the sequence of statements executed before the first image control statement, between
23 the execution of two image control statements, or after the last image control statement is known as
24 a **segment**. The segment executed immediately before the execution of an image control statement
25 includes the evaluation of all expressions within the statement.

26 By execution of image control statements or user-defined ordering (8.5.6), the program can ensure that
27 the execution of the i^{th} segment on image P, P_i , either precedes or succeeds the execution of the j^{th}
28 segment on another image Q, Q_j . If the program does not ensure this, segments P_i and Q_j are unordered;
29 depending on the relative execution speeds of the images, some or all of the execution of the segment P_i
30 may take place at the same time as some or all of the execution of the segment Q_j .

NOTE 8.30

The set of all segments on all images is partially ordered: the segment P_i precedes segment Q_j if and only if there is a sequence of segments starting with P_i and ending with Q_j such that each segment of the sequence precedes the next either because they are on the same image or because of the execution of image control statements.

- 1 A scalar co-variable that is of type default integer, default logical, default real, or default bits, and has
 2 the VOLATILE attribute (5.3.18) may be referenced during the execution of a segment that is unordered
 3 relative to the execution of a segment in which the co-variable is defined. Otherwise,
- 4 • if a co-array variable is defined on an image in a segment, it shall not be referenced or defined in
 5 a segment on another image unless the segments are ordered,
 - 6 • if the allocation of an allocatable subobject of a co-array or the pointer association of a pointer
 7 subobject of a co-array is changed on an image in a segment, that subobject shall not be referenced
 8 or defined in a segment on another image unless the segments are ordered, and
 - 9 • if a procedure invocation on image P is in execution in segments P_i, P_{i+1}, \dots, P_k and defines a
 10 non-co-array dummy argument, the argument associated entity shall not be referenced or defined
 11 on another image Q in a segment Q_j unless Q_j precedes P_i or succeeds P_k .

NOTE 8.31

Apart from the effects of volatile variables, the processor may optimize the execution of a segment as if it were the only image in execution.

NOTE 8.32

The model upon which the interpretation of a program is based is that there is a permanent memory location for each co-array variable and that all images can access it. In practice, an image may make a copy of a non-volatile co-array variable (in cache or a register, for example) and, as an optimization, defer copying a changed value back to the permanent location while it is still being used. Since the variable is not volatile, it is safe to defer this transfer until the end of the current segment and thereafter to reload from permanent memory any co-array variable that was not defined within the segment. It would not be safe to defer these actions beyond the end of the current segment since another image might reference the variable then.

- 12 If an image P writes a record during the execution of P_i to a file that is opened for direct access with a
 13 TEAM= specifier, no other image Q shall read or write the record during execution of a segment that
 14 is unordered with P_i . Furthermore, it shall not read the record in a segment that succeeds P_i unless
- 15 • after image P writes the record, it executes a FLUSH statement (9.8) for the file during the
 16 execution of a segment P_k , where $k \geq i$, and
 - 17 • before image Q reads the record, it executes a FLUSH statement for the file during the execution
 18 of a segment Q_j that succeeds P_k .

NOTE 8.33

The incorrect sequencing of image control statements can halt execution indefinitely. For example, one image might be executing a SYNC ALL statement while another is executing an ALLOCATE statement for a co-array; or one image might be executing a blocking QUERY statement for which an image in its image set never executes the corresponding NOTIFY statement.

1 8.5.2 SYNC ALL statement

2 R857 *sync-all-stmt* is SYNC ALL [([*sync-stat-list*])]

3 R858 *sync-stat* is STAT = *stat-variable*

4 or ERRMSG = *errmsg-variable*

5 C847 No specifier shall appear more than once in a given *sync-stat-list*.

6 The STAT= and ERRMSG= specifiers for image execution control statements are described in 8.5.7.

7 Execution of a SYNC ALL statement performs a synchronization of all images. Execution on an image,
8 M, of the segment following the SYNC ALL statement is delayed until each other image has executed a
9 SYNC ALL statement as many times as has image M. The segments that executed before the SYNC ALL
10 statement on an image precede the segments that execute after the SYNC ALL statement on another
11 image.

NOTE 8.34

If synchronization is required when the images commence statement execution, a SYNC ALL statement should be the first executable statement of the main program. This is necessary if the code relies on the initialization of a co-array variable on another image.

NOTE 8.35

The processor might have special hardware or employ an optimized algorithm to make the SYNC ALL statement execute efficiently.

Here is a simple example of its use. Image 1 reads data and broadcasts it to other images:

```

REAL :: P[*]
...
SYNC ALL
IF (THIS_IMAGE()==1) THEN
  READ (*,*) P
  DO I = 2, NUM_IMAGES()
    P[I] = P
  END DO
END IF
SYNC ALL

```

12 8.5.3 SYNC TEAM statement

13 R859 *sync-team-stmt* is SYNC TEAM (*image-team* [, *sync-stat-list*])

14 R860 *image-team* is *scalar-variable*

15 C848 The *image-team* shall be a scalar variable of type IMAGE.TEAM from the intrinsic module
16 ISO_FORTRAN_ENV.

17 Execution of a SYNC TEAM statement performs a **team synchronization**, which is a synchronization
18 of the images in a team. The team is specified by the value of *image-team* and shall include the executing
19 image. All images of the team shall execute a SYNC TEAM statement with a value of *image-team* that
20 was constructed by corresponding invocations of the intrinsic function FORM_TEAM for the team. They
21 do not commence executing subsequent statements until all images in the team have executed a SYNC
22 TEAM statement for the team an equal number of times since FORM_TEAM was invoked for the team.
23 If images M and T are any two members of the team, the segments that execute before the statement

- 1 on image M precede the segments that execute after the statement on image T.
- 2 Execution of an OPEN statement with a TEAM= specifier, a CLOSE statement for a unit whose connect
 3 team consists of more than one image, or a CALL statement for a collective subroutine is interpreted as if
 4 an execution of a SYNC TEAM statement for the team occurred at the beginning and end of execution
 5 of the statement. The team is the set of images identified by the TEAM= specifier for the OPEN
 6 statement, the unit's connect team for the CLOSE statement, the IMAGES argument for the FORM-
 7 TEAM intrinsic procedure, or the argument of type IMAGE_TEAM for all other collective subroutines.
 8 This is called **implicit team synchronization**.

J3 internal note

Unresolved Technical Issue 086

What happens if *image-team* identifies a team that does not include the image executing the SYNC TEAM statement? (For example, if it is NULL_IMAGE_TEAM.) It would be friendly for this to raise an error in the statement rather than making the program non-conformant.

What happens if *image-team* is undefined (that is, has never been given a value)? It would be friendly for this to raise an error in the statement rather than making the program non-conformant.

Given that IMAGE_TEAM is a derived type, why not specify that it has default initialization and that this initializes it to the same value as NULL_IMAGE_TEAM.

J3 internal note

Unresolved Technical Issue 088

I assume that if a team consists of one image, the only effect of SYNC TEAM is to split the segment.

So why not make CLOSE (of a connected unit) always collective? The compiler is going to have to split the segment anyway, just in case the connect-team has more than one image in it.

NOTE 8.36

In this example the images are divided into two teams, one for an ocean calculation and one for an atmosphere calculation.

```

USE, INTRINSIC :: ISO_FORTRAN_ENV
TYPE(IMAGE_TEAM) :: TEAM
INTEGER :: N2, STEP, NSTEPS
LOGICAL :: OCEAN

N2 = NUM_IMAGES()/2
OCEAN = (THIS_IMAGE()<=N2)
IF(OCEAN) THEN
  CALL FORM_TEAM(TEAM, (/ (I, I=1,N2) /) )
ELSE
  CALL FORM_TEAM(TEAM, (/ (I, I=N2+1,NUM_IMAGES()) /) )
END IF
: ! Initial calculation
SYNC ALL
DO STEP = 1, NSTEPS
  IF (OCEAN) THEN
    DO
      : ! Ocean calculation
      SYNC TEAM(TEAM)
      IF ( ... ) EXIT ! Ready to swap data
    END DO
  ELSE

```


NOTE 8.36 (cont.)

```

DO
  : ! Atmosphere calculation
  SYNC TEAM(Team)
  IF ( ... ) EXIT ! Ready to swap data
END DO
END IF
SYNC ALL
: ! Swap data
END DO

```

In the inner loops, each set of images first works entirely with its own data and each image synchronizes with the rest of its team. The number of synchronizations for the ocean team might differ from the number for the atmosphere team. The SYNC ALL statement that follows is needed to ensure that both teams have done their calculations before data are swapped.

1 8.5.4 SYNC IMAGES statement

2 R861 *sync-images-stmt* is SYNC IMAGES (*image-set* [, *sync-stat-list*])

3 R862 *image-set* is *int-expr*
4 or *

5 C849 An *image-set* that is an *int-expr* shall be scalar or of rank one.

6 If *image-set* is an array expression, the value of each element shall be positive and not greater than the
7 number of images, and there shall be no repeated values.

8 If *image-set* is a scalar expression, its value shall be positive and not greater than the number of images.

9 An *image-set* that is an asterisk specifies all images.

10 Execution of a SYNC IMAGES statement performs a synchronization of the image with each of the
11 other images in the *image-set*. Execution on an image, M, of the segment following the SYNC IMAGES
12 statement is delayed until each other image T in the *image-set* has executed a SYNC IMAGES statement
13 with M in its *image-set* as many times as image M has executed a SYNC IMAGES statement with T in
14 the *image-set*. The segments that executed before the SYNC IMAGES statement on image M precede
15 the segments that execute after the SYNC IMAGES statement on image T.

NOTE 8.37

A SYNC IMAGES statement that specifies the single image value THIS_IMAGE() in its image set is allowed. This simplifies writing programs for an arbitrary number of images by allowing correct execution in the limiting case of the number of images being equal to one.

NOTE 8.38

Execution of a SYNC IMAGES(*) statement is not equivalent to the execution of a SYNC ALL statement. SYNC ALL causes all images to wait for each other. SYNC IMAGES statements are not required to specify the same image set on all the images participating in the synchronization. In the following example, image 1 will wait for each of the other images to complete its use of the data. The other images wait for image 1 to set up the data, but do not wait on any of the other images.

```
IF (THIS_IMAGE() == 1) then
```

NOTE 8.38 (cont.)

```

! Set up co-array data needed by all other images
SYNC IMAGES(*)
ELSE
  SYNC IMAGES(1)
! Use the data set up by image 1
END IF

```

NOTE 8.39

Execution of a SYNC TEAM statement causes all the images of the team to wait for each other. There might, however, be situations where this is not efficient. In the following example, each image synchronizes with its neighbor.

```

INTEGER :: ME, NE, STEP, NSTEPS
NE = NUM_IMAGES()
ME = THIS_IMAGE()
! Initial calculation
SYNC ALL
DO STEP = 1, NSTEPS
  IF (ME > 1) SYNC IMAGES(ME-1)
  ! Perform calculation
  IF (ME < NE) SYNC IMAGES(ME+1)
END DO
SYNC ALL

```

The calculation starts on image 1 since all the others will be waiting on SYNC IMAGES(ME-1). When this is done, image 2 can start and image 1 can perform its second calculation. This continues until they are all executing different steps at the same time. Eventually, image 1 will finish and then the others will finish one by one.

The SYNC IMAGES syntax involves *image-set* rather than *image-team* to allow the set of images to vary from image to image.

1 8.5.5 NOTIFY and QUERY statements

```

2 R863 notify-stmt           is NOTIFY ( image-set [ , sync-stat-list ] )
3 R864 query-stmt          is QUERY ( image-set [ , query-spec-list ] )
4 R865 query-spec         is READY = scalar-logical-variable
5                               or sync-stat

```

6 C850 (R864) No specifier shall appear more than once in a given *query-spec-list*.

7 Execution on image M of a NOTIFY statement with a different image T in its *image-set* increments by
8 1 a record of the number of times, $N_{M \rightarrow T}$, image M executed such a NOTIFY statement.

9 A QUERY statement is **blocking** if and only if it has no READY= specifier. A QUERY statement
10 is **satisfied** on completion of its execution if and only if it is a blocking QUERY statement or it set
11 the variable specified by its READY= specifier to true. Execution on image M of a satisfied QUERY
12 statement with a different image T included in its image set increases by 1 a record of the number of
13 times, $Q_{M \leftarrow T}$, image M executed such a QUERY statement. This increase is made after its value has
14 been compared with $N_{T \rightarrow M}$.

15 If a READY= specifier appears, execution on image M of a QUERY statement causes the *scalar-logical-*

- 1 *variable* to become defined. The value is false if, for a different image T in the image set, $N_{T \rightarrow M} \leq$
 2 $Q_{M \leftarrow T}$. Otherwise, the value is true.
- 3 If a READY= specifier does not appear, increasing $Q_{M \leftarrow T}$ and completing execution of the statement
 4 is delayed until, for each different T in the image set, $N_{T \rightarrow M} > Q_{M \leftarrow T}$.
- 5 A NOTIFY statement execution on image T and a satisfied QUERY statement on image M correspond
 6 if and only if
- 7 • the NOTIFY statement's image set includes image M,
 - 8 • the QUERY statement's image set includes image T, and
 - 9 • after execution of both statements has completed, $N_{T \rightarrow M} = Q_{M \leftarrow T}$.
- 10 Segments on an image executed before the execution of a NOTIFY statement precede the segments on
 11 other images that follow its corresponding QUERY statements.

NOTE 8.40

The NOTIFY and QUERY statements can be used to order statement executions between a producer and consumer image.

```

INTEGER,PARAMETER :: PRODUCER = 1, CONSUMER = 2
INTEGER :: VALUE[*]
LOGICAL :: READY

SELECT CASE (THIS_IMAGE())
CASE (PRODUCER)
  VALUE[CONSUMER] = 3
  NOTIFY (CONSUMER)
CASE (CONSUMER)
  WaitLoop: DO
    QUERY (PRODUCER,READY=READY)
    IF (READY) EXIT WaitLoop
    ! Statements neither referencing VALUE[CONSUMER], nor causing it to
    ! become defined or undefined
  END DO WaitLoop
  ! references to VALUE
CASE DEFAULT
  ! Statements neither referencing VALUE[CONSUMER], nor causing it to
  ! become defined or undefined
END SELECT

```

Unlike SYNC IMAGES statements, the number of notifications and corresponding queries may be unequal. A program can complete with an excess number of notifies.

NOTE 8.41

NOTIFY/QUERY pairs can be used in place of SYNC ALL and SYNC IMAGES to achieve better load balancing and allow one image to proceed with calculations while another image is catching up. For example,

```

IF (THIS_IMAGE()==1) THEN
  DO I=1,100

```

NOTE 8.41 (cont.)

```

...      ! Primary processing of column I
NOTIFY(2) ! Done with column I
END DO
SYNC IMAGES(2)
ELSE IF (THIS_IMAGE()==2) THEN
DO I=1,100
  QUERY(1) ! Wait until image 1 is done with column I
  ...      ! Secondary processing of column I
END DO
SYNC IMAGES(1)
END IF

```

1 8.5.6 SYNC MEMORY statement

- 2 The SYNC MEMORY statement provides a means of dividing a segment on an image into two segments,
3 each of which can be ordered in some user-defined way with respect to segments on other images.

4 R866 *sync-memory-stmt* is SYNC MEMORY [([*sync-stat-list*])]

NOTE 8.42

SYNC MEMORY usually suppresses compiler optimizations that might reorder memory operations across the segment boundary defined by the SYNC MEMORY statement and ensures that all memory operations initiated in the preceding segments in its image complete before any memory operations in the subsequent segment in its image are initiated. It needs to do this unless it can establish that failure to do so could not alter processing on another image.

- 5 All of the other image control statements include the effect of executing a SYNC MEMORY statement.
6 In addition, the other image control statements cause some form of cooperation with other images for
7 the purpose of ordering execution between images.

NOTE 8.43

A common example of user-written code that can be used in conjunction with SYNC MEMORY to implement specialized schemes for segment ordering is the spin-wait loop. For example:

```

LOGICAL, VOLATILE :: LOCKED[*] = .TRUE.
INTEGER :: IAM, P, Q

IAM = THIS_IMAGE()
IF (IAM == P) THEN
  ! Preceding segment
  SYNC MEMORY                ! A
  LOCKED[Q] = .FALSE.        ! segment  $P_i$ 
  SYNC MEMORY                ! B
ELSE IF (IAM == Q) THEN
  DO WHILE (LOCKED); END DO ! segment  $Q_j$ 
  SYNC MEMORY                ! C
  ! Subsequent segment
END IF

```

Here, image Q does not complete the segment Q_j until image P executes segment P_i . This ensures that executions of segments before P_i on image P precede executions of segments on image Q after

NOTE 8.43 (cont.) Q_j .

The first SYNC MEMORY statement (A) ensures that the compiler does not reorder the following statement (locking) with the previous statements, since the lock should be freed only after the work has been completed.

The definition of LOCKED[Q] might be deferred to the end of segment P_i . The second SYNC MEMORY statement (B) ends that segment immediately after the definition, minimizing any delay in releasing the lock in segment Q_j .

The third SYNC MEMORY statement (C) marks the beginning of a new segment, informing the compiler that the values of co-array variables referenced in that segment might have been changed by other images in preceding segments, so need to be loaded from memory.

NOTE 8.44

As a second example, the user might have access to an external procedure that performs synchronization between images. That library procedure will not necessarily be aware of the mechanisms used by the processor to manage remote data references and definitions, and therefore not, by itself, be able to ensure the correct memory state before and after its reference. The SYNC MEMORY statement provides the needed memory ordering that enables the safe use of the external synchronization routine. For example:

```

INTEGER :: IAM
REAL      :: X[*]

IAM = THIS_IMAGE()
IF (IAM == 1) X = 1.0
SYNC MEMORY
CALL EXTERNAL_SYNC()
SYNC MEMORY
IF (IAM == 2) WRITE(*,*) X[1]

```

where executing the subroutine EXTERNAL_SYNC has an image synchronization effect similar to executing a SYNC ALL statement.

1 8.5.7 STAT= and ERRMSG= specifiers in image execution control statements

- 2 If the STAT= specifier appears, successful execution of the SYNC ALL, SYNC TEAM, SYNC IMAGES,
3 SYNC MEMORY, NOTIFY, or QUERY statement causes the specified variable to become defined with
4 the value zero. If an error occurs during execution of one of these statements, the variable becomes
5 defined with a processor-dependent positive integer value. If an error condition occurs during execution
6 of a SYNC ALL, SYNC TEAM, SYNC IMAGES, SYNC MEMORY, NOTIFY, or QUERY statement
7 that does not contain the STAT= specifier, execution of all images is terminated.
- 8 If the ERRMSG= specifier appears and an error condition occurs during execution of the SYNC ALL,
9 SYNC TEAM, SYNC IMAGES, SYNC MEMORY, NOTIFY, or QUERY statement, the processor shall
10 assign an explanatory message to the specified variable. If no such condition occurs, the processor shall
11 not change the value of the variable.

NOTE 8.45

Which errors, if any, are diagnosed is processor dependent. The processor might check that a valid set of images has been provided, with no out-of-range or repeated values. It might test for network time-outs. While the overall program would probably not be able to recover from the failure of an image, it could perhaps provide information on what failed and be able to save some of the program state to a file.

1 **9 Input/output statements**

2 **Input statements** provide the means of transferring data from external media to internal storage or
3 from an internal file to internal storage. This process is called **reading**. **Output statements** provide
4 the means of transferring data from internal storage to external media or from internal storage to an
5 internal file. This process is called **writing**. Some input/output statements specify that editing of the
6 data is to be performed.

7 In addition to the statements that transfer data, there are auxiliary input/output statements to ma-
8 nipulate the external medium, or to describe or inquire about the properties of the connection to the
9 external medium.

10 The input/output statements are the OPEN, CLOSE, READ, WRITE, PRINT, BACKSPACE, END-
11 FILE, REWIND, FLUSH, WAIT, and INQUIRE statements.

12 The READ statement is a **data transfer input statement**. The WRITE statement and the PRINT
13 statement are **data transfer output statements**. The OPEN statement and the CLOSE state-
14 ment are **file connection statements**. The INQUIRE statement is a **file inquiry statement**. The
15 BACKSPACE, ENDFILE, and REWIND statements are **file positioning statements**.

16 A file is composed of either a sequence of file storage units (9.2.4) or a sequence of records, which
17 provide an extra level of organization to the file. A file composed of records is called a **record file**. A
18 file composed of file storage units is called a **stream file**. A processor may allow a file to be viewed
19 both as a record file and as a stream file; in this case the relationship between the file storage units when
20 viewed as a stream file and the records when viewed as a record file is processor dependent.

21 A file is either an external file (9.2) or an internal file (9.3).

22 **9.1 Records**

23 A **record** is a sequence of values or a sequence of characters. For example, a line on a terminal is usually
24 considered to be a record. However, a record does not necessarily correspond to a physical entity. There
25 are three kinds of records:

- 26 (1) formatted;
- 27 (2) unformatted;
- 28 (3) endfile.

NOTE 9.1

What is called a “record” in Fortran is commonly called a “logical record”. There is no concept in Fortran of a “physical record.”

29 **9.1.1 Formatted record**

30 A **formatted record** consists of a sequence of characters that are capable of representation in the
31 processor; however, a processor may prohibit some control characters (3.1) from appearing in a formatted
32 record. The length of a formatted record is measured in characters and depends primarily on the number
33 of characters put into the record when it is written. However, it may depend on the processor and the
34 external medium. The length may be zero. Formatted records may be read or written only by formatted
35 input/output statements.

1 Formatted records may be prepared by means other than Fortran.

2 9.1.2 Unformatted record

3 An **unformatted record** consists of a sequence of values in a processor-dependent form and may contain
4 data of any type or may contain no data. The length of an unformatted record is measured in file storage
5 units (9.2.4) and depends on the output list (9.5.3) used when it is written, as well as on the processor
6 and the external medium. The length may be zero. Unformatted records may be read or written only
7 by unformatted input/output statements.

8 9.1.3 Endfile record

9 An **endfile record** is written explicitly by the ENDFILE statement; the file shall be connected for
10 sequential access. An endfile record is written implicitly to a file connected for sequential access when
11 the most recent data transfer statement referring to the file is a data transfer output statement, no
12 intervening file positioning statement referring to the file has been executed, and

- 13 (1) a REWIND or BACKSPACE statement references the unit to which the file is connected,
14 or
- 15 (2) the unit is closed, either explicitly by a CLOSE statement, implicitly by a program termi-
16 nation not caused by an error condition, or implicitly by another OPEN statement for the
17 same unit.

18 An endfile record may occur only as the last record of a file. An endfile record does not have a length
19 property.

NOTE 9.2

An endfile record does not necessarily have any physical embodiment. The processor may use a record count or other means to register the position of the file at the time an ENDFILE statement is executed, so that it can take appropriate action when that position is reached again during a read operation. The endfile record, however it is implemented, is considered to exist for the BACKSPACE statement (9.7.2).

20 9.2 External files

21 An **external file** is any file that exists in a medium external to the program.

22 At any given time, there is a processor-dependent set of allowed **access methods**, a processor-dependent
23 set of allowed **forms**, a processor-dependent set of allowed **actions**, and a processor-dependent set of
24 allowed **record lengths** for a file.

NOTE 9.3

For example, the processor-dependent set of allowed actions for a printer would likely include the write action, but not the read action.

25 A file may have a name; a file that has a name is called a **named file**. The name of a named file is
26 represented by a character string value. The set of allowable names for a file is processor dependent.

NOTE 9.4

Named files that are opened with the TEAM= specifier (9.4.5.17) have the same name on each image of the team. Apart from this, whether a named file on one image is the same as a file with the same name on another image is processor dependent. For code portability, if different files are

NOTE 9.4 (cont.)

needed on each image, different file names should be used. One technique is to incorporate the image number as part of the name.

1 An external file that is connected to a unit has a **position** property (9.2.3).

NOTE 9.5

For more explanatory information on external files, see C.6.1.

2 9.2.1 File existence

3 At any given time, there is a processor-dependent set of external files that **exist** for a program. A file
4 may be known to the processor, yet not exist for a program at a particular time.

NOTE 9.6

Security reasons may prevent a file from existing for a program. A newly created file may exist
but contain no records.

5 To create a file means to cause a file to exist that did not exist previously. To delete a file means to
6 terminate the existence of the file.

7 All input/output statements may refer to files that exist. An INQUIRE, OPEN, CLOSE, WRITE,
8 PRINT, REWIND, FLUSH, or ENDFILE statement also may refer to a file that does not exist. Execu-
9 tion of a WRITE, PRINT, or ENDFILE statement referring to a preconnected file that does not exist
10 creates the file.

11 9.2.2 File access

12 There are three methods of accessing the data of an external file: sequential, direct, and stream. Some
13 files may have more than one allowed access method; other files may be restricted to one access method.

NOTE 9.7

For example, a processor may allow only sequential access to a file on magnetic tape. Thus, the
set of allowed access methods depends on the file and the processor.

14 The method of accessing a file is determined when the file is connected to a unit (9.4.3) or when the file
15 is created if the file is preconnected (9.4.4).

16 9.2.2.1 Sequential access

17 **Sequential access** is a method of accessing the records of an external record file in order.

18 When connected for sequential access, an external file has the following properties.

- 19 (1) The order of the records is the order in which they were written if the direct access method
20 is not a member of the set of allowed access methods for the file. If the direct access method
21 is also a member of the set of allowed access methods for the file, the order of the records
22 is the same as that specified for direct access. In this case, the first record accessible by
23 sequential access is the record whose record number is 1 for direct access. The second record
24 accessible by sequential access is the record whose record number is 2 for direct access, etc.
25 A record that has not been written since the file was created shall not be read.
- 26 (2) The records of the file are either all formatted or all unformatted, except that the last record
27 of the file may be an endfile record. Unless the previous reference to the file was a data

- 1 transfer output statement, the last record, if any, of the file shall be an endfile record.
- 2 (3) The records of the file shall be read or written only by sequential access input/output
3 statements.
- 4 (4) Each record shall be read or written by a single image. The processor shall ensure that once
5 an image commences transferring the data of a record to the file, no other image transfers
6 data to the file until the whole record has been transferred.

7 9.2.2.2 Direct access

8 **Direct access** is a method of accessing the records of an external record file in arbitrary order.

9 When connected for direct access, an external file has the following properties.

- 10 (1) Each record of the file is uniquely identified by a positive integer called the **record number**.
11 The record number of a record is specified when the record is written. Once established,
12 the record number of a record can never be changed. The order of the records is the order
13 of their record numbers.

NOTE 9.8

A record cannot be deleted; however, a record may be rewritten.

- 14 (2) The records of the file are either all formatted or all unformatted. If the sequential access
15 method is also a member of the set of allowed access methods for the file, its endfile record,
16 if any, is not considered to be part of the file while it is connected for direct access. If the
17 sequential access method is not a member of the set of allowed access methods for the file,
18 the file shall not contain an endfile record.
- 19 (3) The records of the file shall be read or written only by direct access input/output statements.
- 20 (4) All records of the file have the same length.
- 21 (5) Records need not be read or written in the order of their record numbers. Any record may
22 be written into the file while it is connected to a unit. For example, it is permissible to write
23 record 3, even though records 1 and 2 have not been written. Any record may be read from
24 the file while it is connected to a unit, provided that the record has been written since the
25 file was created, and if a READ statement for this connection is permitted.
- 26 (6) The records of the file shall not be read or written using list-directed formatting (10.10),
27 namelist formatting (10.11), or a nonadvancing input/output statement (9.2.3.1).

28 9.2.2.3 Stream access

29 **Stream access** is a method of accessing the file storage units (9.2.4) of an external stream file.

30 The properties of an external file connected for stream access depend on whether the connection is for
31 unformatted or formatted access.

32 When connected for unformatted stream access, an external file has the following properties.

- 33 (1) The file storage units of the file shall be read or written only by stream access input/output
34 statements.
- 35 (2) Each file storage unit in the file is uniquely identified by a positive integer called the position.
36 The first file storage unit in the file is at position 1. The position of each subsequent file
37 storage unit is one greater than that of its preceding file storage unit.
- 38 (3) If it is possible to position the file, the file storage units need not be read or written in
39 order of their position. For example, it might be permissible to write the file storage unit
40 at position 3, even though the file storage units at positions 1 and 2 have not been written.
41 Any file storage unit may be read from the file while it is connected to a unit, provided that

1 the file storage unit has been written since the file was created, and if a READ statement
2 for this connection is permitted.

3 When connected for formatted stream access, an external file has the following properties.

- 4 (1) Some file storage units of the file may contain record markers; this imposes a record structure
5 on the file in addition to its stream structure. There might or might not be a record marker
6 at the end of the file. If there is no record marker at the end of the file, the final record is
7 incomplete.
- 8 (2) No maximum length (9.4.5.13) is applicable to these records.
- 9 (3) Writing an empty record with no record marker has no effect.

NOTE 9.9

Because the record structure is determined from the record markers that are stored in the file itself, an incomplete record at the end of the file is necessarily not empty.

- 10 (4) The file storage units of the file shall be read or written only by formatted stream access
11 input/output statements.
- 12 (5) Each file storage unit in the file is uniquely identified by a positive integer called the position.
13 The first file storage unit in the file is at position 1. The relationship between positions of
14 successive file storage units is processor dependent; not all positive integers need correspond
15 to valid positions.
- 16 (6) If it is possible to position the file, the file position can be set to a position that was
17 previously identified by the POS= specifier in an INQUIRE statement.

NOTE 9.10

There may be some character positions in the file that do not correspond to characters written; this is because on some processors a record marker may be written to the file as a carriage-return/line-feed or other sequence. The means of determining the position in a file connected for stream access is via the POS= specifier in an INQUIRE statement (9.9.2.21).

- 18 (7) A processor may prohibit some control characters (3.1) from appearing in a formatted stream
19 file.

20 9.2.3 File position

21 Execution of certain input/output statements affects the position of an external file. Certain circum-
22 stances can cause the position of a file to become indeterminate.

23 The **initial point** of a file is the position just before the first record or file storage unit. The **terminal**
24 **point** is the position just after the last record or file storage unit. If there are no records or file storage
25 units in the file, the initial point and the terminal point are the same position.

26 If a record file is positioned within a record, that record is the **current record**; otherwise, there is no
27 current record.

28 Let n be the number of records in the file. If $1 < i \leq n$ and a file is positioned within the i th record or
29 between the $(i - 1)$ th record and the i th record, the $(i - 1)$ th record is the **preceding record**. If $n \geq 1$
30 and the file is positioned at its terminal point, the preceding record is the n th and last record. If $n = 0$
31 or if a file is positioned at its initial point or within the first record, there is no preceding record.

32 If $1 \leq i < n$ and a file is positioned within the i th record or between the i th and $(i + 1)$ th record, the
33 $(i + 1)$ th record is the **next record**. If $n \geq 1$ and the file is positioned at its initial point, the first record
34 is the next record. If $n = 0$ or if a file is positioned at its terminal point or within the n th (last) record,
35 there is no next record.

1 For a file connected for stream access, the file position is either between two file storage units, at the
2 initial point of the file, at the terminal point of the file, or undefined.

3 **9.2.3.1 Advancing and nonadvancing input/output**

4 An **advancing input/output statement** always positions a record file after the last record read or
5 written, unless there is an error condition.

6 A **nonadvancing input/output statement** may position a record file at a character position within
7 the current record, or a subsequent record (10.8.2). Using nonadvancing input/output, it is possible to
8 read or write a record of the file by a sequence of input/output statements, each accessing a portion
9 of the record. It is also possible to read variable-length records and be notified of their lengths. If a
10 nonadvancing output statement leaves a file positioned within a current record and no further output
11 statement is executed for the file before it is closed or a BACKSPACE, ENDFILE, or REWIND statement
12 is executed for it, the effect is as if the output statement were the corresponding advancing output
13 statement.

14 **9.2.3.2 File position prior to data transfer**

15 The positioning of the file prior to data transfer depends on the method of access: sequential, direct, or
16 stream.

17 For sequential access on input, if there is a current record, the file position is not changed. Otherwise,
18 the file is positioned at the beginning of the next record and this record becomes the current record.
19 Input shall not occur if there is no next record or if there is a current record and the last data transfer
20 statement accessing the file performed output.

21 If the file contains an endfile record, the file shall not be positioned after the endfile record prior to data
22 transfer. However, a REWIND or BACKSPACE statement may be used to reposition the file.

23 For sequential access on output, if there is a current record, the file position is not changed and the
24 current record becomes the last record of the file. Otherwise, a new record is created as the next record
25 of the file; this new record becomes the last and current record of the file and the file is positioned at
26 the beginning of this record.

27 For direct access, the file is positioned at the beginning of the record specified by the REC= specifier.
28 This record becomes the current record.

29 For stream access, the file is positioned immediately before the file storage unit specified by the POS=
30 specifier; if there is no POS= specifier, the file position is not changed.

31 File positioning for child data transfer statements is described in 9.5.4.7.

32 **9.2.3.3 File position after data transfer**

33 If an error condition (9.10) occurred, the position of the file is indeterminate. If no error condition
34 occurred, but an end-of-file condition (9.10) occurred as a result of reading an endfile record, the file is
35 positioned after the endfile record.

36 For unformatted stream access, if no error condition occurred, the file position is not changed. For
37 unformatted stream output, if the file position exceeds the previous terminal point of the file, the
38 terminal point is set to the file position.

NOTE 9.11

An unformatted stream output statement with a POS= specifier and an empty output list can have the effect of extending the terminal point of a file without actually writing any data.
--

- 1 For formatted stream input, if an end-of-file condition occurred, the file position is not changed.
- 2 For nonadvancing input, if no error condition or end-of-file condition occurred, but an end-of-record
3 condition (9.10) occurred, the file is positioned after the record just read. If no error condition, end-of-
4 file condition, or end-of-record condition occurred in a nonadvancing input statement, the file position
5 is not changed. If no error condition occurred in a nonadvancing output statement, the file position is
6 not changed.
- 7 In all other cases, the file is positioned after the record just read or written and that record becomes the
8 preceding record.
- 9 For a formatted stream output statement, if no error condition occurred, the terminal point of the file
10 is set to the highest-numbered position to which data was transferred by the statement.

NOTE 9.12

The highest-numbered position might not be the current one if the output involved T or TL edit descriptors (10.8.1.1) and the statement is a nonadvancing output statement.

11 9.2.4 File storage units

12 A **file storage unit** is the basic unit of storage in a stream file or an unformatted record file. It is the
13 unit of file position for stream access, the unit of record length for unformatted files, and the unit of file
14 size for all external files.

15 Every value in a stream file or an unformatted record file shall occupy an integer number of file storage
16 units; if the stream or record file is unformatted, this number shall be the same for all scalar values of
17 the same type and type parameters. The number of file storage units required for an item of a given type
18 and type parameters may be determined using the IOLENGTH= specifier of the INQUIRE statement
19 (9.9.3).

20 For a file connected for unformatted stream access, the processor shall not have alignment restrictions
21 that prevent a value of any type from being stored at any positive integer file position.

22 The number of bits in a file storage unit is given by the constant FILE_STORAGE_SIZE (13.8.3.6)
23 defined in the intrinsic module ISO_FORTRAN_ENV. It is recommended that the file storage unit be
24 an 8-bit octet where this choice is practical.

NOTE 9.13

The requirement that every data value occupy an integer number of file storage units implies that data items inherently smaller than a file storage unit will require padding. This suggests that the file storage unit be small to avoid wasted space. Ideally, the file storage unit would be chosen such that padding is never required. A file storage unit of one bit would always meet this goal, but would likely be impractical because of the alignment requirements.

The prohibition on alignment restrictions prohibits the processor from requiring data alignments larger than the file storage unit.

The 8-bit octet is recommended as a good compromise that is small enough to accommodate the requirements of many applications, yet not so small that the data alignment requirements are likely to cause significant performance problems.

25 9.3 Internal files

1 Internal files provide a means of transferring and converting data from internal storage to internal
2 storage.

3 An internal file is a record file with the following properties.

- 4 (1) The file is a variable of default, ASCII, or ISO 10646 character type that is not an array
5 section with a vector subscript.
- 6 (2) A record of an internal file is a scalar character variable.
- 7 (3) If the file is a scalar character variable, it consists of a single record whose length is the same
8 as the length of the scalar character variable. If the file is a character array, it is treated
9 as a sequence of character array elements. Each array element, if any, is a record of the
10 file. The ordering of the records of the file is the same as the ordering of the array elements
11 in the array (6.2.2.2) or the array section (6.2.2.3). Every record of the file has the same
12 length, which is the length of an array element in the array.
- 13 (4) A record of the internal file becomes defined by writing the record. If the number of
14 characters written in a record is less than the length of the record, the remaining portion
15 of the record is filled with blanks. The number of characters to be written shall not exceed
16 the length of the record.
- 17 (5) A record may be read only if the record is defined.
- 18 (6) A record of an internal file may become defined (or undefined) by means other than an
19 output statement. For example, the character variable may become defined by a character
20 assignment statement.
- 21 (7) An internal file is always positioned at the beginning of the first record prior to data transfer,
22 except for child data transfer statements (9.5.4.7). This record becomes the current record.
- 23 (8) The initial value of a connection mode (9.4.1) is the value that would be implied by an
24 initial OPEN statement without the corresponding keyword.
- 25 (9) Reading and writing records shall be accomplished only by sequential access formatted
26 input/output statements.
- 27 (10) An internal file shall not be specified as the unit in a file connection statement or a file
28 positioning statement.

29 9.4 File connection

30 A **unit**, specified by an *io-unit*, provides a means for referring to a file.

31 R901	<i>io-unit</i>	is	<i>file-unit-number</i>
32		or	*
33		or	<i>internal-file-variable</i>
34 R902	<i>file-unit-number</i>	is	<i>scalar-int-expr</i>
35 R903	<i>internal-file-variable</i>	is	<i>char-variable</i>

36 C901 (R903) The *char-variable* shall not be an array section with a vector subscript.

37 C902 (R903) The *char-variable* shall be of type default character, ASCII character, or ISO 10646
38 character.

39 A unit is either an external unit or an internal unit. An **external unit** is used to refer to an external file
40 and is specified by an asterisk or a *file-unit-number*. The value of *file-unit-number* shall be nonnegative,
41 equal to one of the named constants INPUT_UNIT, OUTPUT_UNIT, or ERROR_UNIT of the ISO_
42 FORTRAN_ENV module (13.8.3), or a NEWUNIT value (9.4.5.10). An **internal unit** is used to refer
43 to an internal file and is specified by an *internal-file-variable* or a *file-unit-number* whose value is equal
44 to the **unit** argument of an active derived-type input/output procedure (9.5.4.7). The value of a *file-*
45 *unit-number* shall identify a valid unit.

- 1 The external unit identified by a particular value of a *scalar-int-expr* is the same external unit in all
 2 program units of the program, and on all images.

NOTE 9.14

In the example:

```

SUBROUTINE A
  READ (6) X
  ...
SUBROUTINE B
  N = 6
  REWIND N
  
```

the value 6 used in both program units identifies the same external unit.

- 3 An asterisk identifies particular processor-dependent external units that are preconnected for format-
 4 ted sequential access (9.5.4.2). These units are also identified by unit numbers defined by the named
 5 constants INPUT_UNIT and OUTPUT_UNIT of the ISO_FORTRAN_ENV module (13.8.3).

- 6 This standard identifies a processor-dependent external unit for the purpose of error reporting. This
 7 unit shall be preconnected for sequential formatted output. The processor may define this to be the
 8 same as the output unit identified by an asterisk. This unit is also identified by a unit number defined
 9 by the named constant ERROR_UNIT of the ISO_FORTRAN_ENV intrinsic module.

10 9.4.1 Connection modes

- 11 A connection for formatted input/output has several changeable modes: the blank interpretation mode
 12 (10.8.6), delimiter mode (10.10.4, 10.11.4.1), sign mode (10.8.4), decimal edit mode (10.8.8), I/O round-
 13 ing mode (10.7.2.3.7), pad mode (9.5.4.4.2), and scale factor (10.8.5). A connection for unformatted
 14 input/output has no changeable modes.

- 15 Values for the modes of a connection are established when the connection is initiated. If the connection
 16 is initiated by an OPEN statement, the values are as specified, either explicitly or implicitly, by the
 17 OPEN statement. If the connection is initiated other than by an OPEN statement (that is, if the file is
 18 an internal file or preconnected file) the values established are those that would be implied by an initial
 19 OPEN statement without the corresponding keywords.

- 20 The scale factor cannot be explicitly specified in an OPEN statement; it is implicitly 0.

- 21 The modes of a connection to an external file may be changed by a subsequent OPEN statement that
 22 modifies the connection.

- 23 The modes of a connection may be temporarily changed by a corresponding keyword specifier in a
 24 data transfer statement or by an edit descriptor. Keyword specifiers take effect at the beginning of
 25 execution of the data transfer statement. Edit descriptors take effect when they are encountered in
 26 format processing. When a data transfer statement terminates, the values for the modes are reset to the
 27 values in effect immediately before the data transfer statement was executed.

28 9.4.2 Unit existence

- 29 At any given time, there is a processor-dependent set of external units that **exist** for a program.

- 30 All input/output statements may refer to units that exist. The CLOSE, INQUIRE, and WAIT state-
 31 ments also may refer to units that do not exist.

1 9.4.3 Connection of a file to a unit

2 An external unit has a property of being **connected** or not connected. If connected, it refers to an
3 external file, and the connection is for a team of images. An external unit may become connected by
4 preconnection or by the execution of an OPEN statement. The property of connection is symmetric; the
5 unit is connected to a file if and only if the file is connected to the unit.

J3 internal note

Unresolved Technical Issue 039

The second sentence now contains irrelevant commentary, and merely implies contradiction with the first sentence, rather than a direct contradiction.

Since (i/o units now are or soon will be) not global, what relevance has the TEAM= specifier to the property of connection? No more than whether the unit is connected for formatted/unformatted, direct/sequential/stream, etc.

Not doing the edit to the second sentence (i.e. deleting “, and the connection is for a team of images”) would be sufficient to resolve this issue here; the discussion of TEAM= also warrants reviewing.

6 Every input/output statement except an OPEN, CLOSE, INQUIRE, or WAIT statement shall refer to
7 a unit that is connected to a file and thereby make use of or affect that file.

8 A file may be connected and not exist (9.2.1).

NOTE 9.15

An example is a preconnected external file that has not yet been written.

9 A unit shall not be connected to more than one file at the same time, and a file shall not be connected to
10 more than one unit at the same time. However, means are provided to change the status of an external
11 unit and to connect a unit to a different file.

12 This standard defines means of portable interoperability with C. C streams are described in 7.19.2 of the C
13 International Standard. Whether a unit may be connected to a file that is also connected to a C stream
14 is processor dependent. If the processor allows a unit to be connected to a file that is also connected to
15 a C stream, the results of performing input/output operations on such a file are processor dependent.
16 It is processor dependent whether the files connected to the units INPUT_UNIT, OUTPUT_UNIT,
17 and ERROR_UNIT correspond to the predefined C text streams standard input, standard output, and
18 standard error. If a procedure defined by means of Fortran and a procedure defined by means other than
19 Fortran perform input/output operations on the same external file, the results are processor dependent.
20 A procedure defined by means of Fortran and a procedure defined by means other than Fortran can
21 perform input/output operations on different external files without interference.

22 After an external unit has been disconnected by the execution of a CLOSE statement, it may be con-
23 nected again within the same program to the same file or to a different file. After an external file has
24 been disconnected by the execution of a CLOSE statement, it may be connected again within the same
25 program to the same unit or to a different unit.

NOTE 9.16

The only means of referencing a file that has been disconnected is by the appearance of its name in an OPEN or INQUIRE statement. There might be no means of reconnecting an unnamed file once it is disconnected.

26 An internal unit is always connected to the internal file designated by the variable that identifies the
27 unit.

NOTE 9.17

For more explanatory information on file connection properties, see C.6.5.
--

1 **9.4.4 Preconnection**

2 **Preconnection** means that the unit is connected to a file at the beginning of execution of the program
 3 and therefore it may be specified in input/output statements without the prior execution of an OPEN
 4 statement.

5 **9.4.5 OPEN statement**

6 An **OPEN statement** initiates or modifies the connection between an external file and a specified unit.
 7 The OPEN statement may be used to connect an existing file to a unit, create a file that is preconnected,
 8 create a file and connect it to a unit, or change certain modes of a connection between a file and a unit.

9 An external unit may be connected by an OPEN statement in any program unit of a program and, once
 10 connected, a reference to it may appear in any program unit of the program.

11 If the file to be connected to the unit does not exist but is the same as the file to which the unit is
 12 preconnected, the modes specified by an OPEN statement become a part of the connection.

13 If the file to be connected to the unit is not the same as the file to which the unit is connected, the effect
 14 is as if a CLOSE statement without a STATUS= specifier had been executed for the unit immediately
 15 prior to the execution of an OPEN statement.

16 If a unit is connected to a file that exists, execution of an OPEN statement for that unit is permitted.
 17 If the FILE= specifier is not included in such an OPEN statement, the file to be connected to the unit
 18 is the same as the file to which the unit is already connected.

19 If the file to be connected to the unit is the same as the file to which the unit is connected, only the
 20 specifiers for changeable modes (9.4.1) may have values different from those currently in effect. If the
 21 POSITION= specifier appears in such an OPEN statement, the value specified shall not disagree with
 22 the current position of the file. If the STATUS= specifier is included in such an OPEN statement, it shall
 23 be specified with the value OLD. Execution of such an OPEN statement causes any new values of the
 24 specifiers for changeable modes to be in effect, but does not cause any change in any of the unspecified
 25 specifiers and the position of the file is unaffected. The ERR=, IOSTAT=, and IOMSG= specifiers from
 26 any previously executed OPEN statement have no effect on any currently executed OPEN statement.

27 A STATUS= specifier with a value of OLD is always allowed when the file to be connected to the unit is
 28 the same as the file to which the unit is connected. In this case, if the status of the file was SCRATCH
 29 before execution of the OPEN statement, the file will still be deleted when the unit is closed, and the
 30 file is still considered to have a status of SCRATCH.

31 If a file is already connected to a unit, execution of an OPEN statement on that file and a different unit
 32 is not permitted.

33 R904	<i>open-stmt</i>	is	OPEN (<i>connect-spec-list</i>)
34 R905	<i>connect-spec</i>	is	[UNIT =] <i>file-unit-number</i>
35		or	ACCESS = <i>scalar-default-char-expr</i>
36		or	ACTION = <i>scalar-default-char-expr</i>
37		or	ASYNCHRONOUS = <i>scalar-default-char-expr</i>
38		or	BLANK = <i>scalar-default-char-expr</i>
39		or	DECIMAL = <i>scalar-default-char-expr</i>
40		or	DELIM = <i>scalar-default-char-expr</i>
41		or	ENCODING = <i>scalar-default-char-expr</i>

- 1 or ERR = *label*
 2 or FILE = *file-name-expr*
 3 or FORM = *scalar-default-char-expr*
 4 or IOMSG = *iomsg-variable*
 5 or IOSTAT = *scalar-int-variable*
 6 or NEWUNIT = *scalar-int-variable*
 7 or PAD = *scalar-default-char-expr*
 8 or POSITION = *scalar-default-char-expr*
 9 or RECL = *scalar-int-expr*
 10 or ROUND = *scalar-default-char-expr*
 11 or SIGN = *scalar-default-char-expr*
 12 or STATUS = *scalar-default-char-expr*
 13 or TEAM = *image-team*
 14 R906 *file-name-expr* is *scalar-default-char-expr*
 15 R907 *iomsg-variable* is *scalar-default-char-variable*
- 16 C903 No specifier shall appear more than once in a given *connect-spec-list*.
- 17 C904 (R904) If the NEWUNIT= specifier does not appear, a *file-unit-number* shall be specified; if
 18 the optional characters UNIT= are omitted, the *file-unit-number* shall be the first item in the
 19 *connect-spec-list*.
- 20 C905 (R904) The *label* used in the ERR= specifier shall be the statement label of a branch target
 21 statement that appears in the same scoping unit as the OPEN statement.
- 22 C906 (R904) If a NEWUNIT= specifier appears, a *file-unit-number* shall not appear.
- 23 If the STATUS= specifier has the value NEW or REPLACE, the FILE= specifier shall appear. If the
 24 STATUS= specifier has the value SCRATCH, the FILE= specifier shall not appear. If the STATUS=
 25 specifier has the value OLD, the FILE= specifier shall appear unless the unit is connected and the file
 26 connected to the unit exists.
- 27 If the NEWUNIT= specifier appears in an OPEN statement, either the FILE= specifier shall appear,
 28 or the STATUS= specifier shall appear with a value of SCRATCH. The unit identified by a NEWUNIT
 29 value shall not be preconnected.
- 30 A specifier that requires a *scalar-default-char-expr* may have a limited list of character values. These
 31 values are listed for each such specifier. Any trailing blanks are ignored. The value specified is without
 32 regard to case. Some specifiers have a default value if the specifier is omitted.
- 33 The IOSTAT=, ERR=, and IOMSG= specifiers are described in 9.10.

NOTE 9.18

An example of an OPEN statement is:

```
OPEN (10, FILE = 'employee.names', ACTION = 'READ', PAD = 'YES')
```

NOTE 9.19

For more explanatory information on the OPEN statement, see C.6.4.

34 9.4.5.1 ACCESS= specifier in the OPEN statement

- 35 The *scalar-default-char-expr* shall evaluate to SEQUENTIAL, DIRECT, or STREAM. The ACCESS=
 36 specifier specifies the access method for the connection of the file as being sequential, direct, or stream.
 37 If this specifier is omitted, the default value is SEQUENTIAL. For an existing file, the specified access

1 method shall be included in the set of allowed access methods for the file. For a new file, the processor
2 creates the file with a set of allowed access methods that includes the specified method.

3 **9.4.5.2 ACTION= specifier in the OPEN statement**

4 The *scalar-default-char-expr* shall evaluate to READ, WRITE, or READWRITE. READ specifies that
5 the WRITE, PRINT, and ENDFILE statements shall not refer to this connection. WRITE specifies
6 that READ statements shall not refer to this connection. READWRITE permits any input/output
7 statements to refer to this connection. If this specifier is omitted, the default value is processor dependent.
8 If READWRITE is included in the set of allowable actions for a file, both READ and WRITE also shall
9 be included in the set of allowed actions for that file. For an existing file, the specified action shall be
10 included in the set of allowed actions for the file. For a new file, the processor creates the file with a set
11 of allowed actions that includes the specified action.

12 **9.4.5.3 ASYNCHRONOUS= specifier in the OPEN statement**

13 The *scalar-default-char-expr* shall evaluate to YES or NO. If YES is specified, asynchronous input/output
14 on the unit is allowed. If NO is specified, asynchronous input/output on the unit is not allowed. If this
15 specifier is omitted, the default value is NO.

16 **9.4.5.4 BLANK= specifier in the OPEN statement**

17 The *scalar-default-char-expr* shall evaluate to NULL or ZERO. The BLANK= specifier is permitted only
18 for a connection for formatted input/output. It specifies the current value of the blank interpretation
19 mode (10.8.6, 9.5.2.6) for input for this connection. This mode has no effect on output. It is a changeable
20 mode (9.4.1). If this specifier is omitted in an OPEN statement that initiates a connection, the default
21 value is NULL.

22 **9.4.5.5 DECIMAL= specifier in the OPEN statement**

23 The *scalar-default-char-expr* shall evaluate to COMMA or POINT. The DECIMAL= specifier is per-
24 mitted only for a connection for formatted input/output. It specifies the current value of the decimal
25 edit mode (10.6, 10.8.8, 9.5.2.7) for this connection. This is a changeable mode (9.4.1). If this specifier
26 is omitted in an OPEN statement that initiates a connection, the default value is POINT.

27 **9.4.5.6 DELIM= specifier in the OPEN statement**

28 The *scalar-default-char-expr* shall evaluate to APOSTROPHE, QUOTE, or NONE. The DELIM= spec-
29 ifier is permitted only for a connection for formatted input/output. It specifies the current value of the
30 delimiter mode (9.5.2.8) for list-directed (10.10.4) and namelist (10.11.4.1) output for the connection.
31 This mode has no effect on input. It is a changeable mode (9.4.1). If this specifier is omitted in an
32 OPEN statement that initiates a connection, the default value is NONE.

33 **9.4.5.7 ENCODING= specifier in the OPEN statement**

34 The *scalar-default-char-expr* shall evaluate to UTF-8 or DEFAULT. The ENCODING= specifier is
35 permitted only for a connection for formatted input/output. The value UTF-8 specifies that the encoding
36 form of the file is UTF-8 as specified by ISO/IEC 10646-1:2000. Such a file is called a **Unicode** file,
37 and all characters therein are of ISO 10646 character type. The value UTF-8 shall not be specified if
38 the processor does not support the ISO 10646 character type. The value DEFAULT specifies that the
39 encoding form of the file is processor-dependent. If this specifier is omitted in an OPEN statement that
40 initiates a connection, the default value is DEFAULT.

1 9.4.5.8 FILE= specifier in the OPEN statement

2 The value of the FILE= specifier is the name of the file to be connected to the specified unit. Any trailing
3 blanks are ignored. The *file-name-expr* shall be a name that is allowed by the processor. If this specifier
4 is omitted and the unit is not connected to a file, the STATUS= specifier shall be specified with a value
5 of SCRATCH; in this case, the connection is made to a processor-dependent file. The interpretation of
6 case is processor dependent.

7 9.4.5.9 FORM= specifier in the OPEN statement

8 The *scalar-default-char-expr* shall evaluate to FORMATTED or UNFORMATTED. The FORM= spec-
9 ifier determines whether the file is being connected for formatted or unformatted input/output. If this
10 specifier is omitted, the default value is UNFORMATTED if the file is being connected for direct access
11 or stream access, and the default value is FORMATTED if the file is being connected for sequential
12 access. For an existing file, the specified form shall be included in the set of allowed forms for the file.
13 For a new file, the processor creates the file with a set of allowed forms that includes the specified form.

14 9.4.5.10 NEWUNIT= specifier in the OPEN statement

15 The variable is defined with a processor determined NEWUNIT value if no error occurs during the
16 execution of the OPEN statement. If an error occurs, the processor shall not change the value of the
17 variable.

18 A NEWUNIT value is a negative number, and shall not be equal to -1, any of the named constants
19 ERROR_UNIT, INPUT_UNIT, or OUTPUT_UNIT from the ISO_FORTRAN_ENV intrinsic module
20 (13.8.3), any value used by the processor for the unit argument to a user-defined derived-type in-
21 put/output procedure, nor any previous NEWUNIT value that identifies a file that is currently con-
22 nected.

23 9.4.5.11 PAD= specifier in the OPEN statement

24 The *scalar-default-char-expr* shall evaluate to YES or NO. The PAD= specifier is permitted only for a
25 connection for formatted input/output. It specifies the current value of the pad mode (9.5.4.4.2, 9.5.2.10)
26 for input for this connection. This mode has no effect on output. It is a changeable mode (9.4.1). If this
27 specifier is omitted in an OPEN statement that initiates a connection, the default value is YES.

NOTE 9.20

For nondefault character types, the blank padding character is processor dependent.

28 9.4.5.12 POSITION= specifier in the OPEN statement

29 The *scalar-default-char-expr* shall evaluate to ASIS, REWIND, or APPEND. The connection shall be
30 for sequential or stream access. A new file is positioned at its initial point. REWIND positions an
31 existing file at its initial point. APPEND positions an existing file such that the endfile record is the
32 next record, if it has one. If an existing file does not have an endfile record, APPEND positions the
33 file at its terminal point. ASIS leaves the position unchanged if the file exists and already is connected.
34 ASIS leaves the position unspecified if the file exists but is not connected. If this specifier is omitted,
35 the default value is ASIS.

36 9.4.5.13 RECL= specifier in the OPEN statement

37 The value of the RECL= specifier shall be positive. It specifies the length of each record in a file being
38 connected for direct access, or specifies the maximum length of a record in a file being connected for
39 sequential access. This specifier shall not appear when a file is being connected for stream access. This
40 specifier shall appear when a file is being connected for direct access. If this specifier is omitted when

1 a file is being connected for sequential access, the default value is processor dependent. If the file is
 2 being connected for formatted input/output, the length is the number of characters for all records that
 3 contain only characters of type default character. When a record contains any nondefault characters,
 4 the appropriate value for the RECL= specifier is processor dependent. If the file is being connected for
 5 unformatted input/output, the length is measured in file storage units. For an existing file, the value of
 6 the RECL= specifier shall be included in the set of allowed record lengths for the file. For a new file,
 7 the processor creates the file with a set of allowed record lengths that includes the specified value.

8 **9.4.5.14 ROUND= specifier in the OPEN statement**

9 The *scalar-default-char-expr* shall evaluate to one of UP, DOWN, ZERO, NEAREST, COMPATIBLE,
 10 or PROCESSOR_DEFINED. The ROUND= specifier is permitted only for a connection for formatted
 11 input/output. It specifies the current value of the I/O rounding mode (10.7.2.3.7, 9.5.2.13) for this
 12 connection. This is a changeable mode (9.4.1). If this specifier is omitted in an OPEN statement that
 13 initiates a connection, the I/O rounding mode is processor dependent; it shall be one of the above modes.

NOTE 9.21

A processor is free to select any I/O rounding mode for the default mode. The mode might correspond to UP, DOWN, ZERO, NEAREST, or COMPATIBLE; or it might be a completely different I/O rounding mode.

14 **9.4.5.15 SIGN= specifier in the OPEN statement**

15 The *scalar-default-char-expr* shall evaluate to one of PLUS, SUPPRESS, or PROCESSOR_DEFINED.
 16 The SIGN= specifier is permitted only for a connection for formatted input/output. It specifies the
 17 current value of the sign mode (10.8.4, 9.5.2.14) for this connection. This is a changeable mode (9.4.1).
 18 If this specifier is omitted in an OPEN statement that initiates a connection, the default value is PRO-
 19 CESSOR_DEFINED.

20 **9.4.5.16 STATUS= specifier in the OPEN statement**

21 The *scalar-default-char-expr* shall evaluate to OLD, NEW, SCRATCH, REPLACE, or UNKNOWN. If
 22 OLD is specified, the file shall exist. If NEW is specified, the file shall not exist.

23 Successful execution of an OPEN statement with NEW specified creates the file and changes the status
 24 to OLD. If REPLACE is specified and the file does not already exist, the file is created and the status is
 25 changed to OLD. If REPLACE is specified and the file does exist, the file is deleted, a new file is created
 26 with the same name, and the status is changed to OLD. If SCRATCH is specified, the file is created
 27 and connected to the specified unit for use by the program but is deleted at the execution of a CLOSE
 28 statement referring to the same unit or at the normal termination of the program.

NOTE 9.22

SCRATCH shall not be specified with a named file.

29 If UNKNOWN is specified, the status is processor dependent. If this specifier is omitted, the default
 30 value is UNKNOWN.

31 **9.4.5.17 TEAM= specifier in the OPEN statement**

32 The *image-team* specifies the **connect team** for the unit, which is the set of images that are permitted
 33 to reference the unit. The team shall include the executing image. If there is no TEAM= specifier, the
 34 connect team consists of only the executing image.

35 All images in the connect team, and no others, shall execute the same OPEN statement with identical

- 1 values for the *connect-specs*. There is an implicit team synchronization.
- 2 If the OPEN statement has a STATUS= specifier with the value SCRATCH, the processor shall connect
3 the same file to the unit on all images in the connect team.
- 4 If the connect team contains more than one image, the OPEN statement shall
- 5 • specify direct access or
- 6 • specify sequential access and have an ACTION= specifier that evaluates to WRITE.

NOTE 9.23

Writing to a sequential file from more than one image without using synchronization is permitted, but is only useful for situations in which the ordering of records is unimportant. An example is for diagnostic output that is labeled with the image index.

- 7 Preconnected units and the units identified by the values INPUT_UNIT, OUTPUT_UNIT, and ERROR_-
8 UNIT in the intrinsic module ISO_FORTRAN_ENV have a connect team consisting of all the images. If
9 an image with index greater than one executes an input/output statement on one of these units, it shall
10 be a WRITE or PRINT statement.

J3 internal note

Unresolved Technical Issue 042

Prohibiting INQUIRE seems ... excessive and unwarranted to say the least. Subgroup agree, but there is still a question-mark over other i/o statements. (BACKSPACE, ENDFILE, WAIT, FLUSH, CLOSE, OPEN, REWIND).

J3 internal note

Unresolved Technical Issue 043

Since READ is prohibited except on image one, in what way does INPUT_UNIT have a connect team (since no-one can do anything)? Surely it would be simpler if its connect term were image 1 only? As implied by the following note...

NOTE 9.24

The input unit identified by * is therefore only available on the image with index one.

11 9.4.6 CLOSE statement

- 12 The **CLOSE statement** is used to terminate the connection of a specified unit to an external file.
- 13 Execution of a CLOSE statement for a unit may occur in any program unit of a program and need not
14 occur in the same program unit as the execution of an OPEN statement referring to that unit.
- 15 Execution of a CLOSE statement performs a wait operation for any pending asynchronous data transfer
16 operations for the specified unit.
- 17 Execution of a CLOSE statement specifying a unit that does not exist or has no file connected to it is
18 permitted and affects no file or unit.
- 19 After a unit has been disconnected by execution of a CLOSE statement, it may be connected again
20 within the same program, either to the same file or to a different file. After a named file has been
21 disconnected by execution of a CLOSE statement, it may be connected again within the same program,
22 either to the same unit or to a different unit, provided that the file still exists.

1 At normal termination of execution of a program, all units that are connected are closed. Each unit
 2 is closed with status KEEP unless the file status prior to termination of execution was SCRATCH, in
 3 which case the unit is closed with status DELETE.

NOTE 9.25

The effect is as though a CLOSE statement without a STATUS= specifier were executed on each connected unit.

4 R908 *close-stmt* is CLOSE (*close-spec-list*)
 5 R909 *close-spec* is [UNIT =] *file-unit-number*
 6 or IOSTAT = *scalar-int-variable*
 7 or IOMSG = *iomsg-variable*
 8 or ERR = *label*
 9 or STATUS = *scalar-default-char-expr*

10 C907 No specifier shall appear more than once in a given *close-spec-list*.

11 C908 A *file-unit-number* shall be specified in a *close-spec-list*; if the optional characters UNIT= are
 12 omitted, the *file-unit-number* shall be the first item in the *close-spec-list*.

13 C909 (R909) The *label* used in the ERR= specifier shall be the statement label of a branch target
 14 statement that appears in the same scoping unit as the CLOSE statement.

15 The *scalar-default-char-expr* has a limited list of character values. Any trailing blanks are ignored. The
 16 value specified is without regard to case.

17 The IOSTAT=, ERR=, and IOMSG= specifiers are described in 9.10.

NOTE 9.26

An example of a CLOSE statement is:

```
CLOSE (10, STATUS = 'KEEP')
```

NOTE 9.27

For more explanatory information on the CLOSE statement, see C.6.6.

18 9.4.6.1 STATUS= specifier in the CLOSE statement

19 The *scalar-default-char-expr* shall evaluate to KEEP or DELETE. The STATUS= specifier determines
 20 the disposition of the file that is connected to the specified unit. KEEP shall not be specified for a file
 21 whose status prior to execution of a CLOSE statement is SCRATCH. If KEEP is specified for a file that
 22 exists, the file continues to exist after the execution of a CLOSE statement. If KEEP is specified for a
 23 file that does not exist, the file will not exist after the execution of a CLOSE statement. If DELETE is
 24 specified, the file will not exist after the execution of a CLOSE statement. If this specifier is omitted, the
 25 default value is KEEP, unless the file status prior to execution of the CLOSE statement is SCRATCH,
 26 in which case the default value is DELETE.

27 9.5 Data transfer statements

28 9.5.1 General

29 The **READ statement** is the data transfer input statement. The **WRITE statement** and the
 30 **PRINT statement** are the data transfer output statements.

1	R910	<i>read-stmt</i>	is READ (<i>io-control-spec-list</i>) [<i>input-item-list</i>]
2			or READ <i>format</i> [, <i>input-item-list</i>]
3	R911	<i>write-stmt</i>	is WRITE (<i>io-control-spec-list</i>) [<i>output-item-list</i>]
4	R912	<i>print-stmt</i>	is PRINT <i>format</i> [, <i>output-item-list</i>]

NOTE 9.28

Examples of data transfer statements are:

```

READ (6, *) SIZE
READ 10, A, B
WRITE (6, 10) A, S, J
PRINT 10, A, S, J
10 FORMAT (2E16.3, I5)

```

5 9.5.2 Control information list

6 9.5.2.1 Syntax

7 A **control information list** is an *io-control-spec-list*. It governs data transfer.

8	R913	<i>io-control-spec</i>	is [UNIT =] <i>io-unit</i>
9			or [FMT =] <i>format</i>
10			or [NML =] <i>namelist-group-name</i>
11			or ADVANCE = <i>scalar-default-char-expr</i>
12			or ASYNCHRONOUS = <i>scalar-char-initialization-expr</i>
13			or BLANK = <i>scalar-default-char-expr</i>
14			or DECIMAL = <i>scalar-default-char-expr</i>
15			or DELIM = <i>scalar-default-char-expr</i>
16			or END = <i>label</i>
17			or EOR = <i>label</i>
18			or ERR = <i>label</i>
19			or ID = <i>scalar-int-variable</i>
20			or IOMSG = <i>iomsg-variable</i>
21			or IOSTAT = <i>scalar-int-variable</i>
22			or PAD = <i>scalar-default-char-expr</i>
23			or POS = <i>scalar-int-expr</i>
24			or REC = <i>scalar-int-expr</i>
25			or ROUND = <i>scalar-default-char-expr</i>
26			or SIGN = <i>scalar-default-char-expr</i>
27			or SIZE = <i>scalar-int-variable</i>

28 C910 No specifier shall appear more than once in a given *io-control-spec-list*.

29 C911 An *io-unit* shall be specified in an *io-control-spec-list*; if the optional characters UNIT= are
30 omitted, the *io-unit* shall be the first item in the *io-control-spec-list*.

31 C912 (R913) A DELIM= or SIGN= specifier shall not appear in a *read-stmt*.

32 C913 (R913) A BLANK=, PAD=, END=, EOR=, or SIZE= specifier shall not appear in a *write-stmt*.

33 C914 (R913) The *label* in the ERR=, EOR=, or END= specifier shall be the statement label of a
34 branch target statement that appears in the same scoping unit as the data transfer statement.

35 C915 (R913) A *namelist-group-name* shall be the name of a namelist group.

- 1 C916 (R913) A *namelist-group-name* shall not appear if an *input-item-list* or an *output-item-list*
2 appears in the data transfer statement.
- 3 C917 (R913) An *io-control-spec-list* shall not contain both a *format* and a *namelist-group-name*.
- 4 C918 (R913) If *format* appears without a preceding FMT=, it shall be the second item in the *io-*
5 *control-spec-list* and the first item shall be *io-unit*.
- 6 C919 (R913) If *namelist-group-name* appears without a preceding NML=, it shall be the second item
7 in the *io-control-spec-list* and the first item shall be *io-unit*.
- 8 C920 (R913) If *io-unit* is not a *file-unit-number*, the *io-control-spec-list* shall not contain a REC=
9 specifier or a POS= specifier.
- 10 C921 (R913) If the REC= specifier appears, an END= specifier shall not appear, a *namelist-group-*
11 *name* shall not appear, and the *format*, if any, shall not be an asterisk.
- 12 C922 (R913) An ADVANCE= specifier may appear only in a formatted sequential or stream in-
13 put/output statement with explicit format specification (10.2) whose control information list
14 does not contain an *internal-file-variable* as the *io-unit*.
- 15 C923 (R913) If an EOR= specifier appears, an ADVANCE= specifier also shall appear.
- 16 C924 (R913) If a SIZE= specifier appears, an ADVANCE= specifier also shall appear.
- 17 C925 (R913) The *scalar-char-initialization-expr* in an ASYNCHRONOUS= specifier shall be of type
18 default character and shall have the value YES or NO.
- 19 C926 (R913) An ASYNCHRONOUS= specifier with a value YES shall not appear unless *io-unit* is a
20 *file-unit-number*.
- 21 C927 (R913) If an ID= specifier appears, an ASYNCHRONOUS= specifier with the value YES shall
22 also appear.
- 23 C928 (R913) If a POS= specifier appears, the *io-control-spec-list* shall not contain a REC= specifier.
- 24 C929 (R913) If a DECIMAL=, BLANK=, PAD=, SIGN=, or ROUND= specifier appears, a *format*
25 or *namelist-group-name* shall also appear.
- 26 C930 (R913) If a DELIM= specifier appears, either *format* shall be an asterisk or *namelist-group-name*
27 shall appear.
- 28 A SIZE= specifier may appear only in an input statement that contains an ADVANCE= specifier with
29 the value NO.
- 30 An EOR= specifier may appear only in an input statement that contains an ADVANCE= specifier with
31 the value NO.
- 32 If the data transfer statement contains a *format* or *namelist-group-name*, the statement is a **formatted**
33 **input/output statement**; otherwise, it is an **unformatted input/output statement**.
- 34 The ADVANCE=, ASYNCHRONOUS=, DECIMAL=, BLANK=, DELIM=, PAD=, SIGN=, and
35 ROUND= specifiers have a limited list of character values. Any trailing blanks are ignored. The
36 values specified are without regard to case.
- 37 The IOSTAT=, ERR=, EOR=, END=, and IOMSG= specifiers are described in 9.10.

NOTE 9.29

An example of a READ statement is:

```
READ (IOSTAT = IOS, UNIT = 6, FMT = '(10F8.2)') A, B
```

1 **9.5.2.2 Format specification in a data transfer statement**

2 The *format* specifier supplies a format specification or specifies list-directed formatting for a formatted
3 input/output statement.

4 R914 *format* **is** *default-char-expr*
5 **or** *label*
6 **or** *

7 C931 (R914) The *label* shall be the label of a FORMAT statement that appears in the same scoping
8 unit as the statement containing the FMT= specifier.

9 The *default-char-expr* shall evaluate to a valid format specification (10.2.1 and 10.2.2).

NOTE 9.30

A *default-char-expr* includes a character constant.

10 If *default-char-expr* is an array, it is treated as if all of the elements of the array were specified in array
11 element order and were concatenated.

12 If *format* is *, the statement is a **list-directed input/output statement**.

NOTE 9.31

An example in which the format is a character expression is:

```
READ (6, FMT = "(" // CHAR_FMT // ")") X, Y, Z
```

where CHAR_FMT is a default character variable.

13 **9.5.2.3 NML= specifier in a data transfer statement**

14 The NML= specifier supplies the *namelist-group-name* (5.6). This name identifies a particular collection
15 of data objects on which transfer is to be performed.

16 If a *namelist-group-name* appears, the statement is a **namelist input/output statement**.

17 **9.5.2.4 ADVANCE= specifier in a data transfer statement**

18 The *scalar-default-char-expr* shall evaluate to YES or NO. The ADVANCE= specifier determines wheth-
19 er advancing input/output occurs for a nonchild input/output statement. If YES is specified for a
20 nonchild input/output statement, advancing input/output occurs. If NO is specified, nonadvancing in-
21 put/output occurs (9.2.3.1). If this specifier is omitted from a nonchild input/output statement that
22 allows the specifier, the default value is YES. A formatted child input/output statement is a nonadvanc-
23 ing input/output statement, and any ADVANCE= specifier is ignored.

24 **9.5.2.5 ASYNCHRONOUS= specifier in a data transfer statement**

25 The ASYNCHRONOUS= specifier determines whether this input/output statement is synchronous or
26 asynchronous. If YES is specified, the statement and the input/output operation are **asynchronous**.

- 1 If NO is specified or if the specifier is omitted, the statement and the input/output operation are
 2 **synchronous**.
- 3 Asynchronous input/output is permitted only for external files opened with an ASYNCHRONOUS=
 4 specifier with the value YES in the OPEN statement.

NOTE 9.32

Both synchronous and asynchronous input/output are allowed for files opened with an ASYNCHRONOUS= specifier of YES. For other files, only synchronous input/output is allowed; this includes files opened with an ASYNCHRONOUS= specifier of NO, files opened without an ASYNCHRONOUS= specifier, preconnected files accessed without an OPEN statement, and internal files.

The ASYNCHRONOUS= specifier value in a data transfer statement is an initialization expression because it effects compiler optimizations and, therefore, needs to be known at compile time.

- 5 The processor may perform an asynchronous data transfer operation asynchronously, but it is not re-
 6 quired to do so. For each external file, records and file storage units read or written by asynchronous
 7 data transfer statements are read, written, and processed in the same order as they would have been if
 8 the data transfer statements were synchronous.

- 9 If a variable is used in an asynchronous data transfer statement as

- 10 (1) an item in an input/output list,
 11 (2) a group object in a namelist, or
 12 (3) a SIZE= specifier

- 13 the base object of the *data-ref* is implicitly given the ASYNCHRONOUS attribute in the scoping unit
 14 of the data transfer statement. This attribute may be confirmed by explicit declaration.

- 15 When an asynchronous input/output statement is executed, the set of storage units specified by the
 16 item list or NML= specifier, plus the storage units specified by the SIZE= specifier, is defined to be the
 17 pending input/output storage sequence for the data transfer operation.

NOTE 9.33

A pending input/output storage sequence is not necessarily a contiguous set of storage units.

- 18 A pending input/output storage sequence **affector** is a variable of which any part is associated with a
 19 storage unit in a pending input/output storage sequence.

9.5.2.6 BLANK= specifier in a data transfer statement

- 21 The *scalar-default-char-expr* shall evaluate to NULL or ZERO. The BLANK= specifier temporarily
 22 changes (9.4.1) the blank interpretation mode (10.8.6, 9.4.5.4) for the connection. If the specifier is
 23 omitted, the mode is not changed.

9.5.2.7 DECIMAL= specifier in a data transfer statement

- 25 The *scalar-default-char-expr* shall evaluate to COMMA or POINT. The DECIMAL= specifier temporarily
 26 changes (9.4.1) the decimal edit mode (10.6, 10.8.8, 9.4.5.5) for the connection. If the specifier is
 27 omitted, the mode is not changed.

9.5.2.8 DELIM= specifier in a data transfer statement

- 28

1 The *scalar-default-char-expr* shall evaluate to APOSTROPHE, QUOTE, or NONE. The DELIM= spec-
2 ifier temporarily changes (9.4.1) the delimiter mode (10.10.4, 10.11.4.1, 9.4.5.6) for the connection. If
3 the specifier is omitted, the mode is not changed.

4 **9.5.2.9 ID= specifier in a data transfer statement**

5 Successful execution of an asynchronous data transfer statement containing an ID= specifier causes the
6 variable specified in the ID= specifier to become defined with a processor-dependent value. This value
7 is referred to as the identifier of the data transfer operation. It can be used in a subsequent WAIT or
8 INQUIRE statement to identify the particular data transfer operation.

9 If an error occurs during the execution of a data transfer statement containing an ID= specifier, the
10 variable specified in the ID= specifier becomes undefined.

11 A child data transfer statement shall not specify the ID= specifier.

12 **9.5.2.10 PAD= specifier in a data transfer statement**

13 The *scalar-default-char-expr* shall evaluate to YES or NO. The PAD= specifier temporarily changes
14 (9.4.1) the pad mode (9.5.4.4.2, 9.4.5.11) for the connection. If the specifier is omitted, the mode is not
15 changed.

16 **9.5.2.11 POS= specifier in a data transfer statement**

17 The POS= specifier specifies the file position in file storage units. This specifier may appear in a data
18 transfer statement only if the statement specifies a unit connected for stream access. A child data
19 transfer statement shall not specify this specifier.

20 A processor may prohibit the use of POS= with particular files that do not have the properties necessary
21 to support random positioning. A processor may also prohibit positioning a particular file to any
22 position prior to its current file position if the file does not have the properties necessary to support such
23 positioning.

NOTE 9.34

A file that represents connection to a device or data stream might not be positionable.

24 If the file is connected for formatted stream access, the file position specified by POS= shall be equal to
25 either 1 (the beginning of the file) or a value previously returned by a POS= specifier in an INQUIRE
26 statement for the file.

27 **9.5.2.12 REC= specifier in a data transfer statement**

28 The REC= specifier specifies the number of the record that is to be read or written. This specifier
29 may appear only in an input/output statement that specifies a unit connected for direct access; it
30 shall not appear in a child data transfer statement. If the control information list contains a REC=
31 specifier, the statement is a **direct access input/output statement**. A child data transfer statement
32 is a direct access data transfer statement if the parent is a direct access data transfer statement. Any
33 other data transfer statement is a **sequential access input/output statement** or a **stream access**
34 **input/output statement**, depending on whether the file connection is for sequential access or stream
35 access.

36 **9.5.2.13 ROUND= specifier in a data transfer statement**

37 The *scalar-default-char-expr* shall evaluate to one of the values specified in 9.4.5.14. The ROUND=
38 specifier temporarily changes (9.4.1) the I/O rounding mode (10.7.2.3.7, 9.4.5.14) for the connection. If

1 the specifier is omitted, the mode is not changed.

2 9.5.2.14 SIGN= specifier in a data transfer statement

3 The *scalar-default-char-expr* shall evaluate to PLUS, SUPPRESS, or PROCESSOR_DEFINED. The
4 SIGN= specifier temporarily changes (9.4.1) the sign mode (10.8.4, 9.4.5.15) for the connection. If the
5 specifier is omitted, the mode is not changed.

6 9.5.2.15 SIZE= specifier in a data transfer statement

7 When a synchronous nonadvancing input statement terminates, the variable specified in the SIZE=
8 specifier becomes defined with the count of the characters transferred by data edit descriptors during
9 execution of the current input statement. Blanks inserted as padding (9.5.4.4.2) are not counted.

10 For asynchronous nonadvancing input, the storage units specified in the SIZE= specifier become defined
11 with the count of the characters transferred when the corresponding wait operation is executed.

12 9.5.3 Data transfer input/output list

13 An input/output list specifies the entities whose values are transferred by a data transfer input/output
14 statement.

15	R915	<i>input-item</i>	is	<i>variable</i>
16			or	<i>io-implied-do</i>
17	R916	<i>output-item</i>	is	<i>expr</i>
18			or	<i>io-implied-do</i>
19	R917	<i>io-implied-do</i>	is	(<i>io-implied-do-object-list</i> , <i>io-implied-do-control</i>)
20	R918	<i>io-implied-do-object</i>	is	<i>input-item</i>
21			or	<i>output-item</i>
22	R919	<i>io-implied-do-control</i>	is	<i>do-variable</i> = <i>scalar-int-expr</i> , ■
23				■ <i>scalar-int-expr</i> [, <i>scalar-int-expr</i>]

24 C932 (R915) A variable that is an *input-item* shall not be a whole assumed-size array.

25 C933 (R915) A variable that is an *input-item* shall not be a procedure pointer.

26 C934 (R919) The *do-variable* shall be a named scalar variable of type integer.

27 C935 (R918) In an *input-item-list*, an *io-implied-do-object* shall be an *input-item*. In an *output-item-*
28 *list*, an *io-implied-do-object* shall be an *output-item*.

29 C936 (R916) An expression that is an *output-item* shall not have a value that is a procedure pointer.

30 An *input-item* shall not appear as, nor be associated with, the *do-variable* of any *io-implied-do* that
31 contains the *input-item*.

NOTE 9.35

A constant, an expression involving operators or function references that does not have a pointer result, or an expression enclosed in parentheses shall not appear as an input list item.

32 If an input item is a pointer, it shall be associated with a definable target and data are transferred from
33 the file to the associated target. If an output item is a pointer, it shall be associated with a target and
34 data are transferred from the target to the file.

NOTE 9.36

Data transfers always involve the movement of values between a file and internal storage. A pointer as such cannot be read or written. Therefore, a pointer shall not appear as an item in an input/output list unless it is associated with a target that can receive a value (input) or can deliver a value (output).

- 1 If an input item or an output item is allocatable, it shall be allocated.
- 2 A list item shall not be polymorphic unless it is processed by a user-defined derived-type input/output
3 procedure (9.5.4.7).
- 4 The *do-variable* of an *io-implied-do* that is in another *io-implied-do* shall not appear as, nor be associated
5 with, the *do-variable* of the containing *io-implied-do*.
- 6 The following rules describing whether to expand an input/output list item are re-applied to each
7 expanded list item until none of the rules apply.
- 8 • If an array appears as an input/output list item, it is treated as if the elements, if any, were
9 specified in array element order (6.2.2.2). However, no element of that array may affect the value
10 of any expression in the *input-item*, nor may any element appear more than once in an *input-item*.

NOTE 9.37

For example:

```

INTEGER A (100), J (100)
...
READ *, A (A)                ! Not allowed
READ *, A (LBOUND (A, 1) : UBOUND (A, 1)) ! Allowed
READ *, A (J)                ! Allowed if no two elements
                                !   of J have the same value
A(1) = 1; A(10) = 10
READ *, A (A (1) : A (10))   ! Not allowed

```

- 11 • If a list item of derived type in an unformatted input/output statement is not processed by a
12 user-defined derived-type input/output procedure (9.5.4.7), and if any subobject of that list item
13 would be processed by a user-defined derived-type input/output procedure, the list item is treated
14 as if all of the components of the object were specified in the list in component order (4.5.4.6);
15 those components shall be accessible in the scoping unit containing the input/output statement
16 and shall not be pointers or allocatable.
- 17 • An effective input/output list item of derived type in an unformatted input/output statement is
18 treated as a single value in a processor-dependent form unless the list item or a subobject thereof
19 is processed by a user-defined derived-type input/output procedure (9.5.4.7).

NOTE 9.38

The appearance of a derived-type object as an input/output list item in an unformatted input/output statement is not equivalent to the list of its components.

Unformatted input/output involving derived-type list items forms the single exception to the rule that the appearance of an aggregate list item (such as an array) is equivalent to the appearance of its expanded list of component parts. This exception permits the processor greater latitude in improving efficiency or in matching the processor-dependent sequence of values for a derived-type object to similar sequences for aggregate objects used by means other than Fortran. However,

NOTE 9.38 (cont.)

formatted input/output of all list items and unformatted input/output of list items other than those of derived types adhere to the above rule.

- 1 • If a list item of derived type in a formatted input/output statement is not processed by a user-
 2 defined derived-type input/output procedure, that list item is treated as if all of the components
 3 of the list item were specified in the list in component order; those components shall be accessible
 4 in the scoping unit containing the input/output statement and shall not be pointers or allocatable.
- 5 • If a derived-type list item is not treated as a list of its individual components, that list item's
 6 ultimate components shall not have the POINTER or ALLOCATABLE attribute unless that list
 7 item is processed by a user-defined derived-type input/output procedure.
- 8 • The scalar objects resulting when a data transfer statement's list items are expanded according
 9 to the rules in this subclause for handling array and derived-type list items are called **effective**
 10 **items**. Zero-sized arrays and *io-implied-dos* with an iteration count of zero do not contribute to
 11 the effective list items. A scalar character item of zero length is an effective list item.

NOTE 9.39

In a formatted input/output statement, edit descriptors are associated with effective list items, which are always scalar. The rules in 9.5.3 determine the set of effective list items corresponding to each actual list item in the statement. These rules might have to be applied repetitively until all of the effective list items are scalar items.

- 12 • For an *io-implied-do*, the loop initialization and execution are the same as for a DO construct
 13 (8.1.7.5).

NOTE 9.40

An example of an output list with an implied DO is:

```
WRITE (LP, FMT = '(10F8.2)') (LOG (A (I)), I = 1, N + 9, K), G
```

- 14 • An input/output list shall not contain an item of nondefault character type if the input/output
 15 statement specifies an internal file of default character type. An input/output list shall not con-
 16 tain an item of nondefault character type other than ISO 10646 or ASCII character type if the
 17 input/output statement specifies an internal file of ISO 10646 character type. An input/output
 18 list shall not contain a character item of any character type other than ASCII character type if
 19 the input/output statement specifies an internal file of ASCII character type.

20 9.5.4 Execution of a data transfer input/output statement

21 Execution of a WRITE or PRINT statement for a file that does not exist creates the file unless an error
 22 condition occurs.

23 The effect of executing a synchronous data transfer input/output statement shall be as if the following
 24 operations were performed in the order specified.

- 25 (1) Determine the direction of data transfer.
 26 (2) Identify the unit.
 27 (3) Perform a wait operation for all pending input/output operations for the unit. If an error,
 28 end-of-file, or end-of-record condition occurs during any of the wait operations, steps 4
 29 through 8 are skipped for the current data transfer statement.
 30 (4) Establish the format if one is specified.

- 1 (5) If the statement is not a child data transfer statement (9.5.4.7),
 2 (a) position the file prior to data transfer (9.2.3.2), and
 3 (b) for formatted data transfer, set the left tab limit (10.8.1.1).
 4 (6) Transfer data between the file and the entities specified by the input/output list (if any) or
 5 namelist.
 6 (7) Determine whether an error, end-of-file, or end-of-record condition has occurred.
 7 (8) Position the file after data transfer (9.2.3.3) unless the statement is a child data transfer
 8 statement (9.5.4.7).
 9 (9) Cause any variable specified in a SIZE= specifier to become defined.
 10 (10) If an error, end-of-file, or end-of-record condition occurred, processing continues as specified
 11 in 9.10; otherwise any variable specified in an IOSTAT= specifier is assigned the value zero.

12 The effect of executing an asynchronous data transfer input/output statement shall be as if the following
 13 operations were performed in the order specified.

- 14 (1) Determine the direction of data transfer.
 15 (2) Identify the unit.
 16 (3) Establish the format if one is specified.
 17 (4) Position the file prior to data transfer (9.2.3.2) and, for formatted data transfer, set the left
 18 tab limit (10.8.1.1).
 19 (5) Establish the set of storage units identified by the input/output list. For a READ statement,
 20 this might require some or all of the data in the file to be read if an input variable is used
 21 as a *scalar-int-expr* in an *io-implied-do-control* in the input/output list, as a *subscript*,
 22 *substring-range*, *stride*, or is otherwise referenced.
 23 (6) Initiate an asynchronous data transfer between the file and the entities specified by the
 24 input/output list (if any) or namelist. The asynchronous data transfer may complete (and
 25 an error, end-of-file, or end-of-record condition may occur) during the execution of this data
 26 transfer statement or during a later wait operation.
 27 (7) Determine whether an error, end-of-file, or end-of-record condition has occurred. The con-
 28 ditions may occur during the execution of this data transfer statement or during the corre-
 29 sponding wait operation, but not both.
 30 Also, any of these conditions that would have occurred during the corresponding wait oper-
 31 ation for a previously pending data transfer operation that does not have an ID= specifier
 32 may occur during the execution of this data transfer statement.

J3 internal note

Unresolved Technical Issue 091

Contradiction; two paragraphs earlier, we say it occurs during the execution of the original i/o statement or “during a later wait operation”. There is no wait operation in an async i/o statement. The best fix for this is to probably add an additional action to the list: “Optionally, perform a wait operation for a pending input/output operation for this unit.”

Also, note the difference in specification between this item (para 1) and the text four paragraphs after the list on wait operations. These should be consistent, and preferably specified only once.

- 33 (8) Position the file as if the data transfer had finished (9.2.3.3).
 34 (9) Cause any variable specified in a SIZE= specifier to become undefined.
 35 (10) If an error, end-of-file, or end-of-record condition occurred, processing continues as specified
 36 in 9.10; otherwise any variable specified in an IOSTAT= specifier is assigned the value zero.

37 For an asynchronous data transfer statement, the data transfers may occur during execution of the
 38 statement, during execution of the corresponding wait operation, or anywhere between. The data transfer

- 1 operation is considered to be pending until a corresponding wait operation is performed.
- 2 For asynchronous output, a pending input/output storage sequence affector (9.5.2.5) shall not be re-
3 fined, become undefined, or have its pointer association status changed.
- 4 For asynchronous input, a pending input/output storage sequence affector shall not be referenced, be-
5 come defined, become undefined, become associated with a dummy argument that has the VALUE
6 attribute, or have its pointer association status changed.
- 7 Error, end-of-file, and end-of-record conditions in an asynchronous data transfer operation may occur
8 during execution of either the data transfer statement or the corresponding wait operation. If an ID=
9 specifier does not appear in the initiating data transfer statement, the conditions may occur during the
10 execution of any subsequent data transfer or wait operation for the same unit. When a condition occurs
11 for a previously executed asynchronous data transfer statement, a wait operation is performed for all
12 pending data transfer operations on that unit. When a condition occurs during a subsequent statement,
13 any actions specified by IOSTAT=, IOMSG=, ERR=, END=, and EOR= specifiers for that statement
14 are taken.

NOTE 9.41

Because end-of-file and error conditions for asynchronous data transfer statements without an ID= specifier may be reported by the processor during the execution of a subsequent data transfer statement, it may be impossible for the user to determine which input/output statement caused the condition. Reliably detecting which READ statement caused an end-of-file condition requires that all asynchronous READ statements for the unit include an ID= specifier.

15 9.5.4.1 Direction of data transfer

- 16 Execution of a READ statement causes values to be transferred from a file to the entities specified by
17 the input list, if any, or specified within the file itself for namelist input. Execution of a WRITE or
18 PRINT statement causes values to be transferred to a file from the entities specified by the output list
19 and format specification, if any, or by the *namelist-group-name* for namelist output.

20 9.5.4.2 Identifying a unit

- 21 A data transfer input/output statement that contains an input/output control list includes a UNIT=
22 specifier that identifies an external or internal unit. A READ statement that does not contain an
23 input/output control list specifies a particular processor-dependent unit, which is the same as the unit
24 identified by * in a READ statement that contains an input/output control list. The PRINT statement
25 specifies some other processor-dependent unit, which is the same as the unit identified by * in a WRITE
26 statement. Thus, each data transfer input/output statement identifies an external or internal unit.
- 27 The unit identified by an unformatted data transfer statement shall be an external unit.
- 28 The unit identified by a data transfer input/output statement shall be connected to a file when execution
29 of the statement begins.

NOTE 9.42

The unit may be preconnected.

30 9.5.4.3 Establishing a format

- 31 If the input/output control list contains * as a format, list-directed formatting is established. If *namelist-*
32 *group-name* appears, namelist formatting is established. If no *format* or *namelist-group-name* is speci-
33 fied, unformatted data transfer is established. Otherwise, the format specified by *format* is established.

1 For output to an internal file, a format specification that is in the file or is associated with the file shall
2 not be specified.

3 **9.5.4.4 Data transfer**

4 Data are transferred between the file and the entities specified by the input/output list or namelist.
5 The list items are processed in the order of the input/output list for all data transfer input/output
6 statements except namelist formatted data transfer statements. The list items for a namelist input
7 statement are processed in the order of the entities specified within the input records. The list items
8 for a namelist output statement are processed in the order in which the variables are specified in the
9 *namelist-group-object-list*. Effective items are derived from the input/output list items as described in
10 9.5.3.

11 All values needed to determine which entities are specified by an input/output list item are determined
12 at the beginning of the processing of that item.

13 All values are transmitted to or from the entities specified by a list item prior to the processing of any
14 succeeding list item for all data transfer input/output statements.

NOTE 9.43

In the example,

```
READ (N) N, X (N)
```

the old value of N identifies the unit, but the new value of N is the subscript of X.

15 All values following the *name=* part of the namelist entity (10.11) within the input records are transmit-
16 ted to the matching entity specified in the *namelist-group-object-list* prior to processing any succeeding
17 entity within the input record for namelist input statements. If an entity is specified more than once
18 within the input record during a namelist formatted data transfer input statement, the last occurrence
19 of the entity specifies the value or values to be used for that entity.

20 An input list item, or an entity associated with it, shall not contain any portion of an established format
21 specification.

22 If the input/output item is a pointer, data are transferred between the file and the associated target.

23 If an internal file has been specified, an input/output list item shall not be in the file or associated with
24 the file.

NOTE 9.44

The file is a data object.

25 A DO variable becomes defined and its iteration count established at the beginning of processing of the
26 items that constitute the range of an *io-implied-do*.

27 On output, every entity whose value is to be transferred shall be defined.

28 **9.5.4.4.1 Unformatted data transfer**

29 During unformatted data transfer, data are transferred without editing between the file and the entities
30 specified by the input/output list. If the file is connected for sequential or direct access, exactly one
31 record is read or written.

32 A value in the file is stored in a contiguous sequence of file storage units, beginning with the file storage

- 1 unit immediately following the current file position.
- 2 After each value is transferred, the current file position is moved to a point immediately after the last
3 file storage unit of the value.
- 4 On input from a file connected for sequential or direct access, the number of file storage units required
5 by the input list shall be less than or equal to the number of file storage units in the record.
- 6 On input, if the file storage units transferred do not contain a value with the same type and type
7 parameters as the input list entity, then the resulting value of the entity is processor-dependent except
8 in the following cases.
- 9 (1) A complex entity may correspond to two real values with the same kind type parameter as
10 the complex entity.
- 11 (2) A default character list entity of length n may correspond to n default characters stored in
12 the file, regardless of the length parameters of the entities that were written to these storage
13 units of the file. If the file is connected for stream input, the characters may have been
14 written by formatted stream output.
- 15 On output to a file connected for unformatted direct access, the output list shall not specify more values
16 than can fit into the record. If the file is connected for direct access and the values specified by the
17 output list do not fill the record, the remainder of the record is undefined.
- 18 If the file is connected for unformatted sequential access, the record is created with a length sufficient
19 to hold the values from the output list. This length shall be one of the set of allowed record lengths for
20 the file and shall not exceed the value specified in the RECL= specifier, if any, of the OPEN statement
21 that established the connection.
- 22 If the file is not connected for unformatted input/output, unformatted data transfer is prohibited.
- 23 **9.5.4.4.2 Formatted data transfer**
- 24 During formatted data transfer, data are transferred with editing between the file and the entities
25 specified by the input/output list or by the *namelist-group-name*. Format control is initiated and editing
26 is performed as described in Clause 10.
- 27 The current record and possibly additional records are read or written.
- 28 If the file is not connected for formatted input/output, formatted data transfer is prohibited.
- 29 During advancing input when the pad mode has the value NO, the input list and format specification
30 shall not require more characters from the record than the record contains.
- 31 During advancing input when the pad mode has the value YES, blank characters are supplied by the
32 processor if the input list and format specification require more characters from the record than the
33 record contains.
- 34 During nonadvancing input when the pad mode has the value NO, an end-of-record condition (9.10)
35 occurs if the input list and format specification require more characters from the record than the record
36 contains, and the record is complete (9.2.2.3). If the record is incomplete, an end-of-file condition occurs
37 instead of an end-of-record condition.
- 38 During nonadvancing input when the pad mode has the value YES, blank characters are supplied by
39 the processor if an effective item and its corresponding data edit descriptors require more characters
40 from the record than the record contains. If the record is incomplete, an end-of-file condition occurs;
41 otherwise an end-of-record condition occurs.
- 42 If the file is connected for direct access, the record number is increased by one as each succeeding record

1 is read or written.

2 On output, if the file is connected for direct access or is an internal file and the characters specified by
3 the output list and format do not fill a record, blank characters are added to fill the record.

4 On output, the output list and format specification shall not specify more characters for a record than
5 have been specified by a RECL= specifier in the OPEN statement or the record length of an internal
6 file.

7 **9.5.4.5 List-directed formatting**

8 If list-directed formatting has been established, editing is performed as described in 10.10.

9 **9.5.4.6 Namelist formatting**

10 If namelist formatting has been established, editing is performed as described in 10.11.

11 Every allocatable *namelist-group-object* in the namelist group shall be allocated and every *namelist-*
12 *group-object* that is a pointer shall be associated with a target. If a *namelist-group-object* is polymorphic
13 or has an ultimate component that is allocatable or a pointer, that object shall be processed by a user-
14 defined derived-type input/output procedure (9.5.4.7).

15 **9.5.4.7 User-defined derived-type input/output**

16 User-defined derived-type input/output procedures allow a program to override the default handling of
17 derived-type objects and values in data transfer input/output statements described in 9.5.3.

18 A user-defined derived-type input/output procedure is a procedure accessible by a *dtio-generic-spec*
19 (12.4.3.2). A particular user-defined derived-type input/output procedure is selected as described in
20 9.5.4.7.3.

21 **9.5.4.7.1 Executing user-defined derived-type input/output data transfers**

22 If a derived-type input/output procedure is selected as specified in 9.5.4.7.3, the processor shall call the se-
23 lected user-defined derived-type input/output procedure for any appropriate data transfer input/output
24 statements executed in that scoping unit. The user-defined derived-type input/output procedure controls
25 the actual data transfer operations for the derived-type list item.

26 A data transfer statement that includes a derived-type list item and that causes a user-defined derived-
27 type input/output procedure to be invoked is called a **parent data transfer statement**. A data
28 transfer statement that is executed while a parent data transfer statement is being processed and that
29 specifies the unit passed into a user-defined derived-type input/output procedure is called a **child data**
30 **transfer statement**.

NOTE 9.45

A user-defined derived-type input/output procedure will usually contain child data transfer state-
ments that read values from or write values to the current record or at the current file position.
The effect of executing the user-defined derived-type input/output procedure is similar to that of
substituting the list items from any child data transfer statements into the parent data transfer
statement's list items, along with similar substitutions in the format specification.

NOTE 9.46

A particular execution of a READ, WRITE or PRINT statement can be both a parent and a
child data transfer statement. A user-defined derived-type input/output procedure can indirectly
call itself or another user-defined derived-type input/output procedure by executing a child data

NOTE 9.46 (cont.)

transfer statement containing a list item of derived type, where a matching interface is accessible for that derived type. If a user-defined derived-type input/output procedure calls itself indirectly in this manner, it shall be declared RECURSIVE.

1 A child data transfer statement is processed differently from a nonchild data transfer statement in the
2 following ways.

- 3 • Executing a child data transfer statement does not position the file prior to data transfer.
- 4 • An unformatted child data transfer statement does not position the file after data transfer is
5 complete.
- 6 • Any ADVANCE= specifier in a child input/output statement is ignored.

7 **9.5.4.7.2 User-defined derived-type input/output procedures**

8 For a particular derived type and a particular set of kind type parameter values, there are four possible
9 sets of characteristics for user-defined derived-type input/output procedures; one each for formatted
10 input, formatted output, unformatted input, and unformatted output. The user need not supply all four
11 procedures. The procedures are specified to be used for derived-type input/output by interface blocks
12 (12.4.3.2) or by generic bindings (4.5.5), with a *dtio-generic-spec* (R1208).

13 In the four interfaces, which specify the characteristics of user-defined procedures for derived-type in-
14 put/output, the following syntax term is used:

15 R920 *dtv-type-spec* **is** TYPE(*derived-type-spec*)
16 **or** CLASS(*derived-type-spec*)

17 C937 (R920) If *derived-type-spec* specifies an extensible type, the CLASS keyword shall be used;
18 otherwise, the TYPE keyword shall be used.

19 C938 (R920) All length type parameters of *derived-type-spec* shall be assumed.

20 If the *dtio-generic-spec* is READ (FORMATTED), the characteristics shall be the same as those specified
21 by the following interface:

```

22           SUBROUTINE my_read_routine_formatted                           &
23                                                                            &
24                                                                            &
25                                                                            &
26                                                                            &
27           ! the derived-type variable
28           dtv-type-spec, INTENT(INOUT) :: dtv
29           INTEGER, INTENT(IN) :: unit ! unit number
30           ! the edit descriptor string
31           CHARACTER (LEN=*), INTENT(IN) :: iotype
32           INTEGER, INTENT(IN) :: v_list(:)
33           INTEGER, INTENT(OUT) :: iostat
34           CHARACTER (LEN=*), INTENT(INOUT) :: iomsg
35           END

```

36 If the *dtio-generic-spec* is READ (UNFORMATTED), the characteristics shall be the same as those
37 specified by the following interface:

```

1      SUBROUTINE my_read_routine_unformatted      &
2          (dtv,                                  &
3           unit,                                  &
4           iostat, iomsg)
5      ! the derived-type variable
6      dtv-type-spec, INTENT(INOUT) :: dtv
7      INTEGER, INTENT(IN) :: unit
8      INTEGER, INTENT(OUT) :: iostat
9      CHARACTER (LEN=*), INTENT(INOUT) :: iomsg
10     END

```

11 If the *dtio-generic-spec* is WRITE (FORMATTED), the characteristics shall be the same as those specified by the following interface:

```

13     SUBROUTINE my_write_routine_formatted      &
14         (dtv,                                  &
15          unit,                                  &
16          iotype, v_list,                       &
17          iostat, iomsg)
18     ! the derived-type value/variable
19     dtv-type-spec, INTENT(IN) :: dtv
20     INTEGER, INTENT(IN) :: unit
21     ! the edit descriptor string
22     CHARACTER (LEN=*), INTENT(IN) :: iotype
23     INTEGER, INTENT(IN) :: v_list(:)
24     INTEGER, INTENT(OUT) :: iostat
25     CHARACTER (LEN=*), INTENT(INOUT) :: iomsg
26     END

```

27 If the *dtio-generic-spec* is WRITE (UNFORMATTED), the characteristics shall be the same as those specified by the following interface:

```

29     SUBROUTINE my_write_routine_unformatted  &
30         (dtv,                                  &
31          unit,                                  &
32          iostat, iomsg)
33     ! the derived-type value/variable
34     dtv-type-spec, INTENT(IN) :: dtv
35     INTEGER, INTENT(IN) :: unit
36     INTEGER, INTENT(OUT) :: iostat
37     CHARACTER (LEN=*), INTENT(INOUT) :: iomsg
38     END

```

39 The actual specific procedure names (the *my_..._routine_...* procedure names above) are not significant. In the discussion here and elsewhere, the dummy arguments in these interfaces are referred by the names given above; the names are, however, arbitrary.

42 When a user-defined derived-type input/output procedure is invoked, the processor shall pass a *unit* argument that has a value as follows.

- 44 • If the parent data transfer statement uses a *file-unit-number*, the value of the *unit* argument shall
- 45 be that of the *file-unit-number*.

- 1 • If the parent data transfer statement is a WRITE statement with an asterisk unit or a PRINT
2 statement, the `unit` argument shall have the same value as the `OUTPUT_UNIT` named constant
3 of the `ISO_FORTRAN_ENV` intrinsic module (13.8.3).
- 4 • If the parent data transfer statement is a READ statement with an asterisk unit or a READ
5 statement without an *io-control-spec-list*, the `unit` argument shall have the same value as the
6 `INPUT_UNIT` named constant of the `ISO_FORTRAN_ENV` intrinsic module (13.8.3).
- 7 • Otherwise the parent data transfer statement must access an internal file, in which case the `unit`
8 argument shall have a processor-dependent negative value.

NOTE 9.47

The `unit` argument passed to a user-defined derived-type input/output procedure will be negative when the parent input/output statement specified an internal unit, or specified an external unit that is a `NEWUNIT` value. When an internal unit is used with the `INQUIRE` statement, an error condition will occur, and any variable specified in an `IOSTAT=` specifier will be assigned the value `IOSTAT_INQUIRE_INTERNAL_UNIT` from the `ISO_FORTRAN_ENV` intrinsic module (13.8.3).

- 9 For formatted data transfer, the processor shall pass an `iotype` argument that has the value
- 10 • “LISTDIRECTED” if the parent data transfer statement specified list directed formatting,
 - 11 • “NAMELIST” if the parent data transfer statement specified namelist formatting, or
 - 12 • “DT” concatenated with the *char-literal-constant*, if any, of the edit descriptor, if the parent data
13 transfer statement contained a format specification and the list item’s corresponding edit descriptor
14 was a DT edit descriptor.
- 15 If the parent data transfer statement is a READ statement, the `dtv` dummy argument is argument
16 associated with the effective list item that caused the user-defined derived-type input procedure to be
17 invoked, as if the effective list item were an actual argument in this procedure reference (2.5.6).
- 18 If the parent data transfer statement is a WRITE or PRINT statement, the processor shall provide the
19 value of the effective list item in the `dtv` dummy argument.
- 20 If the *v-list* of the edit descriptor appears in the parent data transfer statement, the processor shall
21 provide the values from it in the `v_list` dummy argument, with the same number of elements in the
22 same order as *v-list*. If there is no *v-list* in the edit descriptor or if the data transfer statement specifies
23 list-directed or namelist formatting, the processor shall provide `v_list` as a zero-sized array.

NOTE 9.48

The user’s procedure may choose to interpret an element of the `v_list` argument as a field width, but this is not required. If it does, it would be appropriate to fill an output field with “*”s if the width is too small.

- 24 The `iostat` argument is used to report whether an error, end-of-record, or end-of-file condition (9.10)
25 occurs. If an error condition occurs, the user-defined derived-type input/output procedure shall assign
26 a positive value to the `iostat` argument. Otherwise, if an end-of-file condition occurs, the user-defined
27 derived-type input procedure shall assign the value of the named constant `IOSTAT_END` (13.8.3.10)
28 to the `iostat` argument. Otherwise, if an end-of-record condition occurs, the user-defined derived-
29 type input procedure shall assign the value of the named constant `IOSTAT_EOR` (13.8.3.11) to `iostat`.
30 Otherwise, the user-defined derived-type input/output procedure shall assign the value zero to the
31 `iostat` argument.

- 1 If the user-defined derived-type input/output procedure returns a nonzero value for the `iostat` argument,
2 the procedure shall also return an explanatory message in the `iomsg` argument. Otherwise, the procedure
3 shall not change the value of the `iomsg` argument.

NOTE 9.49

The values of the `iostat` and `iomsg` arguments set in a user-defined derived-type input/output procedure need not be passed to all of the parent data transfer statements.

- 4 If the `iostat` argument of the user-defined derived-type input/output procedure has a nonzero value
5 when that procedure returns, and the processor therefore terminates execution of the program as de-
6 scribed in 9.10, the processor shall make the value of the `iomsg` argument available in a processor-
7 dependent manner.
- 8 When a parent READ statement is active, an input/output statement shall not read from any external
9 unit other than the one specified by the dummy argument `unit` and shall not perform output to any
10 external unit.
- 11 When a parent WRITE or PRINT statement is active, an input/output statement shall not perform
12 output to any external unit other than the one specified by the dummy argument `unit` and shall not
13 read from any external unit.
- 14 When a parent data transfer statement is active, a data transfer statement that specifies an internal file
15 is permitted.
- 16 OPEN, CLOSE, BACKSPACE, ENDFILE, and REWIND statements shall not be executed while a
17 parent data transfer statement is active.
- 18 A user-defined derived-type input/output procedure may use a FORMAT with a DT edit descriptor for
19 handling a component of the derived type that is itself a derived type. A child data transfer statement
20 that is a list directed or namelist input/output statement may contain a list item of derived type.
- 21 Because a child data transfer statement does not position the file prior to data transfer, the child data
22 transfer statement starts transferring data from where the file was positioned by the parent data transfer
23 statement's most recently processed effective list item or record positioning edit descriptor. This is not
24 necessarily at the beginning of a record.
- 25 A record positioning edit descriptor, such as TL and TR, used on `unit` by a child data transfer statement
26 shall not cause the record position to be positioned before the record position at the time the user-defined
27 derived-type input/output procedure was invoked.

NOTE 9.50

A robust user-defined derived-type input/output procedure may wish to use INQUIRE to determine the settings of BLANK=, PAD=, ROUND=, DECIMAL=, and DELIM= for an external unit. The INQUIRE provides values as specified in 9.9.

- 28 Neither a parent nor child data transfer statement shall be asynchronous.
- 29 A user-defined derived-type input/output procedure, and any procedures invoked therefrom, shall not
30 define, nor cause to become undefined, any storage location referenced by any input/output list item,
31 the corresponding format, or any specifier in any active parent data transfer statement, except through
32 the `dtv` argument.

NOTE 9.51

A child data transfer statement shall not specify the ID=, POS=, or REC= specifiers in an input/output control list.

NOTE 9.52

A simple example of derived type formatted output follows. The derived type variable `chairman` has two components. The type and an associated write formatted procedure are defined in a module so as to be accessible from wherever they might be needed. It would also be possible to check that `iotype` indeed has the value 'DT' and to set `iostat` and `iomsg` accordingly.

```

MODULE p

  TYPE :: person
    CHARACTER (LEN=20) :: name
    INTEGER :: age
  CONTAINS
    PROCEDURE,PRIVATE :: pwf
    GENERIC             :: WRITE(FORMATTED) => pwf
  END TYPE person

CONTAINS

  SUBROUTINE pwf (dtv,unit,iotype,vlist,iostat,iomsg)
! argument definitions
    CLASS(person), INTENT(IN) :: dtv
    INTEGER, INTENT(IN) :: unit
    CHARACTER (LEN=*), INTENT(IN) :: iotype
    INTEGER, INTENT(IN) :: vlist(:)
    INTEGER, INTENT(OUT) :: iostat
    CHARACTER (LEN=*), INTENT(INOUT) :: iomsg
! local variable
    CHARACTER (LEN=9) :: pfmt

!  vlist(1) and (2) are to be used as the field widths of the two
!  components of the derived type variable.  First set up the format to
!  be used for output.
    WRITE(pfmt,'(A,I2,A,I2,A)') '(A', vlist(1), ',I', vlist(2), '))'

!  now the basic output statement
    WRITE(unit, FMT=pfmt, IOSTAT=iostat) dtv%name, dtv%age

  END SUBROUTINE pwf

END MODULE p

PROGRAM
  USE p
  INTEGER id, members
  TYPE (person) :: chairman
  ...
  WRITE(6, FMT="(I2, DT (15,6), I5)" ) id, chairman, members
! this writes a record with four fields, with lengths 2, 15, 6, 5

```

NOTE 9.52 (cont.)

```
! respectively
END PROGRAM
```

NOTE 9.53

In the following example, the variables of the derived type `node` form a linked list, with a single value at each node. The subroutine `pwf` is used to write the values in the list, one per line.

```
MODULE p

  TYPE node
    INTEGER :: value = 0
    TYPE (NODE), POINTER :: next_node => NULL ( )
  CONTAINS
    PROCEDURE,PRIVATE :: pwf
    GENERIC             :: WRITE(FORMATTED) => pwf
  END TYPE node

CONTAINS

  RECURSIVE SUBROUTINE pwf (dtv,unit,iotype,vlist,iostat,iomsg)
! Write the chain of values, each on a separate line in I9 format.
    CLASS(node), INTENT(IN) :: dtv
    INTEGER, INTENT(IN) :: unit
    CHARACTER (LEN=*), INTENT(IN) :: iotype
    INTEGER, INTENT(IN) :: vlist(:)
    INTEGER, INTENT(OUT) :: iostat
    CHARACTER (LEN=*), INTENT(INOUT) :: iomsg

    WRITE(unit,'(i9 /)', IOSTAT = iostat) dtv%value
    IF(iostat/=0) RETURN
    IF(ASSOCIATED(dtv%next_node)) WRITE(unit,'(dt)', IOSTAT=iostat) dtv%next_node
  END SUBROUTINE pwf

END MODULE p
```

1 9.5.4.7.3 Resolving derived-type input/output procedure references

2 A suitable generic interface for user-defined derived-type input/output of an effective item is one that
 3 has a *dtio-generic-spec* that is appropriate to the direction (read or write) and form (formatted or
 4 unformatted) of the data transfer as specified in 9.5.4.7, and has a specific interface whose `dtv` argument
 5 is compatible with the effective item according to the rules for argument association in 12.5.2.5.

6 When an effective item (9.5.3) that is of derived-type is encountered during a data transfer, user-defined
 7 derived-type input/output occurs if both of the following conditions are true.

- 8 (1) The circumstances of the input/output are such that user-defined derived-type input/output
 9 is permitted; that is, either
 - 10 (a) the transfer was initiated by a list-directed, namelist, or unformatted input/output
 11 statement, or
 - 12 (b) a format specification is supplied for the input/output statement, and the edit de-

1 descriptor corresponding to the effective item is a DT edit descriptor.

- 2 (2) A suitable user-defined derived-type input/output procedure is available; that is, either
 3 (a) the declared type of the effective item has a suitable generic type-bound procedure,
 4 or
 5 (b) a suitable generic interface is accessible.

6 If (2a) is true, the procedure referenced is determined as for explicit type-bound procedure references
 7 (12.5); that is, the binding with the appropriate specific interface is located in the declared type of the
 8 effective item, and the corresponding binding in the dynamic type of the effective item is selected.

9 If (2a) is false and (2b) is true, the reference is to the procedure identified by the appropriate specific
 10 interface in the interface block. This reference shall not be to a dummy procedure that is not present,
 11 or to a disassociated procedure pointer.

12 9.5.5 Termination of data transfer statements

13 Termination of an input/output data transfer statement occurs when

- 14 (1) format processing encounters a data edit descriptor and there are no remaining elements in
 15 the *input-item-list* or *output-item-list*,
 16 (2) unformatted or list-directed data transfer exhausts the *input-item-list* or *output-item-list*,
 17 (3) namelist output exhausts the *namelist-group-object-list*,
 18 (4) an error condition occurs,
 19 (5) an end-of-file condition occurs,
 20 (6) a slash (/) is encountered as a value separator (10.10, 10.11) in the record being read during
 21 list-directed or namelist input, or
 22 (7) an end-of-record condition occurs during execution of a nonadvancing input statement
 23 (9.10).

24 9.6 Waiting on pending data transfer

25 9.6.1 Wait operation

26 Execution of an asynchronous data transfer statement in which neither an error, end-of-record, nor end-
 27 of-file condition occurs initiates a pending data transfer operation. There may be multiple pending data
 28 transfer operations for the same or multiple units simultaneously. A pending data transfer operation
 29 remains pending until a corresponding wait operation is performed. A wait operation may be performed
 30 by a WAIT, INQUIRE, FLUSH, CLOSE, data transfer, or file positioning statement.

31 A **wait operation** completes the processing of a pending data transfer operation. Each wait operation
 32 completes only a single data transfer operation, although a single statement may perform multiple wait
 33 operations.

34 If the actual data transfer is not yet complete, the wait operation first waits for its completion. If the
 35 data transfer operation is an input operation that completed without error, the storage units of the
 36 input/output storage sequence then become defined with the values as described in 9.5.2.15 and 9.5.4.4.

37 If any error, end-of-file, or end-of-record conditions occur, the applicable actions specified by the IO-
 38 STAT=, IOMSG=, ERR=, END=, and EOR= specifiers of the statement that performs the wait oper-
 39 ation are taken.

40 If an error or end-of-file condition occurs during a wait operation for a unit, the processor performs a
 41 wait operation for all pending data transfer operations for that unit.

NOTE 9.54

Error, end-of-file, and end-of-record conditions may be raised either during the data transfer statement that initiates asynchronous input/output, a subsequent asynchronous data transfer statement for the same unit, or during the wait operation. If such conditions are raised during a data transfer statement, they trigger actions according to the IOSTAT=, ERR=, END=, and EOR= specifiers of that statement; if they are raised during the wait operation, the actions are in accordance with the specifiers of the statement that performs the wait operation.

1 After completion of the wait operation, the data transfer operation and its input/output storage sequence
2 are no longer considered to be pending.

3 **9.6.2 WAIT statement**

4 A WAIT statement performs a wait operation for specified pending asynchronous data transfer opera-
5 tions.

NOTE 9.55

The CLOSE, INQUIRE, and file positioning statements may also perform wait operations.

6	R921	<i>wait-stmt</i>	is	WAIT (<i>wait-spec-list</i>)
7	R922	<i>wait-spec</i>	is	[UNIT =] <i>file-unit-number</i>
8			or	END = <i>label</i>
9			or	EOR = <i>label</i>
10			or	ERR = <i>label</i>
11			or	ID = <i>scalar-int-expr</i>
12			or	IOMSG = <i>iomsg-variable</i>
13			or	IOSTAT = <i>scalar-int-variable</i>

14 C939 No specifier shall appear more than once in a given *wait-spec-list*.

15 C940 A *file-unit-number* shall be specified in a *wait-spec-list*; if the optional characters UNIT= are
16 omitted, the *file-unit-number* shall be the first item in the *wait-spec-list*.

17 C941 (R922) The *label* in the ERR=, EOR=, or END= specifier shall be the statement label of a
18 branch target statement that appears in the same scoping unit as the WAIT statement.

19 The IOSTAT=, ERR=, EOR=, END=, and IOMSG= specifiers are described in 9.10.

20 The value of the expression specified in the ID= specifier shall be the identifier of a pending data transfer
21 operation for the specified unit. If the ID= specifier appears, a wait operation for the specified data
22 transfer operation is performed. If the ID= specifier is omitted, wait operations for all pending data
23 transfers for the specified unit are performed.

24 Execution of a WAIT statement specifying a unit that does not exist, has no file connected to it, or is
25 not open for asynchronous input/output is permitted, provided that the WAIT statement has no ID=
26 specifier; such a WAIT statement does not cause an error or end-of-file condition to occur.

NOTE 9.56

An EOR= specifier has no effect if the pending data transfer operation is not a nonadvancing read.
And END= specifier has no effect if the pending data transfer operation is not a READ.

27 **9.7 File positioning statements**

1 9.7.1 Syntax

2	R923	<i>backspace-stmt</i>	is BACKSPACE <i>file-unit-number</i>
3			or BACKSPACE (<i>position-spec-list</i>)
4	R924	<i>endfile-stmt</i>	is ENDFILE <i>file-unit-number</i>
5			or ENDFILE (<i>position-spec-list</i>)
6	R925	<i>rewind-stmt</i>	is REWIND <i>file-unit-number</i>
7			or REWIND (<i>position-spec-list</i>)

8 A unit that is connected for direct access shall not be referred to by a BACKSPACE, ENDFILE, or
 9 REWIND statement. A unit that is connected for unformatted stream access shall not be referred to
 10 by a BACKSPACE statement. A unit that is connected with an ACTION= specifier having the value
 11 READ shall not be referred to by an ENDFILE statement.

12	R926	<i>position-spec</i>	is [UNIT =] <i>file-unit-number</i>
13			or IOMSG = <i>iomsg-variable</i>
14			or IOSTAT = <i>scalar-int-variable</i>
15			or ERR = <i>label</i>

16 C942 No specifier shall appear more than once in a given *position-spec-list*.

17 C943 A *file-unit-number* shall be specified in a *position-spec-list*; if the optional characters UNIT=
 18 are omitted, the *file-unit-number* shall be the first item in the *position-spec-list*.

19 C944 (R926) The *label* in the ERR= specifier shall be the statement label of a branch target statement
 20 that appears in the same scoping unit as the file positioning statement.

21 The IOSTAT=, ERR=, and IOMSG= specifiers are described in 9.10.

22 Execution of a file positioning statement performs a wait operation for all pending asynchronous data
 23 transfer operations for the specified unit.

24 9.7.2 BACKSPACE statement

25 Execution of a BACKSPACE statement causes the file connected to the specified unit to be positioned
 26 before the current record if there is a current record, or before the preceding record if there is no current
 27 record. If the file is at its initial point, the position of the file is not changed.

NOTE 9.57

If the preceding record is an endfile record, the file is positioned before the endfile record.

28 If a BACKSPACE statement causes the implicit writing of an endfile record, the file is positioned before
 29 the record that precedes the endfile record.

30 Backspacing a file that is connected but does not exist is prohibited.

31 Backspacing over records written using list-directed or namelist formatting is prohibited.

32 A BACKSPACE statement shall not reference a unit whose connect team has more than one image.

NOTE 9.58

An example of a BACKSPACE statement is:

```
BACKSPACE (10, IOSTAT = N)
```

1 9.7.3 ENDFILE statement

2 Execution of an ENDFILE statement for a file connected for sequential access writes an endfile record
 3 as the next record of the file. The file is then positioned after the endfile record, which becomes the last
 4 record of the file. If the file also may be connected for direct access, only those records before the endfile
 5 record are considered to have been written. Thus, only those records may be read during subsequent
 6 direct access connections to the file.

7 After execution of an ENDFILE statement for a file connected for sequential access, a BACKSPACE
 8 or REWIND statement shall be used to reposition the file prior to execution of any data transfer
 9 input/output statement or ENDFILE statement.

10 Execution of an ENDFILE statement for a file connected for stream access causes the terminal point of
 11 the file to become equal to the current file position. Only file storage units before the current position are
 12 considered to have been written; thus only those file storage units may be subsequently read. Subsequent
 13 stream output statements may be used to write further data to the file.

14 Execution of an ENDFILE statement for a file that is connected but does not exist creates the file; if
 15 the file is connected for sequential access, it is created prior to writing the endfile record.

16 An ENDFILE statement shall not reference a unit whose connect team has more than one image.

NOTE 9.59

An example of an ENDFILE statement is:

```
ENDFILE K
```

17 9.7.4 REWIND statement

18 Execution of a REWIND statement causes the specified file to be positioned at its initial point.

NOTE 9.60

If the file is already positioned at its initial point, execution of this statement has no effect on the
 position of the file.

19 Execution of a REWIND statement for a file that is connected but does not exist is permitted and has
 20 no effect on any file.

21 A REWIND statement shall not reference a unit whose connect team has more than one image.

NOTE 9.61

An example of a REWIND statement is:

```
REWIND 10
```

22 9.8 FLUSH statement

23 The form of the FLUSH statement is:

24	R927	<i>flush-stmt</i>	is	FLUSH <i>file-unit-number</i>
25			or	FLUSH (<i>flush-spec-list</i>)
26	R928	<i>flush-spec</i>	is	[UNIT =] <i>file-unit-number</i>
27			or	IOSTAT = <i>scalar-int-variable</i>

1 or IOMSG = *iomsg-variable*
 2 or ERR = *label*

3 C945 No specifier shall appear more than once in a given *flush-spec-list*.

4 C946 A *file-unit-number* shall be specified in a *flush-spec-list*; if the optional characters UNIT= are
 5 omitted from the unit specifier, the *file-unit-number* shall be the first item in the *flush-spec-list*.

6 C947 (R928) The *label* in the ERR= specifier shall be the statement label of a branch target statement
 7 that appears in the same scoping unit as the FLUSH statement.

8 The IOSTAT=, IOMSG= and ERR= specifiers are described in 9.10. The IOSTAT= variable shall
 9 be set to a processor-dependent positive value if an error occurs, to zero if the processor-dependent
 10 flush operation was successful, or to a processor-dependent negative value if the flush operation is not
 11 supported for the unit specified.

12 Execution of a FLUSH statement causes data written to an external unit to be made available to other
 13 images of the unit's connect team which execute a FLUSH statement in a subsequent segment for that
 14 unit. It also causes data written to an external file to be available to other processes, or causes data
 15 placed in an external file by means other than Fortran to be available to a READ statement. The action
 16 is processor dependent.

17 Execution of a FLUSH statement for a file that is connected but does not exist is permitted and has no
 18 effect on any file. A FLUSH statement has no effect on file position.

19 Execution of a FLUSH statement performs a wait operation for all pending asynchronous data transfer
 20 operations for the specified unit.

NOTE 9.62

Because this standard does not specify the mechanism of file storage, the exact meaning of the flush operation is not precisely defined. The intention is that the flush operation should make all data written to a file available to other processes or devices, or make data recently added to a file by other processes or devices available to the program via a subsequent read operation. This is commonly called "flushing I/O buffers".

NOTE 9.63

An example of a FLUSH statement is:

```
FLUSH( 10, IOSTAT=N)
```

21 9.9 File inquiry statement

22 9.9.1 Forms of the INQUIRE statement

23 The INQUIRE statement may be used to inquire about properties of a particular named file or of the
 24 connection to a particular unit. There are three forms of the INQUIRE statement: **inquire by file**,
 25 which uses the FILE= specifier, **inquire by unit**, which uses the UNIT= specifier, and **inquire by**
 26 **output list**, which uses only the IOLENGTH= specifier. All specifier value assignments are performed
 27 according to the rules for assignment statements.

28 For inquiry by unit, the unit specified need not exist or be connected to a file. If it is connected to a
 29 file, the inquiry is being made about the connection and about the file connected.

30 An INQUIRE statement may be executed before, while, or after a file is connected to a unit. All values

1 assigned by an INQUIRE statement are those that are current at the time the statement is executed.

2 R929 *inquire-stmt* **is** INQUIRE (*inquire-spec-list*)
 3 **or** INQUIRE (IOLENGTH = *scalar-int-variable*) ■
 4 ■ *output-item-list*

NOTE 9.64

Examples of INQUIRE statements are:

```
INQUIRE (IOLENGTH = IOL) A (1:N)
INQUIRE (UNIT = JOAN, OPENED = LOG_01, NAMED = LOG_02, &
FORM = CHAR_VAR, IOSTAT = IOS)
```

5 9.9.2 Inquiry specifiers

6 Unless constrained, the following inquiry specifiers may be used in either of the inquire by file or inquire
 7 by unit forms of the INQUIRE statement:

8 R930 *inquire-spec* **is** [UNIT =] *file-unit-number*
 9 **or** FILE = *file-name-expr*
 10 **or** ACCESS = *scalar-default-char-variable*
 11 **or** ACTION = *scalar-default-char-variable*
 12 **or** ASYNCHRONOUS = *scalar-default-char-variable*
 13 **or** BLANK = *scalar-default-char-variable*
 14 **or** DECIMAL = *scalar-default-char-variable*
 15 **or** DELIM = *scalar-default-char-variable*
 16 **or** DIRECT = *scalar-default-char-variable*
 17 **or** ENCODING = *scalar-default-char-variable*
 18 **or** ERR = *label*
 19 **or** EXIST = *scalar-default-logical-variable*
 20 **or** FORM = *scalar-default-char-variable*
 21 **or** FORMATTED = *scalar-default-char-variable*
 22 **or** ID = *scalar-int-expr*
 23 **or** IOMSG = *iomsg-variable*
 24 **or** IOSTAT = *scalar-int-variable*
 25 **or** NAME = *scalar-default-char-variable*
 26 **or** NAMED = *scalar-default-logical-variable*
 27 **or** NEXTREC = *scalar-int-variable*
 28 **or** NUMBER = *scalar-int-variable*
 29 **or** OPENED = *scalar-default-logical-variable*
 30 **or** PAD = *scalar-default-char-variable*
 31 **or** PENDING = *scalar-default-logical-variable*
 32 **or** POS = *scalar-int-variable*
 33 **or** POSITION = *scalar-default-char-variable*
 34 **or** READ = *scalar-default-char-variable*
 35 **or** READWRITE = *scalar-default-char-variable*
 36 **or** RECL = *scalar-int-variable*
 37 **or** ROUND = *scalar-default-char-variable*
 38 **or** SEQUENTIAL = *scalar-default-char-variable*
 39 **or** SIGN = *scalar-default-char-variable*
 40 **or** SIZE = *scalar-int-variable*
 41 **or** STREAM = *scalar-default-char-variable*
 42 **or** TEAM = *image-team*
 43 **or** UNFORMATTED = *scalar-default-char-variable*

- 1 **or** WRITE = *scalar-default-char-variable*
- 2 C948 No specifier shall appear more than once in a given *inquire-spec-list*.
- 3 C949 An *inquire-spec-list* shall contain one FILE= specifier or one UNIT= specifier, but not both.
- 4 C950 In the inquire by unit form of the INQUIRE statement, if the optional characters UNIT= are
5 omitted, the *file-unit-number* shall be the first item in the *inquire-spec-list*.
- 6 C951 If an ID= specifier appears in an *inquire-spec-list*, a PENDING= specifier shall also appear.
- 7 C952 (R928) The *label* in the ERR= specifier shall be the statement label of a branch target statement
8 that appears in the same scoping unit as the INQUIRE statement.
- 9 If *file-unit-number* identifies an internal unit (9.5.4.7.2), an error condition occurs.
- 10 When a returned value of a specifier other than the NAME= specifier is of type character, the value
11 returned is in upper case.
- 12 If an error condition occurs during execution of an INQUIRE statement, all of the inquiry specifier
13 variables become undefined, except for variables in the IOSTAT= and IOMSG= specifiers (if any).
- 14 The IOSTAT=, ERR=, and IOMSG= specifiers are described in 9.10.

15 **9.9.2.1 FILE= specifier in the INQUIRE statement**

16 The value of the *file-name-expr* in the FILE= specifier specifies the name of the file being inquired about.
17 The named file need not exist or be connected to a unit. The value of the *file-name-expr* shall be of a
18 form acceptable to the processor as a file name. Any trailing blanks are ignored. The interpretation of
19 case is processor dependent.

20 **9.9.2.2 ACCESS= specifier in the INQUIRE statement**

21 The *scalar-default-char-variable* in the ACCESS= specifier is assigned the value SEQUENTIAL if the
22 connection is for sequential access, DIRECT if the connection is for direct access, or STREAM if the
23 connection is for stream access. If there is no connection, it is assigned the value UNDEFINED.

24 **9.9.2.3 ACTION= specifier in the INQUIRE statement**

25 The *scalar-default-char-variable* in the ACTION= specifier is assigned the value READ if the connection
26 is for input only, WRITE if the connection is for output only, and READWRITE if the connection is for
27 both input and output. If there is no connection, the *scalar-default-char-variable* is assigned the value
28 UNDEFINED.

29 **9.9.2.4 ASYNCHRONOUS= specifier in the INQUIRE statement**

30 The *scalar-default-char-variable* in the ASYNCHRONOUS= specifier is assigned the value YES if the
31 connection allows asynchronous input/output; it is assigned the value NO if the connection does not
32 allow asynchronous input/output. If there is no connection, the *scalar-default-char-variable* is assigned
33 the value UNDEFINED.

34 **9.9.2.5 BLANK= specifier in the INQUIRE statement**

35 The *scalar-default-char-variable* in the BLANK= specifier is assigned the value ZERO or NULL, corre-
36 sponding to the blank interpretation mode in effect for a connection for formatted input/output. If there
37 is no connection, or if the connection is not for formatted input/output, the *scalar-default-char-variable*
38 is assigned the value UNDEFINED.

1 **9.9.2.6 DECIMAL= specifier in the INQUIRE statement**

2 The *scalar-default-char-variable* in the DECIMAL= specifier is assigned the value COMMA or POINT,
3 corresponding to the decimal edit mode in effect for a connection for formatted input/output. If there
4 is no connection, or if the connection is not for formatted input/output, the *scalar-default-char-variable*
5 is assigned the value UNDEFINED.

6 **9.9.2.7 DELIM= specifier in the INQUIRE statement**

7 The *scalar-default-char-variable* in the DELIM= specifier is assigned the value APOSTROPHE, QUOTE,
8 or NONE, corresponding to the delimiter mode in effect for a connection for formatted input/output.
9 If there is no connection or if the connection is not for formatted input/output, the *scalar-default-char-*
10 *variable* is assigned the value UNDEFINED.

11 **9.9.2.8 DIRECT= specifier in the INQUIRE statement**

12 The *scalar-default-char-variable* in the DIRECT= specifier is assigned the value YES if DIRECT is
13 included in the set of allowed access methods for the file, NO if DIRECT is not included in the set of
14 allowed access methods for the file, and UNKNOWN if the processor is unable to determine whether
15 DIRECT is included in the set of allowed access methods for the file.

16 **9.9.2.9 ENCODING= specifier in the INQUIRE statement**

17 The *scalar-default-char-variable* in the ENCODING= specifier is assigned the value UTF-8 if the con-
18 nection is for formatted input/output with an encoding form of UTF-8, and is assigned the value UN-
19 DEFINED if the connection is for unformatted input/output. If there is no connection, it is assigned
20 the value UTF-8 if the processor is able to determine that the encoding form of the file is UTF-8; if
21 the processor is unable to determine the encoding form of the file, the variable is assigned the value
22 UNKNOWN.

NOTE 9.65

The value assigned may be something other than UTF-8, UNDEFINED, or UNKNOWN if the processor supports other specific encoding forms (e.g. UTF-16BE).

23 **9.9.2.10 EXIST= specifier in the INQUIRE statement**

24 Execution of an INQUIRE by file statement causes the *scalar-default-logical-variable* in the EXIST=
25 specifier to be assigned the value true if there exists a file with the specified name; otherwise, false is
26 assigned. Execution of an INQUIRE by unit statement causes true to be assigned if the specified unit
27 exists; otherwise, false is assigned.

28 **9.9.2.11 FORM= specifier in the INQUIRE statement**

29 The *scalar-default-char-variable* in the FORM= specifier is assigned the value FORMATTED if the
30 connection is for formatted input/output, and is assigned the value UNFORMATTED if the connection
31 is for unformatted input/output. If there is no connection, it is assigned the value UNDEFINED.

32 **9.9.2.12 FORMATTED= specifier in the INQUIRE statement**

33 The *scalar-default-char-variable* in the FORMATTED= specifier is assigned the value YES if FOR-
34 MATTED is included in the set of allowed forms for the file, NO if FORMATTED is not included in
35 the set of allowed forms for the file, and UNKNOWN if the processor is unable to determine whether
36 FORMATTED is included in the set of allowed forms for the file.

1 **9.9.2.13 ID= specifier in the INQUIRE statement**

2 The value of the expression specified in the ID= specifier shall be the identifier of a pending data transfer
3 operation for the specified unit. This specifier interacts with the PENDING= specifier (9.9.2.20).

4 **9.9.2.14 NAME= specifier in the INQUIRE statement**

5 The *scalar-default-char-variable* in the NAME= specifier is assigned the value of the name of the file if
6 the file has a name; otherwise, it becomes undefined.

NOTE 9.66

If this specifier appears in an INQUIRE by file statement, its value is not necessarily the same as the name given in the FILE= specifier. However, the value returned shall be suitable for use as the value of the *file-name-expr* in the FILE= specifier in an OPEN statement.

The processor may return a file name qualified by a user identification, device, directory, or other relevant information.

7 The case of the characters assigned to *scalar-default-char-variable* is processor dependent.

8 **9.9.2.15 NAMED= specifier in the INQUIRE statement**

9 The *scalar-default-logical-variable* in the NAMED= specifier is assigned the value true if the file has a
10 name; otherwise, it is assigned the value false.

11 **9.9.2.16 NEXTREC= specifier in the INQUIRE statement**

12 The *scalar-int-variable* in the NEXTREC= specifier is assigned the value $n + 1$, where n is the record
13 number of the last record read from or written to the connection for direct access. If there is a connection
14 but no records have been read or written since the connection, the *scalar-int-variable* is assigned the
15 value 1. If there is no connection, the connection is not for direct access, or the position is indeterminate
16 because of a previous error condition, the *scalar-int-variable* becomes undefined. If there are pending
17 data transfer operations for the specified unit, the value assigned is computed as if all the pending data
18 transfers had already completed.

J3 internal note

Unresolved Technical Issue 044

NEXTREC= does not have a well-defined meaning in the context of a unit opened with the TEAM= specifier.

19 **9.9.2.17 NUMBER= specifier in the INQUIRE statement**

20 The *scalar-int-variable* in the NUMBER= specifier is assigned the value of the external unit number of
21 the unit that is connected to the file. If there is no unit connected to the file, the value -1 is assigned.

22 **9.9.2.18 OPENED= specifier in the INQUIRE statement**

23 Execution of an INQUIRE by file statement causes the *scalar-default-logical-variable* in the OPENED=
24 specifier to be assigned the value true if the file specified is connected to a unit; otherwise, false is
25 assigned. Execution of an INQUIRE by unit statement causes the *scalar-default-logical-variable* to be
26 assigned the value true if the specified unit is connected to a file; otherwise, false is assigned.

27 **9.9.2.19 PAD= specifier in the INQUIRE statement**

1 The *scalar-default-char-variable* in the PAD= specifier is assigned the value YES or NO, corresponding
2 to the pad mode in effect for a connection for formatted input/output. If there is no connection or if
3 the connection is not for formatted input/output, the *scalar-default-char-variable* is assigned the value
4 UNDEFINED.

5 **9.9.2.20 PENDING= specifier in the INQUIRE statement**

6 The PENDING= specifier is used to determine whether previously pending asynchronous data transfers
7 are complete. A data transfer operation is previously pending if it is pending at the beginning of
8 execution of the INQUIRE statement.

9 If an ID= specifier appears and the specified data transfer operation is complete, then the variable
10 specified in the PENDING= specifier is assigned the value false and the INQUIRE statement performs
11 the wait operation for the specified data transfer.

12 If the ID= specifier is omitted and all previously pending data transfer operations for the specified unit
13 are complete, then the variable specified in the PENDING= specifier is assigned the value false and the
14 INQUIRE statement performs wait operations for all previously pending data transfers for the specified
15 unit.

16 In all other cases, the variable specified in the PENDING= specifier is assigned the value true and no
17 wait operations are performed; in this case the previously pending data transfers remain pending after
18 the execution of the INQUIRE statement.

NOTE 9.67

The processor has considerable flexibility in defining when it considers a transfer to be complete.
Any of the following approaches could be used:

- (1) The INQUIRE statement could consider an asynchronous data transfer to be incomplete until after the corresponding wait operation. In this case PENDING= would always return true unless there were no previously pending data transfers for the unit.
- (2) The INQUIRE statement could wait for all specified data transfers to complete and then always return false for PENDING=.
- (3) The INQUIRE statement could actually test the state of the specified data transfer operations.

19 **9.9.2.21 POS= specifier in the INQUIRE statement**

20 The *scalar-int-variable* in the POS= specifier is assigned the number of the file storage unit immediately
21 following the current position of a file connected for stream access. If the file is positioned at its terminal
22 position, the variable is assigned a value one greater than the number of the highest-numbered file storage
23 unit in the file. If the file is not connected for stream access or if the position of the file is indeterminate
24 because of previous error conditions, the variable becomes undefined.

25 **9.9.2.22 POSITION= specifier in the INQUIRE statement**

26 The *scalar-default-char-variable* in the POSITION= specifier is assigned the value REWIND if the
27 connection was opened for positioning at its initial point, APPEND if the connection was opened for
28 positioning before its endfile record or at its terminal point, and ASIS if the connection was opened
29 without changing its position. If there is no connection or if the file is connected for direct access, the
30 *scalar-default-char-variable* is assigned the value UNDEFINED. If the file has been repositioned since
31 the connection, the *scalar-default-char-variable* is assigned a processor-dependent value, which shall not
32 be REWIND unless the file is positioned at its initial point and shall not be APPEND unless the file is
33 positioned so that its endfile record is the next record or at its terminal point if it has no endfile record.

1 **9.9.2.23 READ= specifier in the INQUIRE statement**

2 The *scalar-default-char-variable* in the READ= specifier is assigned the value YES if READ is included
3 in the set of allowed actions for the file, NO if READ is not included in the set of allowed actions for
4 the file, and UNKNOWN if the processor is unable to determine whether READ is included in the set
5 of allowed actions for the file.

6 **9.9.2.24 READWRITE= specifier in the INQUIRE statement**

7 The *scalar-default-char-variable* in the READWRITE= specifier is assigned the value YES if READ-
8 WRITE is included in the set of allowed actions for the file, NO if READWRITE is not included in
9 the set of allowed actions for the file, and UNKNOWN if the processor is unable to determine whether
10 READWRITE is included in the set of allowed actions for the file.

11 **9.9.2.25 RECL= specifier in the INQUIRE statement**

12 The *scalar-int-variable* in the RECL= specifier is assigned the value of the record length of a connection
13 for direct access, or the value of the maximum record length of a connection for sequential access. If
14 the connection is for formatted input/output, the length is the number of characters for all records that
15 contain only characters of type default character. If the connection is for unformatted input/output,
16 the length is measured in file storage units. If there is no connection, or if the connection is for stream
17 access, the *scalar-int-variable* becomes undefined.

18 **9.9.2.26 ROUND= specifier in the INQUIRE statement**

19 The *scalar-default-char-variable* in the ROUND= specifier is assigned the value UP, DOWN, ZERO,
20 NEAREST, COMPATIBLE, or PROCESSOR_DEFINED, corresponding to the I/O rounding mode in
21 effect for a connection for formatted input/output. If there is no connection or if the connection is not
22 for formatted input/output, the *scalar-default-char-variable* is assigned the value UNDEFINED. The
23 processor shall return the value PROCESSOR_DEFINED only if the I/O rounding mode currently in
24 effect behaves differently than the UP, DOWN, ZERO, NEAREST, and COMPATIBLE modes.

25 **9.9.2.27 SEQUENTIAL= specifier in the INQUIRE statement**

26 The *scalar-default-char-variable* in the SEQUENTIAL= specifier is assigned the value YES if SEQUEN-
27 TIAL is included in the set of allowed access methods for the file, NO if SEQUENTIAL is not included
28 in the set of allowed access methods for the file, and UNKNOWN if the processor is unable to determine
29 whether SEQUENTIAL is included in the set of allowed access methods for the file.

30 **9.9.2.28 SIGN= specifier in the INQUIRE statement**

31 The *scalar-default-char-variable* in the SIGN= specifier is assigned the value PLUS, SUPPRESS, or
32 PROCESSOR_DEFINED, corresponding to the sign mode in effect for a connection for formatted in-
33 put/output. If there is no connection, or if the connection is not for formatted input/output, the
34 *scalar-default-char-variable* is assigned the value UNDEFINED.

35 **9.9.2.29 SIZE= specifier in the INQUIRE statement**

36 The *scalar-int-variable* in the SIZE= specifier is assigned the size of the file in file storage units. If the
37 file size cannot be determined, the variable is assigned the value -1.

38 For a file that may be connected for stream access, the file size is the number of the highest-numbered
39 file storage unit in the file.

40 For a file that may be connected for sequential or direct access, the file size may be different from the
41 number of storage units implied by the data in the records; the exact relationship is processor-dependent.

1 9.9.2.30 STREAM= specifier in the INQUIRE statement

2 The *scalar-default-char-variable* in the STREAM= specifier is assigned the value YES if STREAM is
3 included in the set of allowed access methods for the file, NO if STREAM is not included in the set of
4 allowed access methods for the file, and UNKNOWN if the processor is unable to determine whether
5 STREAM is included in the set of allowed access methods for the file.

6 9.9.2.31 TEAM= specifier in the INQUIRE statement

7 The *image-team* in the TEAM= specifier is assigned the value of the connect team if the file or unit is
8 connected; otherwise it is assigned the value NULL_IMAGE_TEAM (13.8.3.13).

NOTE 9.68

The indices of the images in a team may be obtained by using the intrinsic function TEAM-
IMAGES (13.7.172).

9 9.9.2.32 UNFORMATTED= specifier in the INQUIRE statement

10 The *scalar-default-char-variable* in the UNFORMATTED= specifier is assigned the value YES if UN-
11 FORMATTED is included in the set of allowed forms for the file, NO if UNFORMATTED is not included
12 in the set of allowed forms for the file, and UNKNOWN if the processor is unable to determine whether
13 UNFORMATTED is included in the set of allowed forms for the file.

14 9.9.2.33 WRITE= specifier in the INQUIRE statement

15 The *scalar-default-char-variable* in the WRITE= specifier is assigned the value YES if WRITE is included
16 in the set of allowed actions for the file, NO if WRITE is not included in the set of allowed actions for
17 the file, and UNKNOWN if the processor is unable to determine whether WRITE is included in the set
18 of allowed actions for the file.

19 9.9.3 Inquire by output list

20 The *scalar-int-variable* in the IOLENGTH= specifier is assigned the processor-dependent number of file
21 storage units that would be required to store the data of the output list in an unformatted file. The
22 value shall be suitable as a RECL= specifier in an OPEN statement that connects a file for unformatted
23 direct access when there are input/output statements with the same input/output list.

24 The output list in an INQUIRE statement shall not contain any derived-type list items that require
25 a user-defined derived-type input/output procedure as described in subclause 9.5.3. If a derived-type
26 list item appears in the output list, the value returned for the IOLENGTH= specifier assumes that no
27 user-defined derived-type input/output procedure will be invoked.

28 9.10 Error, end-of-record, and end-of-file conditions

29 The set of input/output error conditions is processor dependent.

30 An **end-of-record condition** occurs when a nonadvancing input statement attempts to transfer data
31 from a position beyond the end of the current record, unless the file is a stream file and the current
32 record is at the end of the file (an end-of-file condition occurs instead).

33 An **end-of-file condition** occurs when

- 34 (1) an endfile record is encountered during the reading of a file connected for sequential access,
- 35 (2) an attempt is made to read a record beyond the end of an internal file, or

1 (3) an attempt is made to read beyond the end of a stream file.

2 An end-of-file condition may occur at the beginning of execution of an input statement. An end-of-file
3 condition also may occur during execution of a formatted input statement when more than one record
4 is required by the interaction of the input list and the format. An end-of-file condition also may occur
5 during execution of a stream input statement.

6 **9.10.1 Error conditions and the ERR= specifier**

7 If an error condition occurs during execution of an input/output statement, the position of the file
8 becomes indeterminate.

9 If an error condition occurs during execution of an input/output statement that contains neither an
10 ERR= nor IOSTAT= specifier, execution of the program is terminated. If an error condition occurs
11 during execution of an input/output statement that contains either an ERR= specifier or an IOSTAT=
12 specifier then

- 13 (1) processing of the input/output list, if any, terminates,
- 14 (2) if the statement is a data transfer statement or the error occurs during a wait operation, all
15 *do-variables* in the statement that initiated the transfer become undefined,
- 16 (3) if an IOSTAT= specifier appears, the *scalar-int-variable* in the IOSTAT= specifier becomes
17 defined as specified in 9.10.4,
- 18 (4) if an IOMSG= specifier appears, the *iomsg-variable* becomes defined as specified in 9.10.5,
- 19 (5) if the statement is a READ statement and it contains a SIZE= specifier, the *scalar-int-*
20 *variable* in the SIZE= specifier becomes defined as specified in 9.5.2.15,
- 21 (6) if the statement is a READ statement or the error condition occurs in a wait operation for
22 a transfer initiated by a READ statement, all input items or namelist group objects in the
23 statement that initiated the transfer become undefined, and
- 24 (7) if an ERR= specifier appears, a branch to the statement labeled by the *label* in the ERR=
25 specifier occurs.

26 **9.10.2 End-of-file condition and the END= specifier**

27 If an end-of-file condition occurs during execution of an input/output statement that contains neither
28 an END= specifier nor an IOSTAT= specifier, execution of the program is terminated. If an end-of-file
29 condition occurs during execution of an input/output statement that contains either an END= specifier
30 or an IOSTAT= specifier, and an error condition does not occur then

- 31 (1) processing of the input list, if any, terminates,
- 32 (2) if the statement is a data transfer statement or the error occurs during a wait operation, all
33 *do-variables* in the statement that initiated the transfer become undefined,
- 34 (3) if the statement is a READ statement or the end-of-file condition occurs in a wait operation
35 for a transfer initiated by a READ statement, all input list items or namelist group objects
36 in the statement that initiated the transfer become undefined,
- 37 (4) if the file specified in the input statement is an external record file, it is positioned after the
38 endfile record,
- 39 (5) if an IOSTAT= specifier appears, the *scalar-int-variable* in the IOSTAT= specifier becomes
40 defined as specified in 9.10.4,
- 41 (6) if an IOMSG= specifier appears, the *iomsg-variable* becomes defined as specified in 9.10.5,
42 and
- 43 (7) if an END= specifier appears, a branch to the statement labeled by the *label* in the END=
44 specifier occurs.

1 9.10.3 End-of-record condition and the EOR= specifier

2 If an end-of-record condition occurs during execution of an input/output statement that contains neither
 3 an EOR= specifier nor an IOSTAT= specifier, execution of the program is terminated. If an end-of-
 4 record condition occurs during execution of an input/output statement that contains either an EOR=
 5 specifier or an IOSTAT= specifier, and an error condition does not occur then

- 6 (1) If the pad mode has the value
 - 7 (a) YES, the record is padded with blanks to satisfy the effective list item (9.5.4.4.2)
 8 and corresponding data edit descriptors that require more characters than the record
 9 contains,
 - 10 (b) NO, the input list item becomes undefined,
- 11 (2) processing of the input list, if any, terminates,
- 12 (3) if the statement is a data transfer statement or the end-of-record condition occurs dur-
 13 ing a wait operation, all *do-variables* in the statement that initiated the transfer become
 14 undefined,
- 15 (4) the file specified in the input statement is positioned after the current record,
- 16 (5) if an IOSTAT= specifier appears, the *scalar-int-variable* in the IOSTAT= specifier becomes
 17 defined as specified in 9.10.4,
- 18 (6) if an IOMSG= specifier appears, the *iomsg-variable* becomes defined as specified in 9.10.5,
- 19 (7) if a SIZE= specifier appears, the *scalar-int-variable* in the SIZE= specifier becomes defined
 20 as specified in (9.5.2.15), and
- 21 (8) if an EOR= specifier appears, a branch to the statement labeled by the *label* in the EOR=
 22 specifier occurs.

23 9.10.4 IOSTAT= specifier

24 Execution of an input/output statement containing the IOSTAT= specifier causes the *scalar-int-variable*
 25 in the IOSTAT= specifier to become defined with

- 26 (1) a zero value if neither an error condition, an end-of-file condition, nor an end-of-record
 27 condition occurs,
- 28 (2) the processor-dependent positive integer value of the constant IOSTAT_INQUIRE_INTER-
 29 NAL_UNIT from the ISO_FORTRAN_ENV intrinsic module (13.8.3) if a unit number in an
 30 INQUIRE statement identifies an internal file,
- 31 (3) a processor-dependent positive integer value different from IOSTAT_INQUIRE_INTER-
 32 NAL_UNIT if any other error condition occurs,
- 33 (4) the processor-dependent negative integer value of the constant IOSTAT_END (13.8.3.10) if
 34 an end-of-file condition occurs and no error condition occurs, or
- 35 (5) the processor-dependent negative integer value of the constant IOSTAT_EOR (13.8.3.11) if
 36 an end-of-record condition occurs and no error condition or end-of-file condition occurs.

NOTE 9.69

An end-of-file condition may occur only for sequential or stream input and an end-of-record condition may occur only for nonadvancing input.

Consider the example:

```

READ (FMT = "(E8.3)", UNIT = 3, IOSTAT = IOSS) X
IF (IOSS < 0) THEN
  ! Perform end-of-file processing on the file connected to unit 3.
  CALL END_PROCESSING

```


NOTE 9.69 (cont.)

```

ELSE IF (IOSS > 0) THEN
  ! Perform error processing
  CALL ERROR_PROCESSING
END IF

```

1 9.10.5 IOMSG= specifier

2 If an error, end-of-file, or end-of-record condition occurs during execution of an input/output statement,
3 the processor shall assign an explanatory message to *iomsg-variable*. If no such condition occurs, the
4 processor shall not change the value of *iomsg-variable*.

5 9.11 Restrictions on input/output statements

6 If a unit, or a file connected to a unit, does not have all of the properties required for the execution of
7 certain input/output statements, those statements shall not refer to the unit.

8 An input/output statement that is executed while another input/output statement is being executed is
9 called a **recursive input/output statement**.

10 A recursive input/output statement shall not identify an external unit that is identified by another
11 input/output statement being executed except that a child data transfer statement may identify its
12 parent data transfer statement external unit.

13 An input/output statement shall not cause the value of any established format specification to be
14 modified.

15 A recursive input/output statement shall not modify the value of any internal unit except that a recursive
16 WRITE statement may modify the internal unit identified by that recursive WRITE statement.

17 The value of a specifier in an input/output statement shall not depend on any *input-item*, *io-implied-*
18 *do do-variable*, or on the definition or evaluation of any other specifier in the *io-control-spec-list* or
19 *inquire-spec-list* in that statement.

20 The value of any subscript or substring bound of a variable that appears in a specifier in an input/output
21 statement shall not depend on any *input-item*, *io-implied-do do-variable*, or on the definition or evalua-
22 tion of any other specifier in the *io-control-spec-list* or *inquire-spec-list* in that statement.

23 In a data transfer statement, the variable specified in an IOSTAT=, IOMSG=, or SIZE= specifier, if
24 any, shall not be associated with any entity in the data transfer input/output list (9.5.3) or *namelist-*
25 *group-object-list*, nor with a *do-variable* of an *io-implied-do* in the data transfer input/output list.

26 In a data transfer statement, if a variable specified in an IOSTAT=, IOMSG=, or SIZE= specifier is an
27 array element reference, its subscript values shall not be affected by the data transfer, the *io-implied-do*
28 processing, or the definition or evaluation of any other specifier in the *io-control-spec-list*.

29 A variable that may become defined or undefined as a result of its use in a specifier in an INQUIRE
30 statement, or any associated entity, shall not appear in another specifier in the same INQUIRE statement.

31 A STOP statement shall not be executed during execution of an input/output statement.

NOTE 9.70

Restrictions on the evaluation of expressions (7.1.8) prohibit certain side effects.

1 10 Input/output editing

2 10.1 Format specifications

3 A format used in conjunction with an input/output statement provides information that directs the
4 editing between the internal representation of data and the characters of a sequence of formatted records.

5 A R914 (9.5.2.2) in an input/output statement may refer to a FORMAT statement or to a character
6 expression that contains a format specification. A format specification provides explicit editing infor-
7 mation. The R914 alternatively may be an asterisk (*), which indicates list-directed formatting (10.10).
8 Namelist formatting (10.11) may be indicated by specifying a *namelist-group-name* instead of a *format*.

9 10.2 Explicit format specification methods

10 10.2.1 FORMAT statement

11 R1001 *format-stmt* is FORMAT *format-specification*
12 R1002 *format-specification* is ([*format-item-list*])
13 or ([*format-item-list*,] *unlimited-format-item*)

14 C1001 (R1001) The *format-stmt* shall be labeled.

15 C1002 (R1002) The comma used to separate *format-items* in a *format-item-list* may be omitted

- 16 (1) between a P edit descriptor and an immediately following F, E, EN, ES, D, or G edit
17 descriptor (10.8.5), possibly preceded by a repeat specifier,
- 18 (2) before a slash edit descriptor when the optional repeat specification is not present (10.8.2),
- 19 (3) after a slash edit descriptor, or
- 20 (4) before or after a colon edit descriptor (10.8.3)

21 Blank characters may precede the initial left parenthesis of the format specification. Additional blank
22 characters may appear at any point within the format specification, with no effect on the interpretation
23 of the format specification, except within a character string edit descriptor (10.9).

NOTE 10.1

Examples of FORMAT statements are:

```
5   FORMAT (1PE12.4, I10)
9   FORMAT (I12, /, ' Dates: ', 2 (2I3, I5))
```

24 10.2.2 Character format specification

25 A character expression used as a *format* in a formatted input/output statement shall evaluate to a
26 character string whose leading part is a valid format specification.

NOTE 10.2

The format specification begins with a left parenthesis and ends with a right parenthesis.

27 All character positions up to and including the final right parenthesis of the format specification shall be

1 defined at the time the input/output statement is executed, and shall not become redefined or undefined
 2 during the execution of the statement. Character positions, if any, following the right parenthesis that
 3 ends the format specification need not be defined and may contain any character data with no effect on
 4 the interpretation of the format specification.

5 If the *format* is a character array, it is treated as if all of the elements of the array were specified in array
 6 element order and were concatenated. However, if a *format* is a character array element, the format
 7 specification shall be entirely within that array element.

NOTE 10.3

If a character constant is used as a *format* in an input/output statement, care shall be taken that the value of the character constant is a valid format specification. In particular, if a format specification delimited by apostrophes contains a character constant edit descriptor delimited with apostrophes, two apostrophes shall be written to delimit the edit descriptor and four apostrophes shall be written for each apostrophe that occurs within the edit descriptor. For example, the text:

```
2 ISN'T 3
```

may be written by various combinations of output statements and format specifications:

```
WRITE (6, 100) 2, 3
100 FORMAT (1X, I1, 1X, 'ISN''T', 1X, I1)
WRITE (6, '(1X, I1, 1X, ''ISN''''T'', 1X, I1)') 2, 3
WRITE (6, '(A)' ) ' 2 ISN''T 3'
```

Doubling of internal apostrophes usually can be avoided by using quotation marks to delimit the format specification and doubling of internal quotation marks usually can be avoided by using apostrophes as delimiters.

8 10.3 Form of a format item list

9 10.3.1 Syntax

10	R1003	<i>format-item</i>	is [<i>r</i>] <i>data-edit-desc</i>
11			or <i>control-edit-desc</i>
12			or <i>char-string-edit-desc</i>
13			or [<i>r</i>] (<i>format-item-list</i>)
14	R1004	<i>unlimited-format-item</i>	is * (<i>format-item-list</i>)
15	R1005	<i>r</i>	is <i>int-literal-constant</i>

16 C1003 (R1005) *r* shall be positive.

17 C1004 (R1005) *r* shall not have a kind parameter specified for it.

18 The integer literal constant *r* is called a **repeat specification**.

19 10.3.2 Edit descriptors

20 An **edit descriptor** is a **data edit descriptor**, a **control edit descriptor**, or a **character string**
 21 **edit descriptor**.

22	R1006	<i>data-edit-desc</i>	is I <i>w</i> [. <i>m</i>]
23			or B <i>w</i> [. <i>m</i>]
24			or O <i>w</i> [. <i>m</i>]

1		or $Z w [. m]$
2		or $F w . d$
3		or $E w . d [E e]$
4		or $EN w . d [E e]$
5		or $ES w . d [E e]$
6		or $G w [. d [E e]]$
7		or $L w$
8		or $A [w]$
9		or $D w . d$
10		or $DT [char-literal-constant] [(v-list)]$
11	R1007	w is <i>int-literal-constant</i>
12	R1008	m is <i>int-literal-constant</i>
13	R1009	d is <i>int-literal-constant</i>
14	R1010	e is <i>int-literal-constant</i>
15	R1011	v is <i>signed-int-literal-constant</i>
16	C1005	(R1010) e shall be positive.
17	C1006	(R1007) w shall be zero or positive for the I, B, O, Z, F, and G edit descriptors. w shall be
18		positive for all other edit descriptors.
19	C1007	(R1006) For the G edit descriptor, d shall be specified if and only if w is not zero.
20	C1008	(R1006) w , m , d , e , and v shall not have kind parameters specified for them.
21	C1009	(R1006) The <i>char-literal-constant</i> in the DT edit descriptor shall not have a kind parameter
22		specified for it.
23		I, B, O, Z, F, E, EN, ES, G, L, A, D, and DT indicate the manner of editing.
24	R1012	<i>control-edit-desc</i> is <i>position-edit-desc</i>
25		or $[r] /$
26		or $:$
27		or <i>sign-edit-desc</i>
28		or $k P$
29		or <i>blank-interp-edit-desc</i>
30		or <i>round-edit-desc</i>
31		or <i>decimal-edit-desc</i>
32	R1013	k is <i>signed-int-literal-constant</i>
33	C1010	(R1013) k shall not have a kind parameter specified for it.
34		In $k P$, k is called the scale factor .
35	R1014	<i>position-edit-desc</i> is $T n$
36		or $TL n$
37		or $TR n$
38		or $n X$
39	R1015	n is <i>int-literal-constant</i>
40	C1011	(R1015) n shall be positive.
41	C1012	(R1015) n shall not have a kind parameter specified for it.
42	R1016	<i>sign-edit-desc</i> is SS
43		or SP
44		or S

1	R1017	<i>blank-interp-edit-desc</i>	is	BN
2			or	BZ
3	R1018	<i>round-edit-desc</i>	is	RU
4			or	RD
5			or	RZ
6			or	RN
7			or	RC
8			or	RP
9	R1019	<i>decimal-edit-desc</i>	is	DC
10			or	DP

11 T, TL, TR, X, slash, colon, SS, SP, S, P, BN, BZ, RU, RD, RZ, RN, RC, RP, DC, and DP indicate the
12 manner of editing.

13 R1020 *char-string-edit-desc* **is** *char-literal-constant*

14 C1013 (R1020) The *char-literal-constant* shall not have a kind parameter specified for it.

15 Each *rep-char* in a character string edit descriptor shall be one of the characters capable of representation
16 by the processor.

17 The character string edit descriptors provide constant data to be output, and are not valid for input.

18 The edit descriptors are without regard to case except for the characters in the character constants.

19 **10.3.3 Fields**

20 A **field** is a part of a record that is read on input or written on output when format control encounters
21 a data edit descriptor or a character string edit descriptor. The **field width** is the size in characters of
22 the field.

23 **10.4 Interaction between input/output list and format**

24 The start of formatted data transfer using a format specification initiates **format control** (9.5.4.4.2).
25 Each action of format control depends on information jointly provided by the next edit descriptor in the
26 format specification and the next effective item in the input/output list, if one exists.

27 If an input/output list specifies at least one effective list item, at least one data edit descriptor shall
28 exist in the format specification.

NOTE 10.4

An empty format specification of the form () may be used only if the input/output list has no effective list items (9.5.4.4). Zero length character items are effective list items, but zero sized arrays and implied DO lists with an iteration count of zero are not.

29 A format specification is interpreted from left to right. The exceptions are format items preceded by a
30 repeat specification *r*, and format reversion (described below).

31 A format item preceded by a repeat specification is processed as a list of *r* items, each identical to the
32 format item but without the repeat specification and separated by commas.

NOTE 10.5

An omitted repeat specification is treated in the same way as a repeat specification whose value is one.

- 1 To each data edit descriptor interpreted in a format specification, there corresponds one effective item
 2 specified by the input/output list (9.5.3), except that an input/output list item of type complex requires
 3 the interpretation of two F, E, EN, ES, D, or G edit descriptors. For each control edit descriptor or
 4 character edit descriptor, there is no corresponding item specified by the input/output list, and format
 5 control communicates information directly with the record.
- 6 Whenever format control encounters a data edit descriptor in a format specification, it determines
 7 whether there is a corresponding effective item specified by the input/output list. If there is such an
 8 item, it transmits appropriately edited information between the item and the record, and then format
 9 control proceeds. If there is no such item, format control terminates.
- 10 If format control encounters a colon edit descriptor in a format specification and another effective in-
 11 put/output list item is not specified, format control terminates.
- 12 If format control encounters the rightmost parenthesis of a complete format specification and another
 13 effective input/output list item is not specified, format control terminates. However, if another effective
 14 input/output list item is specified, format control then reverts to the beginning of the format item
 15 terminated by the last preceding right parenthesis that is not part of a DT edit descriptor. If there
 16 is no such preceding right parenthesis, format control reverts to the first left parenthesis of the format
 17 specification. If any reversion occurs, the reused portion of the format specification shall contain at
 18 least one data edit descriptor. If format control reverts to a parenthesis that is preceded by a repeat
 19 specification, the repeat specification is reused. Reversion of format control, of itself, has no effect on
 20 the changeable modes (9.4.1). If format control reverts to a parenthesis that is not the beginning of an
 21 *unlimited-format-item*, the file is positioned in a manner identical to the way it is positioned when a
 22 slash edit descriptor is processed (10.8.2).

NOTE 10.6

Example: The format specification:

```
10 FORMAT (1X, 2(F10.3, I5))
```

with an output list of

```
WRITE (10,10) 10.1, 3, 4.7, 1, 12.4, 5, 5.2, 6
```

produces the same output as the format specification:

```
10 FORMAT (1X, F10.3, I5, F10.3, I5/F10.3, I5, F10.3, I5)
```

NOTE 10.7

The effect of an *unlimited-format-item* is as if its enclosed list were preceded by a very large repeat count. There is no file positioning implied by *unlimited-format-item* reversion. This may be used to write what is commonly called a comma separated value record.

For example,

```
WRITE( 10, '( "IARRAY =", *( I0, :, ",") )' ) IARRAY
```

produces a single record with a header and a comma separated list of integer values.

23 10.5 Positioning by format control

1 After each data edit descriptor or character string edit descriptor is processed, the file is positioned after
2 the last character read or written in the current record.

3 After each T, TL, TR, or X edit descriptor is processed, the file is positioned as described in 10.8.1.
4 After each slash edit descriptor is processed, the file is positioned as described in 10.8.2.

5 During formatted stream output, processing of an A edit descriptor can cause file positioning to occur
6 (10.7.4).

7 If format control reverts as described in 10.4, the file is positioned in a manner identical to the way it is
8 positioned when a slash edit descriptor is processed (10.8.2).

9 During a read operation, any unprocessed characters of the current record are skipped whenever the
10 next record is read.

11 **10.6 Decimal symbol**

12 The **decimal symbol** is the character that separates the whole and fractional parts in the decimal
13 representation of a real number in an internal or external file. When the decimal edit mode is POINT,
14 the decimal symbol is a decimal point. When the decimal edit mode is COMMA, the decimal symbol is
15 a comma.

16 **10.7 Data edit descriptors**

17 **10.7.1 General**

18 Data edit descriptors cause the conversion of data to or from its internal representation; during formatted
19 stream output, the A data edit descriptor may also cause file positioning. On input, the specified variable
20 becomes defined unless an error condition, an end-of-file condition, or an end-of-record condition occurs.
21 On output, the specified expression is evaluated.

22 During input from a Unicode file,

- 23 (1) characters in the record that correspond to an ASCII character variable shall have a position
24 in the ISO 10646 character type collating sequence of 127 or less, and
- 25 (2) characters in the record that correspond to a default character variable shall be representable
26 in the default character type.

27 During input from a non-Unicode file,

- 28 (1) characters in the record that correspond to a character variable shall have the kind of the
29 character variable, and
- 30 (2) characters in the record that correspond to a numeric logical, or bits variable shall be of
31 default character type.

32 During output to a Unicode file, all characters transmitted to the record are of ISO 10646 character
33 type. If a character input/output list item or character string edit descriptor contains a character that
34 is not representable in the ISO 10646 character type, the result is processor-dependent.

35 During output to a non-Unicode file, characters transmitted to the record as a result of processing a
36 character string edit descriptor or as a result of evaluating a numeric, logical, bits, or default character
37 data entity, are of type default character.

38 **10.7.2 Numeric and bits editing**

1 10.7.2.1 General rules

2 The I, B, O, Z, F, E, EN, ES, D, and G edit descriptors may be used to specify the input/output of
3 integer, real, complex, and bits data. The following general rules apply.

- 4 (1) On input, leading blanks are not significant. When the input field is not an IEEE excep-
5 tional specification (10.7.2.3.2), the interpretation of blanks, other than leading blanks, is
6 determined by the blank interpretation mode (10.8.6). Plus signs may be omitted. A field
7 containing only blanks is considered to be zero.
- 8 (2) On input, with F, E, EN, ES, D, and G editing, a decimal symbol appearing in the input
9 field overrides the portion of an edit descriptor that specifies the decimal symbol location.
10 The input field may have more digits than the processor uses to approximate the value of
11 the datum.
- 12 (3) On output with I, F, E, EN, ES, D, and G editing, the representation of a positive or zero
13 internal value in the field may be prefixed with a plus sign, as controlled by the S, SP, and
14 SS edit descriptors or the processor. The representation of a negative internal value in the
15 field shall be prefixed with a minus sign.
- 16 (4) On output, the representation is right justified in the field. If the number of characters
17 produced by the editing is smaller than the field width, leading blanks are inserted in the
18 field.
- 19 (5) On output, if the number of characters produced exceeds the field width or if an exponent
20 exceeds its specified length using the *Ew.d Ee*, *ENw.d Ee*, *ESw.d Ee*, or *Gw.d Ee* edit
21 descriptor, the processor shall fill the entire field of width *w* with asterisks. However,
22 the processor shall not produce asterisks if the field width is not exceeded when optional
23 characters are omitted.

NOTE 10.8

When an SP edit descriptor is in effect, a plus sign is not optional.

- 24 (6) On output, with I, B, O, Z, F, and G editing, the specified value of the field width *w* may
25 be zero. In such cases, the processor selects the smallest positive actual field width that
26 does not result in a field filled with asterisks. The specified value of *w* shall not be zero on
27 input.

28 10.7.2.2 Integer editing

29 The *Iw* and *Iw.m*, edit descriptors indicate that the field to be edited occupies *w* positions, except when
30 *w* is zero. When *w* is zero, the processor selects the field width. On input, *w* shall not be zero. The
31 specified input/output list item shall be of type integer or bits. The G, B, O, and Z edit descriptor also
32 may be used to edit integer data (10.7.5.2.1, 10.7.2.4).

33 If the input list item is of type bits, the integer value specified by the input field is converted to type
34 bits according to the model in 13.3. On input, *m* has no effect.

35 In the input field for the I edit descriptor, the character string shall be a *signed-digit-string* (R409) if
36 the list item is of type integer and a *digit-string* (R410) if it is of type bits, except for the interpretation
37 of blanks.

38 The output field for the *Iw* edit descriptor consists of zero or more leading blanks followed by a minus
39 sign if the internal value is negative, or an optional plus sign otherwise, followed by the magnitude of
40 the internal value as a *digit-string* without leading zeros.

NOTE 10.9

A *digit-string* always consists of at least one digit.

1 The output field for the $Iw.m$ edit descriptor is the same as for the Iw edit descriptor, except that the
 2 *digit-string* consists of at least m digits. If necessary, sufficient leading zeros are included to achieve the
 3 minimum of m digits. The value of m shall not exceed the value of w , except when w is zero. If m is
 4 zero and the internal value is zero, the output field consists of only blank characters, regardless of the
 5 sign control in effect. When m and w are both zero, and the internal value is zero, one blank character
 6 is produced.

7 10.7.2.3 Real and complex editing

8 10.7.2.3.1 General

9 The F, E, EN, ES, and D edit descriptors specify the editing of real and complex data. An input/output
 10 list item corresponding to an F, E, EN, ES, or D edit descriptor shall be real or complex. The G, B, O,
 11 and Z edit descriptors also may be used to edit real and complex data (10.7.5.2.2).

12 10.7.2.3.2 F editing

13 The $Fw.d$ edit descriptor indicates that the field occupies w positions, the fractional part of which
 14 consists of d digits. When w is zero, the processor selects the field width. On input, w shall not be zero.

15 A lower-case letter is equivalent to the corresponding upper-case letter in an IEEE exceptional specifi-
 16 cation or the exponent in a numeric input field.

17 The input field is either an IEEE exceptional specification or consists of an optional sign, followed by a
 18 string of one or more digits optionally containing a decimal symbol, including any blanks interpreted as
 19 zeros. The d has no effect on input if the input field contains a decimal symbol. If the decimal symbol
 20 is omitted, the rightmost d digits of the string, with leading zeros assumed if necessary, are interpreted
 21 as the fractional part of the value represented. The string of digits may contain more digits than a
 22 processor uses to approximate the value. The basic form may be followed by an exponent of one of the
 23 following forms:

- 24 • a *sign* followed by a *digit-string*;
- 25 • the letter E followed by zero or more blanks, followed by a *signed-digit-string*;
- 26 • the letter D followed by zero or more blanks, followed by a *signed-digit-string*.

27 An exponent containing a D is processed identically to an exponent containing an E.

NOTE 10.10

If the input field does not contain an exponent, the effect is as if the basic form were followed by
 an exponent with a value of $-k$, where k is the established scale factor (10.8.5).

28 An input field that is an IEEE exceptional specification consists of optional blanks, followed by either

- 29 • an optional sign, followed by the string 'INF' or the string 'INFINITY', or
- 30 • an optional sign, followed by the string 'NaN', optionally followed by zero or more alphanumeric
 31 characters enclosed in parentheses,

32 optionally followed by blanks.

33 The value specified by form (1) is an IEEE infinity; this form shall not be used if the processor does
 34 not support IEEE infinities for the input variable. The value specified by form (2) is an IEEE NaN;

- 1 this form shall not be used if the processor does not support IEEE NaNs for the input variable. The
 2 NaN value is a quiet NaN if the only nonblank characters in the field are 'NaN' or 'NaN()'; otherwise,
 3 the NaN value is processor-dependent. The interpretation of a sign in a NaN input field is processor
 4 dependent.
- 5 For an internal value that is an IEEE infinity, the output field consists of blanks, if necessary, followed
 6 by a minus sign for negative infinity or an optional plus sign otherwise, followed by the letters 'Inf' or
 7 'Infinity', right justified within the field. If w is less than 3, the field is filled with asterisks; otherwise,
 8 if w is less than 8, 'Inf' is produced.
- 9 For an internal value that is an IEEE NaN, the output field consists of blanks, if necessary, followed by
 10 the letters 'NaN' and optionally followed by one to $w - 5$ alphanumeric processor-dependent characters
 11 enclosed in parentheses, right justified within the field. If w is less than 3, the field is filled with asterisks.

NOTE 10.11

The processor-dependent characters following 'NaN' may convey additional information about that particular NaN.

- 12 For an internal value that is neither an IEEE infinity nor a NaN, the output field consists of blanks, if
 13 necessary, followed by a minus sign if the internal value is negative, or an optional plus sign otherwise,
 14 followed by a string of digits that contains a decimal symbol and represents the magnitude of the internal
 15 value, as modified by the established scale factor and rounded (10.7.2.3.7) to d fractional digits. Leading
 16 zeros are not permitted except for an optional zero immediately to the left of the decimal symbol if the
 17 magnitude of the value in the output field is less than one. The optional zero shall appear if there would
 18 otherwise be no digits in the output field.

10.7.2.3.3 E and D editing

- 19 The $Ew.d$, $Dw.d$, and $Ew.d Ee$ edit descriptors indicate that the external field occupies w positions, the
 20 fractional part of which consists of d digits, unless a scale factor greater than one is in effect, and the
 21 exponent part consists of e digits. The e has no effect on input.

- 22 The form and interpretation of the input field is the same as for $Fw.d$ editing (10.7.2.3.2).

- 23 For an internal value that is an IEEE infinity or NaN, the form of the output field is the same as for
 24 $Fw.d$.

- 25 For an internal value that is neither an IEEE infinity nor a NaN, the form of the output field for a scale
 26 factor of zero is

$$27 \quad [\pm] [0].x_1x_2 \dots x_d exp$$

- 28 where:

- 29 • \pm signifies a plus sign or a minus sign;
- 30 • $.$ signifies a decimal symbol (10.6);
- 31 • $x_1x_2 \dots x_d$ are the d most significant digits of the internal value after rounding (10.7.2.3.7);
- 32 • exp is a decimal exponent having one of the forms specified in table 10.1.

Table 10.1: **E and D exponent forms**

Edit Descriptor	Absolute Value of Exponent	Form of Exponent ¹
<i>Ew.d</i>	$ exp \leq 99$	$E\pm z_1 z_2$ or $\pm 0z_1 z_2$
	$99 < exp \leq 999$	$\pm z_1 z_2 z_3$
<i>Ew.d Ee</i>	$ exp \leq 10^e - 1$	$E\pm z_1 z_2 \dots z_e$
<i>Dw.d</i>	$ exp \leq 99$	$D\pm z_1 z_2$ or $E\pm z_1 z_2$ or $\pm 0z_1 z_2$
	$99 < exp \leq 999$	$\pm z_1 z_2 z_3$
(1) where each z is a digit.		

1 The sign in the exponent is produced. A plus sign is produced if the exponent value is zero. The edit
2 descriptor forms *Ew.d* and *Dw.d* shall not be used if $|exp| > 999$.

3 The scale factor k controls the decimal normalization (10.3.2, 10.8.5). If $-d < k \leq 0$, the output field
4 contains exactly $|k|$ leading zeros and $d - |k|$ significant digits after the decimal symbol. If $0 < k < d + 2$,
5 the output field contains exactly k significant digits to the left of the decimal symbol and $d - k + 1$
6 significant digits to the right of the decimal symbol. Other values of k are not permitted.

7 **10.7.2.3.4 EN editing**

8 The EN edit descriptor produces an output field in the form of a real number in engineering notation
9 such that the decimal exponent is divisible by three and the absolute value of the significand (R414) is
10 greater than or equal to 1 and less than 1000, except when the output value is zero. The scale factor
11 has no effect on output.

12 The forms of the edit descriptor are *ENw.d* and *ENw.d Ee* indicating that the external field occupies w
13 positions, the fractional part of which consists of d digits and the exponent part consists of e digits.

14 The form and interpretation of the input field is the same as for *Fw.d* editing (10.7.2.3.2).

15 For an internal value that is an IEEE infinity or NaN, the form of the output field is the same as for
16 *Fw.d*.

17 For an internal value that is neither an IEEE infinity nor a NaN, the form of the output field is

18 $[\pm] yyy . x_1 x_2 \dots x_d exp$

19 where:

- 20 • \pm signifies a plus sign or a minus sign;
- 21 • yyy are the 1 to 3 decimal digits representative of the most significant digits of the internal value
22 after rounding (10.7.2.3.7);
- 23 • yyy is an integer such that $1 \leq yyy < 1000$ or, if the output value is zero, $yyy = 0$;
- 24 • $.$ signifies a decimal symbol (10.6);
- 25 • $x_1 x_2 \dots x_d$ are the d next most significant digits of the internal value after rounding;
- 26 • exp is a decimal exponent, divisible by three, having one of the forms specified in table 10.2.

Table 10.2: EN exponent forms

Edit Descriptor	Absolute Value of Exponent	Form of Exponent ¹
EN <i>w.d</i>	$ exp \leq 99$	$E\pm z_1 z_2$ or $\pm 0z_1 z_2$
	$99 < exp \leq 999$	$\pm z_1 z_2 z_3$
EN <i>w.d Ee</i>	$ exp \leq 10^e - 1$	$E\pm z_1 z_2 \dots z_e$
(1) where each z is a digit.		

- 1 The sign in the exponent is produced. A plus sign is produced if the exponent value is zero. The edit
2 descriptor form EN*w.d* shall not be used if $|exp| > 999$.

NOTE 10.12

Examples:

Internal Value	Output field Using SS, EN12.3
6.421	6.421E+00
-.5	-500.000E-03
.00217	2.170E-03
4721.3	4.721E+03

3 10.7.2.3.5 ES editing

4 The ES edit descriptor produces an output field in the form of a real number in scientific notation such
5 that the absolute value of the significand (R414) is greater than or equal to 1 and less than 10, except
6 when the output value is zero. The scale factor has no effect on output.

7 The forms of the edit descriptor are ES*w.d* and ES*w.d Ee* indicating that the external field occupies w
8 positions, the fractional part of which consists of d digits and the exponent part consists of e digits.

9 The form and interpretation of the input field is the same as for *Fw.d* editing (10.7.2.3.2).

10 For an internal value that is an IEEE infinity or NaN, the form of the output field is the same as for
11 *Fw.d*.

12 For an internal value that is neither an IEEE infinity nor a NaN, the form of the output field is

13 $[\pm] y . x_1 x_2 \dots x_d exp$

14 where:

- 15 • \pm signifies a plus sign or a minus sign;
- 16 • y is a decimal digit representative of the most significant digit of the internal value after rounding
17 (10.7.2.3.7);
- 18 • $.$ signifies a decimal symbol (10.6);
- 19 • $x_1 x_2 \dots x_d$ are the d next most significant digits of the internal value after rounding;
- 20 • exp is a decimal exponent having one of the forms specified in table 10.3.

Table 10.3: **ES exponent forms**

Edit Descriptor	Absolute Value of Exponent	Form of Exponent ¹
<i>ESw.d</i>	$ exp \leq 99$	$E\pm z_1 z_2$ or $\pm 0z_1 z_2$
	$99 < exp \leq 999$	$\pm z_1 z_2 z_3$
<i>ESw.d Ee</i>	$ exp \leq 10^e - 1$	$E\pm z_1 z_2 \dots z_e$
(1) where each z is a digit.		

- 1 The sign in the exponent is produced. A plus sign is produced if the exponent value is zero. The edit
 2 descriptor form *ESw.d* shall not be used if $|exp| > 999$.

NOTE 10.13

Examples:

Internal Value	Output field Using SS, ES12.3
6.421	6.421E+00
-.5	-5.000E-01
.00217	2.170E-03
4721.3	4.721E+03

3 10.7.2.3.6 Complex editing

- 4 A complex datum consists of a pair of separate real data. The editing of a scalar datum of complex type
 5 is specified by two edit descriptors each of which specifies the editing of real data. The first of the edit
 6 descriptors specifies the real part; the second specifies the imaginary part. The two edit descriptors may
 7 be different. Control and character string edit descriptors may be processed between the edit descriptor
 8 for the real part and the edit descriptor for the imaginary part.

9 10.7.2.3.7 Rounding mode

- 10 The rounding mode can be specified by an OPEN statement (9.4.1), a data transfer input/output
 11 statement (9.5.2.13), or an edit descriptor (10.8.7).

- 12 In what follows, the term “decimal value” means the exact decimal number as given by the character
 13 string, while the term “internal value” means the number actually stored in the processor. For example,
 14 in dealing with the decimal constant 0.1, the decimal value is the mathematical quantity 1/10, which
 15 has no exact representation in binary form. Formatted output of real data involves conversion from an
 16 internal value to a decimal value; formatted input involves conversion from a decimal value to an internal
 17 value.

- 18 When the I/O rounding mode is UP, the value resulting from conversion shall be the smallest repre-
 19 sentable value that is greater than or equal to the original value. When the I/O rounding mode is
 20 DOWN, the value resulting from conversion shall be the largest representable value that is less than or
 21 equal to the original value. When the I/O rounding mode is ZERO, the value resulting from conversion
 22 shall be the value closest to the original value and no greater in magnitude than the original value. When
 23 the I/O rounding mode is NEAREST, the value resulting from conversion shall be the closer of the two
 24 nearest representable values if one is closer than the other. If the two nearest representable values are
 25 equidistant from the original value, it is processor dependent which one of them is chosen. When the
 26 I/O rounding mode is COMPATIBLE, the value resulting from conversion shall be the closer of the
 27 two nearest representable values or the value away from zero if halfway between them. When the I/O
 28 rounding mode is PROCESSOR_DEFINED, rounding during conversion shall be a processor dependent
 29 default mode, which may correspond to one of the other modes.

- 30 On processors that support IEEE rounding on conversions, NEAREST shall correspond to round to

1 nearest, as specified in the IEEE International Standard.

NOTE 10.14

On processors that support IEEE rounding on conversions, the I/O rounding modes COMPATIBLE and NEAREST will produce the same results except when the datum is halfway between the two representable values. In that case, NEAREST will pick the even value, but COMPATIBLE will pick the value away from zero. The I/O rounding modes UP, DOWN, and ZERO have the same effect as those specified in the IEEE International Standard for round toward $+\infty$, round toward $-\infty$, and round toward 0, respectively.

2 **10.7.2.4 Bits editing**

3 The *Bw*, *Bw.m*, *Ow*, *Ow.m*, *Zw*, and *Zw.m* edit descriptors indicate that the field to be edited occupies
4 *w* positions, except when *w* is zero. When *w* is zero, the processor selects the field width. On input,
5 *w* shall not be zero. The input/output list item shall be of type bits, integer, real, complex, or logical.
6 The G edit descriptor (10.7.5.3) or the I edit descriptor (10.7.2.2) also can be used to edit bits data.

7 If the input list item is not of type bits, the input field is edited as if the input list item were of type
8 bits and bits compatible (12.5.2.4) with the actual list item.

9 On input, *m* has no effect.

10 In the input field for the B, O, and Z edit descriptors the character string shall consist of binary, octal,
11 or hexadecimal digits (as in R426, R427, R428) in the respective input field. The lower-case hexadecimal
12 digits a through f in a hexadecimal input field are equivalent to the corresponding upper-case hexadecimal
13 digits.

14 If the output list item, *x*, is not of type bits, it is interpreted as if it were of type bits with the value
15 BITS(*x*).

16 The output field for the *Bw*, *Ow*, and *Zw* descriptors consists of zero or more leading blanks followed by
17 the internal value in a form identical to the digits of a binary, octal, or hexadecimal constant, respectively,
18 with the same value and without leading zeros.

NOTE 10.15

A binary, octal, or hexadecimal constant always consists of at least one digit or hexadecimal digit.

19 R1021 *hex-digit-string* **is** *hex-digit* [*hex-digit*] ...

20 The output field for the *Bw.m*, *Ow.m*, and *Zw.m* edit descriptor is the same as for the *Bw*, *Ow*, and *Zw*
21 edit descriptor, except that the *digit-string* or *hex-digit-string* consists of at least *m* digits. If necessary,
22 sufficient leading zeros are included to achieve the minimum of *m* digits. The value of *m* shall not exceed
23 the value of *w*, except when *w* is zero. If *m* is zero and the internal value consists of all zero bits, the
24 output field consists of only blank characters. When *m* and *w* are both zero, and the internal value
25 consists of all zero bits, one blank character is produced.

26 **10.7.3 Logical editing**

27 The *Lw* edit descriptor indicates that the field occupies *w* positions. The specified input/output list
28 item shall be of type logical. The G edit descriptor also may be used to edit logical data (10.7.5.4).

29 The input field consists of optional blanks, optionally followed by a period, followed by a T for true or
30 F for false. The T or F may be followed by additional characters in the field, which are ignored.

31 A lower-case letter is equivalent to the corresponding upper-case letter in a logical input field.

NOTE 10.16

The logical constants `.TRUE.` and `.FALSE.` are acceptable input forms.

1 The output field consists of $w - 1$ blanks followed by a T or F, depending on whether the internal value
2 is true or false, respectively.

3 10.7.4 Character editing

4 The `A[w]` edit descriptor is used with an input/output list item of type character. The `G` edit descriptor
5 also may be used to edit character data (10.7.5.5). The kind type parameter of all characters transferred
6 and converted under control of one `A` or `G` edit descriptor is implied by the kind of the corresponding
7 list item.

8 If a field width w is specified with the `A` edit descriptor, the field consists of w characters. If a field
9 width w is not specified with the `A` edit descriptor, the number of characters in the field is the length of
10 the corresponding list item, regardless of the value of the kind type parameter.

11 Let len be the length of the input/output list item. If the specified field width w for an `A` edit descriptor
12 corresponding to an input item is greater than or equal to len , the rightmost len characters will be taken
13 from the input field. If the specified field width w is less than len , the w characters will appear left
14 justified with $len - w$ trailing blanks in the internal value.

15 If the specified field width w for an `A` edit descriptor corresponding to an output item is greater than
16 len , the output field will consist of $w - len$ blanks followed by the len characters from the internal value.
17 If the specified field width w is less than or equal to len , the output field will consist of the leftmost w
18 characters from the internal value.

NOTE 10.17

For nondefault character types, the blank padding character is processor dependent.

19 If the file is connected for stream access, the output may be split across more than one record if it
20 contains newline characters. A newline character is a nonblank character returned by the intrinsic
21 function `NEW_LINE`. Beginning with the first character of the output field, each character that is not
22 a newline is written to the current record in successive positions; each newline character causes file
23 positioning at that point as if by slash editing (the current record is terminated at that point, a new
24 empty record is created following the current record, this new record becomes the last and current record
25 of the file, and the file is positioned at the beginning of this new record).

NOTE 10.18

If the intrinsic function `NEW_LINE` returns a blank character for a particular character kind, then
the processor does not support using a character of that kind to cause record termination in a
formatted stream file.

26 10.7.5 Generalized editing**27 10.7.5.1 Overview**

28 The `Gw`, `Gw.d` and `Gw.d Ee` edit descriptors are used with an input/output list item of any intrinsic
29 type. When w is nonzero, these edit descriptors indicate that the external field occupies w positions;
30 for real or complex data the fractional part consists of a maximum of d digits and the exponent part
31 consists of e digits. When these edit descriptors are used to specify the input/output of integer, logical,
32 bits, or character data, d and e have no effect. When w is zero the processor selects the field width. On
33 input, w shall not be zero.

1 **10.7.5.2 Generalized numeric editing**

2 When used to specify the input/output of integer, real, and complex data, the *Gw*, *Gw.d* and *Gw.d Ee*
3 edit descriptors follow the general rules for numeric editing (10.7.2).

NOTE 10.19

The *Gw.d Ee* edit descriptor follows any additional rules for the *Ew.d Ee* edit descriptor.

4 **10.7.5.2.1 Generalized integer editing**

5 When used to specify the input/output of integer data, the *Gw.d* and *Gw.d Ee* edit descriptors follow
6 the rules for the *Iw* edit descriptor (10.7.2.2), except that *w* shall not be zero. When used to specify the
7 output of integer data, the *G0* edit descriptor follows the rules for the *I0* edit descriptor.

8 **10.7.5.2.2 Generalized real and complex editing**

9 The form and interpretation of the input field is the same as for *Fw.d* editing (10.7.2.3.2).

10 When used to specify the output of real or complex data, the *G0* edit descriptor follows the rules for
11 the *ESw.d Ee* edit descriptor. Reasonable processor-dependent values of *w*, *d*, and *e* are used with each
12 output value.

13 For an internal value that is an IEEE infinity or NaN, the form of the output field for the *Gw.d* and
14 *Gw.d Ee* edit descriptors is the same as for *Fw.d*.

15 Otherwise, the method of representation in the output field depends on the magnitude of the internal
16 value being edited. Let *N* be the magnitude of the internal value and *r* be the rounding mode value
17 defined in the table below. If $0 < N < 0.1 - r \times 10^{-d-1}$ or $N \geq 10^d - r$, or *N* is identically 0 and
18 *d* is 0, *Gw.d* output editing is the same as *k* *PEw.d* output editing and *Gw.d Ee* output editing is
19 the same as *k* *PEw.d Ee* output editing, where *k* is the scale factor (10.8.5) currently in effect. If
20 $0.1 - r \times 10^{-d-1} \leq N < 10^d - r$ or *N* is identically 0 and *d* is not zero, the scale factor has no effect,
21 and the value of *N* determines the editing as follows:

Magnitude of Internal Value	Equivalent Conversion
$N = 0$	$F(w - n).(d - 1), n('b')$
$0.1 - r \times 10^{-d-1} \leq N < 1 - r \times 10^{-d}$	$F(w - n).d, n('b')$
$1 - r \times 10^{-d} \leq N < 10 - r \times 10^{-d+1}$	$F(w - n).(d - 1), n('b')$
$10 - r \times 10^{-d+1} \leq N < 100 - r \times 10^{-d+2}$	$F(w - n).(d - 2), n('b')$
.	.
.	.
.	.
$10^{d-2} - r \times 10^{-2} \leq N < 10^{d-1} - r \times 10^{-1}$	$F(w - n).1, n('b')$
$10^{d-1} - r \times 10^{-1} \leq N < 10^d - r$	$F(w - n).0, n('b')$

22 where *b* is a blank, *n* is 4 for *Gw.d* and *e* + 2 for *Gw.d Ee*, and *r* is defined for each rounding mode as
23 follows:

I/O Rounding Mode	<i>r</i>
COMPATIBLE	0.5
NEAREST	0.5 if the higher value is even -0.5 if the lower value is even
UP	1
DOWN	0

I/O Rounding Mode	<i>r</i>
ZERO	1 if internal value is negative 0 if internal value is positive

1 The value of $w - n$ shall be positive

NOTE 10.20

The scale factor has no effect on output unless the magnitude of the datum to be edited is outside the range that permits effective use of F editing.

2 **10.7.5.3 Generalized bits editing**

3 When used to specify the input/output of bits data, the *Gw.d* and *Gw.d Ee* edit descriptors follow the
4 rules for the *Zw* edit descriptor (10.7.2.4), except that w shall not be zero. When used to specify the
5 output of bits data, the *G0* edit descriptor follows the rules for the *Z0* edit descriptor.

6 **10.7.5.4 Generalized logical editing**

7 When used to specify the input/output of logical data, the *Gw.d* and *Gw.d Ee* edit descriptors follow
8 the rules for the *Lw* edit descriptor (10.7.3). When used to specify the output of logical data, the *G0*
9 edit descriptor follows the rules for the *L1* edit descriptor.

10 **10.7.5.5 Generalized character editing**

11 When used to specify the input/output of character data, the *Gw.d* and *Gw.d Ee* edit descriptors follow
12 the rules for the *Aw* edit descriptor (10.7.4). When used to specify the output of character data, the *G0*
13 edit descriptor follows the rules for the *A* edit descriptor with no field width.

14 **10.7.6 User-defined derived-type editing**

15 The *DT* edit descriptor allows a user-provided procedure to be used instead of the processor's default
16 input/output formatting for processing a list item of derived type.

17 The *DT* edit descriptor may include a character literal constant. The character value "DT" concatenated
18 with the character literal constant is passed to the user-defined derived-type input/output procedure as
19 the *iotype* argument (9.5.4.7). The v values of the edit descriptor are passed to the user-defined
20 derived-type input/output procedure as the *v_list* array argument.

NOTE 10.21

For the edit descriptor *DT'Link List'(10, 4, 2)*, *iotype* is "DTLink List" and *v_list* is (/10, 4, 2/).

21 If a derived-type variable or value corresponds to a *DT* edit descriptor, there shall be an accessible
22 interface to a corresponding derived-type input/output procedure for that derived type (9.5.4.7). A *DT*
23 edit descriptor shall not correspond with a list item that is not of a derived type.

24 **10.8 Control edit descriptors**

25 A control edit descriptor does not cause the transfer of data or the conversion of data to or from internal
26 representation, but may affect the conversions performed by subsequent data edit descriptors.

1 10.8.1 Position editing

2 The T, TL, TR, and X edit descriptors specify the position at which the next character will be transmitted
3 to or from the record. If any character skipped by a T, TL, TR, or X edit descriptor is of type nondefault
4 character, and the unit is an internal file of type default character or an external non-Unicode file, the
5 result of that position editing is processor dependent.

6 The position specified by a T edit descriptor may be in either direction from the current position. On
7 input, this allows portions of a record to be processed more than once, possibly with different editing.

8 The position specified by an X edit descriptor is forward from the current position. On input, a position
9 beyond the last character of the record may be specified if no characters are transmitted from such
10 positions.

NOTE 10.22

An nX edit descriptor has the same effect as a TR_n edit descriptor.

11 On output, a T, TL, TR, or X edit descriptor does not by itself cause characters to be transmitted and
12 therefore does not by itself affect the length of the record. If characters are transmitted to positions at
13 or after the position specified by a T, TL, TR, or X edit descriptor, positions skipped and not previously
14 filled are filled with blanks. The result is as if the entire record were initially filled with blanks.

15 On output, a character in the record may be replaced. However, a T, TL, TR, or X edit descriptor never
16 directly causes a character already placed in the record to be replaced. Such edit descriptors may result
17 in positioning such that subsequent editing causes a replacement.

18 10.8.1.1 T, TL, and TR editing

19 The **left tab limit** affects file positioning by the T and TL edit descriptors. Immediately prior to
20 nonchild data transfer, the left tab limit becomes defined as the character position of the current record
21 or the current position of the stream file. If, during data transfer, the file is positioned to another record,
22 the left tab limit becomes defined as character position one of that record.

23 The T_n edit descriptor indicates that the transmission of the next character to or from a record is to
24 occur at the n th character position of the record, relative to the left tab limit.

25 The TL_n edit descriptor indicates that the transmission of the next character to or from the record is
26 to occur at the character position n characters backward from the current position. However, if n is
27 greater than the difference between the current position and the left tab limit, the TL_n edit descriptor
28 indicates that the transmission of the next character to or from the record is to occur at the left tab
29 limit.

30 The TR_n edit descriptor indicates that the transmission of the next character to or from the record is
31 to occur at the character position n characters forward from the current position.

NOTE 10.23

The n in a T_n , TL_n , or TR_n edit descriptor shall be specified and shall be greater than zero.

32 10.8.1.2 X editing

33 The nX edit descriptor indicates that the transmission of the next character to or from a record is to
34 occur at the position n characters forward from the current position.

NOTE 10.24

The n in an nX edit descriptor shall be specified and shall be greater than zero.

1 10.8.2 Slash editing

2 The slash edit descriptor indicates the end of data transfer to or from the current record.

3 On input from a file connected for sequential or stream access, the remaining portion of the current
4 record is skipped and the file is positioned at the beginning of the next record. This record becomes
5 the current record. On output to a file connected for sequential or stream access, a new empty record
6 is created following the current record; this new record then becomes the last and current record of the
7 file and the file is positioned at the beginning of this new record.

8 For a file connected for direct access, the record number is increased by one and the file is positioned
9 at the beginning of the record that has that record number, if there is such a record, and this record
10 becomes the current record.

NOTE 10.25

A record that contains no characters may be written on output. If the file is an internal file or a file connected for direct access, the record is filled with blank characters.

An entire record may be skipped on input.

11 The repeat specification is optional in the slash edit descriptor. If it is not specified, the default value is
12 one.

13 10.8.3 Colon editing

14 The colon edit descriptor terminates format control if there are no more effective items in the in-
15 put/output list (9.5.3). The colon edit descriptor has no effect if there are more effective items in the
16 input/output list.

17 10.8.4 SS, SP, and S editing

18 The SS, SP, and S edit descriptors temporarily change (9.4.1) the sign mode (9.4.5.15, 9.5.2.14) for the
19 connection. The edit descriptors SS, SP, and S set the sign mode corresponding to the SIGN= specifier
20 values SUPPRESS, PLUS, and PROCESSOR_DEFINED, respectively.

21 The sign mode controls optional plus characters in numeric output fields. When the sign mode is PLUS,
22 the processor shall produce a plus sign in any position that normally contains an optional plus sign.
23 When the sign mode is SUPPRESS, the processor shall not produce a plus sign in such positions. When
24 the sign mode is PROCESSOR_DEFINED, the processor has the option of producing a plus sign or not
25 in such positions, subject to 10.7.2(5).

26 The SS, SP, and S edit descriptors affect only I, F, E, EN, ES, D, and G editing during the execution of
27 an output statement. The SS, SP, and S edit descriptors have no effect during the execution of an input
28 statement.

29 10.8.5 P editing

30 The kP edit descriptor temporarily changes (9.4.1) the scale factor for the connection to k . The scale
31 factor affects the editing of F, E, EN, ES, D, and G edit descriptors for numeric quantities.

32 The scale factor k affects the appropriate editing in the following manner.

- 1 (1) On input, with F, E, EN, ES, D, and G editing (provided that no exponent exists in the
2 field) and F output editing, the scale factor effect is that the externally represented number
3 equals the internally represented number multiplied by 10^k .
- 4 (2) On input, with F, E, EN, ES, D, and G editing, the scale factor has no effect if there is an
5 exponent in the field.
- 6 (3) On output, with E and D editing, the significand (R414) part of the quantity to be produced
7 is multiplied by 10^k and the exponent is reduced by k .
- 8 (4) On output, with G editing, the effect of the scale factor is suspended unless the magnitude
9 of the datum to be edited is outside the range that permits the use of F editing. If the use
10 of E editing is required, the scale factor has the same effect as with E output editing.
- 11 (5) On output, with EN and ES editing, the scale factor has no effect.

12 If UP, DOWN, ZERO, or NEAREST I/O rounding mode is in effect,

- 13 (1) on input, the scale factor is applied to the external decimal value and then this is converted
14 using the current I/O rounding mode, and
- 15 (2) on output, the internal value is converted using the current I/O rounding mode and then
16 the scale factor is applied to the converted decimal value.

17 **10.8.6 BN and BZ editing**

18 The BN and BZ edit descriptors temporarily change (9.4.1) the blank interpretation mode (9.4.5.4,
19 9.5.2.6) for the connection. The edit descriptors BN and BZ set the blank interpretation mode corre-
20 sponding to the BLANK= specifier values NULL and ZERO, respectively.

21 The blank interpretation mode controls the interpretation of nonleading blanks in numeric and bits
22 input fields. Such blank characters are interpreted as zeros when the blank interpretation mode has the
23 value ZERO; they are ignored when the blank interpretation mode has the value NULL. The effect of
24 ignoring blanks is to treat the input field as if blanks had been removed, the remaining portion of the
25 field right justified, and the blanks replaced as leading blanks. However, a field containing only blanks
26 has the value zero.

27 The blank interpretation mode affects only numeric editing, bits editing, generalized numeric editing,
28 and generalized bits editing on input. It has no effect on output.

29 **10.8.7 RU, RD, RZ, RN, RC, and RP editing**

30 The round edit descriptors temporarily change (9.4.1) the connection's I/O rounding mode (9.4.5.14,
31 9.5.2.13, 10.7.2.3.7). The round edit descriptors RU, RD, RZ, RN, RC, and RP set the I/O rounding
32 mode corresponding to the ROUND= specifier values UP, DOWN, ZERO, NEAREST, COMPATIBLE,
33 and PROCESSOR_DEFINED, respectively. The I/O rounding mode affects the conversion of real and
34 complex values in formatted input/output. It affects only D, E, EN, ES, F, and G editing.

35 **10.8.8 DC and DP editing**

36 The decimal edit descriptors temporarily change (9.4.1) the decimal edit mode (9.4.5.5, 9.5.2.7, 10.6)
37 for the connection. The edit descriptors DC and DP set the decimal edit mode corresponding to the
38 DECIMAL= specifier values COMMA and POINT, respectively.

39 The decimal edit mode controls the representation of the decimal symbol (10.6) during conversion of
40 real and complex values in formatted input/output. The decimal edit mode affects only D, E, EN, ES,
41 F, and G editing. If the mode is COMMA during list-directed input/output, the character used as a
42 value separator is a semicolon in place of a comma.

1 10.9 Character string edit descriptors

2 A character string edit descriptor shall not be used on input.

3 The character string edit descriptor causes characters to be written from the enclosed characters of the
4 edit descriptor itself, including blanks. For a character string edit descriptor, the width of the field is
5 the number of characters between the delimiting characters. Within the field, two consecutive delimiting
6 characters are counted as a single character.

NOTE 10.26

A delimiter for a character string edit descriptor is either an apostrophe or quote.

7 10.10 List-directed formatting

8 10.10.1 General

9 List-directed input/output allows data editing according to the type of the list item instead of by a
10 format specification. It also allows data to be free-field, that is, separated by commas (or semicolons)
11 or blanks.

12 10.10.2 Values and value separators

13 The characters in one or more list-directed records constitute a sequence of values and value separators.
14 The end of a record has the same effect as a blank character, unless it is within a character constant. Any
15 sequence of two or more consecutive blanks is treated as a single blank, unless it is within a character
16 constant.

17 Each value is either a null value, c , r^*c , or r^* , where c is a literal constant, optionally signed if integer
18 or real, or an undelimited character constant and r is an unsigned, nonzero, integer literal constant.
19 Neither c nor r shall have kind type parameters specified. The constant c is interpreted as though
20 it had the same kind type parameter as the corresponding list item. The r^*c form is equivalent to r
21 successive appearances of the constant c , and the r^* form is equivalent to r successive appearances of
22 the null value. Neither of these forms may contain embedded blanks, except where permitted within the
23 constant c .

24 A **value separator** is

- 25 (1) a comma optionally preceded by one or more contiguous blanks and optionally followed by
26 one or more contiguous blanks, unless the decimal edit mode is COMMA, in which case a
27 semicolon is used in place of the comma,
- 28 (2) a slash optionally preceded by one or more contiguous blanks and optionally followed by
29 one or more contiguous blanks, or
- 30 (3) one or more contiguous blanks between two nonblank values or following the last nonblank
31 value, where a nonblank value is a constant, an r^*c form, or an r^* form.

NOTE 10.27

Although a slash encountered in an input record is referred to as a separator, it actually causes
termination of list-directed and namelist input statements; it does not actually separate two values.

NOTE 10.28

If no list items are specified in a list-directed input/output statement, one input record is skipped
or one empty output record is written.

1 10.10.3 List-directed input

2 Input forms acceptable to edit descriptors for a given type are acceptable for list-directed formatting,
 3 except as noted below. The form of the input value shall be acceptable for the type of the next effective
 4 item in the list. Blanks are never used as zeros, and embedded blanks are not permitted in constants,
 5 except within character constants and complex constants as specified below.

6 For the r^*c form of an input value, the constant c is interpreted as an undelimited character constant
 7 if the first list item corresponding to this value is of type default, ASCII, or ISO 10646 character, there
 8 is a nonblank character immediately after r^* , and that character is not an apostrophe or a quotation
 9 mark; otherwise, c is interpreted as a literal constant.

NOTE 10.29

The end of a record has the effect of a blank, except when it appears within a character constant.

10 When the next effective item is of type integer, the value in the input record is interpreted as if an Iw
 11 edit descriptor with a suitable value of w were used.

12 When the next effective item is of type bits, the value in the input record is interpreted as if a Zw edit
 13 descriptor with a suitable value of w were used.

14 When the next effective item is of type real, the input form is that of a numeric input field. A numeric
 15 input field is a field suitable for F editing (10.7.2.3.2) that is assumed to have no fractional digits unless
 16 a decimal symbol appears within the field.

17 When the next effective item is of type complex, the input form consists of a left parenthesis followed by
 18 an ordered pair of numeric input fields separated by a comma (if the decimal edit mode is POINT) or
 19 semicolon (if the decimal edit mode is COMMA), and followed by a right parenthesis. The first numeric
 20 input field is the real part of the complex constant and the second is the imaginary part. Each of the
 21 numeric input fields may be preceded or followed by any number of blanks and ends of records. The end
 22 of a record may occur after the real part or before the imaginary part.

23 When the next effective item is of type logical, the input form shall not include value separators among
 24 the optional characters permitted for L editing.

25 When the next effective item is of type character, the input form consists of a possibly delimited sequence
 26 of zero or more *rep-chars* whose kind type parameter is implied by the kind of the effective list item.
 27 Character sequences may be continued from the end of one record to the beginning of the next record,
 28 but the end of record shall not occur between a doubled apostrophe in an apostrophe-delimited character
 29 sequence, nor between a doubled quote in a quote-delimited character sequence. The end of the record
 30 does not cause a blank or any other character to become part of the character sequence. The character
 31 sequence may be continued on as many records as needed. The characters blank, comma, semicolon,
 32 and slash may appear in default, ASCII, or ISO 10646 character sequences.

33 If the next effective item is of type default, ASCII, or ISO 10646 character and

- 34 (1) the character sequence does not contain value separators,
- 35 (2) the character sequence does not cross a record boundary,
- 36 (3) the first nonblank character is not a quotation mark or an apostrophe,
- 37 (4) the leading characters are not *digits* followed by an asterisk, and
- 38 (5) the character sequence contains at least one character,

39 the delimiting apostrophes or quotation marks are not required. If the delimiters are omitted, the
 40 character sequence is terminated by the first blank, comma (if the decimal edit mode is POINT),
 41 semicolon (if the decimal edit mode is COMMA), slash, or end of record; in this case apostrophes and
 quotation marks within the datum are not to be doubled.

1 Let len be the length of the next effective item, and let w be the length of the character sequence. If
 2 len is less than or equal to w , the leftmost len characters of the sequence are transmitted to the next
 3 effective item. If len is greater than w , the sequence is transmitted to the leftmost w characters of the
 4 next effective item and the remaining $len-w$ characters of the next effective item are filled with blanks.
 5 The effect is as though the sequence were assigned to the next effective item in an intrinsic assignment
 6 statement (7.4.1.3).

7 10.10.3.1 Null values

8 A null value is specified by

- 9 (1) the r^* form,
- 10 (2) no characters between consecutive value separators, or
- 11 (3) no characters before the first value separator in the first record read by each execution of a
 12 list-directed input statement.

NOTE 10.30

The end of a record following any other value separator, with or without separating blanks, does not specify a null value in list-directed input.

13 A null value has no effect on the definition status of the next effective item. A null value shall not be
 14 used for either the real or imaginary part of a complex constant, but a single null value may represent
 15 an entire complex constant.

16 A slash encountered as a value separator during execution of a list-directed input statement causes
 17 termination of execution of that input statement after the transference of the previous value. Any
 18 characters remaining in the current record are ignored. If there are additional items in the input list, the
 19 effect is as if null values had been supplied for them. Any *do-variable* in the input list becomes defined
 20 as if enough null values had been supplied for any remaining input list items.

NOTE 10.31

All blanks in a list-directed input record are considered to be part of some value separator except for

- (1) blanks embedded in a character sequence,
- (2) embedded blanks surrounding the real or imaginary part of a complex constant, and
- (3) leading blanks in the first record read by each execution of a list-directed input statement, unless immediately followed by a slash or comma.

NOTE 10.32

List-directed input example:

```

INTEGER I; REAL X (8); CHARACTER (11) P;
COMPLEX Z; LOGICAL G
...
READ *, I, X, P, Z, G
...
```

The input data records are:

```

12345,12345,,2*1.5,4*
ISN'T_BOB'S,(123,0),.TEXAS$
```


NOTE 10.32 (cont.)

The results are:

Variable	Value
I	12345
X (1)	12345.0
X (2)	unchanged
X (3)	1.5
X (4)	1.5
X (5) – X (8)	unchanged
P	ISN'T_BOB'S
Z	(123.0,0.0)
G	true

1 10.10.4 List-directed output

2 The form of the values produced is the same as that required for input, except as noted otherwise. With
 3 the exception of adjacent undelimited character sequences, the values are separated by one or more
 4 blanks or by a comma, or a semicolon if the decimal edit mode is comma, optionally preceded by one or
 5 more blanks and optionally followed by one or more blanks.

6 The processor may begin new records as necessary, but the end of record shall not occur within a constant
 7 except as specified for complex constants and character sequences. The processor shall not insert blanks
 8 within character sequences or within constants, except as specified for complex constants.

9 Logical output values are T for the value true and F for the value false.

10 Integer output constants are produced with the effect of an *Iw* edit descriptor.

11 Bits output constants are produced with the effect of a *Zw.m* edit descriptor with *w* and *m* equal to
 12 CEILING($k/4.0$) where *k* is the kind type parameter value of the list item.

13 Real constants are produced with the effect of either an *F* edit descriptor or an *E* edit descriptor,
 14 depending on the magnitude *x* of the value and a range $10^{d_1} \leq x < 10^{d_2}$, where *d*₁ and *d*₂ are processor-
 15 dependent integers. If the magnitude *x* is within this range or is zero, the constant is produced using
 16 *0PFw.d*; otherwise, *1PEw.d Ee* is used.

17 For numeric output, reasonable processor-dependent values of *w*, *d*, and *e* are used for each of the
 18 numeric constants output.

19 Complex constants are enclosed in parentheses with a separator between the real and imaginary parts,
 20 each produced as defined above for real constants. The separator is a comma if the decimal edit mode is
 21 POINT; it is a semicolon if the decimal edit mode is COMMA. The end of a record may occur between
 22 the separator and the imaginary part only if the entire constant is as long as, or longer than, an entire
 23 record. The only embedded blanks permitted within a complex constant are between the separator and
 24 the end of a record and one blank at the beginning of the next record.

25 Character sequences produced when the delimiter mode has a value of NONE

- 26 (1) are not delimited by apostrophes or quotation marks,
 27 (2) are not separated from each other by value separators,
 28 (3) have each internal apostrophe or quotation mark represented externally by one apostrophe
 29 or quotation mark, and
 30 (4) have a blank character inserted by the processor at the beginning of any record that begins
 31 with the continuation of a character sequence from the preceding record.

- 1 Character sequences produced when the delimiter mode has a value of QUOTE are delimited by quotes,
2 are preceded and followed by a value separator, and have each internal quote represented on the external
3 medium by two contiguous quotes.
- 4 Character sequences produced when the delimiter mode has a value of APOSTROPHE are delimited
5 by apostrophes, are preceded and followed by a value separator, and have each internal apostrophe
6 represented on the external medium by two contiguous apostrophes.
- 7 If two or more successive values in an output record have identical values, the processor has the option
8 of producing a repeated constant of the form $r*c$ instead of the sequence of identical values.
- 9 Slashes, as value separators, and null values are not produced as output by list-directed formatting.
- 10 Except for continuation of delimited character sequences, each output record begins with a blank char-
11 acter.

NOTE 10.33

The length of the output records is not specified and may be processor dependent.

12 10.11 Namelist formatting

13 10.11.1 General

- 14 Namelist input/output allows data editing with NAME=value subsequences. This facilitates documen-
15 tation of input and output files and more flexibility on input.

16 10.11.2 Name-value subsequences

- 17 The characters in one or more namelist records constitute a sequence of **name-value subsequences**,
18 each of which consists of an object designator followed by an equals and followed by one or more values
19 and value separators. The equals may optionally be preceded or followed by one or more contiguous
20 blanks. The end of a record has the same effect as a blank character, unless it is within a character
21 constant. Any sequence of two or more consecutive blanks is treated as a single blank, unless it is within
22 a character constant.

- 23 The name may be any name in the *namelist-group-object-list* (5.6).

- 24 A value separator for namelist formatting is the same as for list-directed formatting (10.10).

25 10.11.3 Namelist input

- 26 Input for a namelist input statement consists of

- 27 (1) optional blanks and namelist comments,
- 28 (2) the character & followed immediately by the *namelist-group-name* as specified in the NAME-
29 LIST statement,
- 30 (3) one or more blanks,
- 31 (4) a sequence of zero or more name-value subsequences separated by value separators, and
- 32 (5) a slash to terminate the namelist input.

NOTE 10.34

A slash encountered in a namelist input record causes the input statement to terminate. A slash cannot be used to separate two values in a namelist input statement.

1 In each name-value subsequence, the name shall be the name of a namelist group object list item with
2 an optional qualification and the name with the optional qualification shall not be a zero-sized array, a
3 zero-sized array section, or a zero-length character string. The optional qualification, if any, shall not
4 contain a vector subscript.

5 A group name or object name is without regard to case.

6 **10.11.3.1 Namelist group object names**

7 Within the input data, each name shall correspond to a particular namelist group object name. Sub-
8 scripts, strides, and substring range expressions used to qualify group object names shall be optionally
9 signed integer literal constants with no kind type parameters specified. If a namelist group object is
10 an array, the input record corresponding to it may contain either the array name or the designator of
11 a subobject of that array, using the syntax of object designators (R603). If the namelist group object
12 name is the name of a variable of derived type, the name in the input record may be either the name of
13 the variable or the designator of one of its components, indicated by qualifying the variable name with
14 the appropriate component name. Successive qualifications may be applied as appropriate to the shape
15 and type of the variable represented.

16 The order of names in the input records need not match the order of the namelist group object items.
17 The input records need not contain all the names of the namelist group object items. The definition
18 status of any names from the *namelist-group-object-list* that do not occur in the input record remains
19 unchanged. In the input record, each object name or subobject designator may be preceded and followed
20 by one or more optional blanks but shall not contain embedded blanks.

21 **10.11.3.2 Namelist group object list items**

22 The name-value subsequences are evaluated serially, in left-to-right order. A namelist group object
23 designator may appear in more than one name-value sequence.

24 When the name in the input record represents an array variable or a variable of derived type, the effect
25 is as if the variable represented were expanded into a sequence of scalar list items, in the same way that
26 formatted input/output list items are expanded (9.5.3). Each input value following the equals shall then
27 be acceptable to format specifications for the type of the list item in the corresponding position in the
28 expanded sequence, except as noted in this subclause. The number of values following the equals shall
29 not exceed the number of list items in the expanded sequence, but may be less; in the latter case, the
30 effect is as if sufficient null values had been appended to match any remaining list items in the expanded
31 sequence.

NOTE 10.35

For example, if the name in the input record is the name of an integer array of size 100, at most 100 values, each of which is either a digit string or a null value, may follow the equals; these values would then be assigned to the elements of the array in array element order.

32 A slash encountered as a value separator during the execution of a namelist input statement causes
33 termination of execution of that input statement after transference of the previous value. If there are
34 additional items in the namelist group object being transferred, the effect is as if null values had been
35 supplied for them.

36 A namelist comment may appear after any value separator except a slash. A namelist comment is also
37 permitted to start in the first nonblank position of an input record except within a character literal
38 constant.

39 Successive namelist records are read by namelist input until a slash is encountered; the remainder of the
40 record is ignored and need not follow the rules for namelist input values.

1 10.11.3.3 Namelist input values

2 Each value is either a null value (10.11.3.4), c , $r*c$, or r^* , where c is a literal constant, optionally signed
3 if integer or real, and r is an unsigned, nonzero, integer literal constant. A kind type parameter shall not
4 be specified for c or r . The constant c is interpreted as though it had the same kind type parameter as
5 the corresponding effective item. The $r*c$ form is equivalent to r successive appearances of the constant
6 c , and the r^* form is equivalent to r successive null values. Neither of these forms may contain embedded
7 blanks, except where permitted within the constant c .

8 The datum c (10.11) is any input value acceptable to format specifications for a given type, except for
9 a restriction on the form of input values corresponding to list items of types logical, integer, bits, and
10 character as specified in this subclause. The form of a real or complex value is dependent on the decimal
11 edit mode in effect (10.6). The form of an input value shall be acceptable for the type of the namelist
12 group object list item. The number and forms of the input values that may follow the equals in a name-
13 value subsequence depend on the shape and type of the object represented by the name in the input
14 record. When the name in the input record is that of a scalar variable of an intrinsic type, the equals
15 shall not be followed by more than one value. Blanks are never used as zeros, and embedded blanks are
16 not permitted in constants except within character constants and complex constants as specified in this
17 subclause.

18 When the next effective namelist group object list item is of type real, the input form of the input value
19 is that of a numeric input field. A numeric input field is a field suitable for F editing (10.7.2.3.2) that is
20 assumed to have no fractional digits unless a decimal symbol appears within the field.

21 When the next effective item is of type complex, the input form of the input value consists of a left
22 parenthesis followed by an ordered pair of numeric input fields separated by a comma (if the decimal
23 edit mode is POINT) or a semicolon (if the decimal edit mode is COMMA), and followed by a right
24 parenthesis. The first numeric input field is the real part of the complex constant and the second part
25 is the imaginary part. Each of the numeric input fields may be preceded or followed by any number of
26 blanks and ends of records. The end of a record may occur between the real part and the comma or
27 semicolon, or between the comma or semicolon and the imaginary part.

28 When the next effective item is of type logical, the input form of the input value shall not include equals
29 or value separators among the optional characters permitted for L editing (10.7.3).

30 When the next effective item is of type integer, the value in the input record is interpreted as if an Iw
31 edit descriptor with a suitable value of w were used.

32 When the next effective item is of type bits, the value in the input record is interpreted as if a Zw edit
33 descriptor with a suitable value of w were used.

34 When the next effective item is of type character, the input form consists of a delimited sequence of zero
35 or more *rep-chars* whose kind type parameter is implied by the kind of the corresponding list item. Such
36 a sequence may be continued from the end of one record to the beginning of the next record, but the
37 end of record shall not occur between a doubled apostrophe in an apostrophe-delimited sequence, nor
38 between a doubled quote in a quote-delimited sequence. The end of the record does not cause a blank or
39 any other character to become part of the sequence. The sequence may be continued on as many records
40 as needed. The characters blank, comma, semicolon, and slash may appear in such character sequences.

NOTE 10.36

A character sequence corresponding to a namelist input item of character type shall be delimited either with apostrophes or with quotes. The delimiter is required to avoid ambiguity between undelimited character sequences and object names. The value of the DELIM= specifier, if any, in the OPEN statement for an external file is ignored during namelist input (9.4.5.6).

1 Let *len* be the length of the next effective item, and let *w* be the length of the character sequence. If
 2 *len* is less than or equal to *w*, the leftmost *len* characters of the sequence are transmitted to the next
 3 effective item. If *len* is greater than *w*, the constant is transmitted to the leftmost *w* characters of the
 4 next effective item and the remaining *len*−*w* characters of the next effective item are filled with blanks.
 5 The effect is as though the sequence were assigned to the next effective item in an intrinsic assignment
 6 statement (7.4.1.3).

7 10.11.3.4 Null values

8 A null value is specified by

- 9 (1) the *r** form,
- 10 (2) blanks between two consecutive nonblank value separators following an equals,
- 11 (3) zero or more blanks preceding the first value separator and following an equals, or
- 12 (4) two consecutive nonblank value separators.

13 A null value has no effect on the definition status of the corresponding input list item. If the namelist
 14 group object list item is defined, it retains its previous value; if it is undefined, it remains undefined. A
 15 null value shall not be used as either the real or imaginary part of a complex constant, but a single null
 16 value may represent an entire complex constant.

NOTE 10.37

The end of a record following a value separator, with or without intervening blanks, does not specify a null value in namelist input.

17 10.11.3.5 Blanks

18 All blanks in a namelist input record are considered to be part of some value separator except for

- 19 (1) blanks embedded in a character constant,
- 20 (2) embedded blanks surrounding the real or imaginary part of a complex constant,
- 21 (3) leading blanks following the equals unless followed immediately by a slash or comma, or a
 22 semicolon if the decimal edit mode is comma, and
- 23 (4) blanks between a name and the following equals.

24 10.11.3.6 Namelist Comments

25 Except within a character literal constant, a “!” character after a value separator or in the first nonblank
 26 position of a namelist input record initiates a comment. The comment extends to the end of the current
 27 input record and may contain any graphic character in the processor-dependent character set. The
 28 comment is ignored. A slash within the namelist comment does not terminate execution of the namelist
 29 input statement. Namelist comments are not allowed in stream input because comments depend on
 30 record structure.

NOTE 10.38

Namelist input example:

```
INTEGER I; REAL X (8); CHARACTER (11) P; COMPLEX Z;
LOGICAL G
NAMELIST / TODAY / G, I, P, Z, X
READ (*, NML = TODAY)
```

The input data records are:

NOTE 10.38 (cont.)

```
&TODAY I = 12345, X(1) = 12345, X(3:4) = 2*1.5, I=6, ! This is a comment.
P = ''ISN'T_BOB'S'', Z = (123,0)/
```

The results stored are:

Variable	Value
I	6
X (1)	12345.0
X (2)	unchanged
X (3)	1.5
X (4)	1.5
X (5) – X (8)	unchanged
P	ISN'T_BOB'S
Z	(123.0,0.0)
G	unchanged

1 10.11.4 Namelist output

2 The form of the output produced is the same as that required for input, except for the forms of real,
3 character, and logical values. The name in the output is in upper case. With the exception of adjacent
4 undelimited character values, the values are separated by one or more blanks or by a comma, or a
5 semicolon if the decimal edit mode is COMMA, optionally preceded by one or more blanks and optionally
6 followed by one or more blanks.

7 Namelist output shall not include namelist comments.

8 The processor may begin new records as necessary. However, except for complex constants and character
9 values, the end of a record shall not occur within a constant, character value, or name, and blanks shall
10 not appear within a constant, character value, or name.

NOTE 10.39

The length of the output records is not specified exactly and may be processor dependent.

11 10.11.4.1 Namelist output editing

12 Values in namelist output records are edited as for list-directed output (10.10.4).

NOTE 10.40

Namelist output records produced with a DELIM= specifier with a value of NONE and which contain a character sequence might not be acceptable as namelist input records.

13 10.11.4.2 Namelist output records

14 If two or more successive values for the same namelist group item in an output record produced have
15 identical values, the processor has the option of producing a repeated constant of the form r^*c instead
16 of the sequence of identical values.

17 The name of each namelist group object list item is placed in the output record followed by an equals
18 and a list of values of the namelist group object list item.

19 An ampersand character followed immediately by a *namelist-group-name* will be produced by namelist
20 formatting at the start of the first output record to indicate which particular group of data objects is
being output. A slash is produced by namelist formatting to indicate the end of the namelist formatting.

- 1 A null value is not produced by namelist formatting.
- 2 Except for new records created by explicit formatting within a user-defined derived-type output pro-
3 cedure or by continuation of delimited character sequences, each output record begins with a blank
4 character.

1 11 Program units

2 11.1 Main program

3 A Fortran **main program** is a program unit that does not contain a SUBROUTINE, FUNCTION,
4 MODULE, SUBMODULE, or BLOCK DATA statement as its first statement.

5 R1101 *main-program* is [*program-stmt*]
6 [*specification-part*]
7 [*execution-part*]
8 [*internal-subprogram-part*]
9 *end-program-stmt*
10 R1102 *program-stmt* is PROGRAM *program-name*
11 R1103 *end-program-stmt* is END [PROGRAM [*program-name*]]

12 C1101 (R1101) In a *main-program*, the *execution-part* shall not contain a RETURN statement or an
13 ENTRY statement.

14 C1102 (R1101) The *program-name* may be included in the *end-program-stmt* only if the optional
15 *program-stmt* is used and, if included, shall be identical to the *program-name* specified in the
16 *program-stmt*.

NOTE 11.1

The program name is global to the program (16.2). For explanatory information about uses for the program name, see subclause C.8.1.

NOTE 11.2

An example of a main program is:

```
PROGRAM ANALYZE
  REAL A, B, C (10,10)      ! Specification part
  CALL FIND                 ! Execution part
CONTAINS
  SUBROUTINE FIND           ! Internal subprogram
  ...
  END SUBROUTINE FIND
END PROGRAM ANALYZE
```

17 The main program may be defined by means other than Fortran; in that case, the program shall not
18 contain a *main-program* program unit.

19 A reference to a Fortran *main-program* shall not appear in any program unit in the program, including
20 itself.

21 11.2 Modules

1 11.2.1 General

2 A **module** contains specifications and definitions that are to be accessible to other program units by use
 3 association. A module that is provided as an inherent part of the processor is an **intrinsic module**. A
 4 **nonintrinsic module** is defined by a module program unit or a means other than Fortran.

5 Procedures and types defined in an intrinsic module are not themselves intrinsic.

6	R1104	<i>module</i>	is	<i>module-stmt</i>
7				[<i>specification-part</i>]
8				[<i>module-subprogram-part</i>]
9				<i>end-module-stmt</i>
10	R1105	<i>module-stmt</i>	is	MODULE <i>module-name</i>
11	R1106	<i>end-module-stmt</i>	is	END [MODULE [<i>module-name</i>]]
12	R1107	<i>module-subprogram-part</i>	is	<i>contains-stmt</i>
13				[<i>module-subprogram</i>] ...
14	R1108	<i>module-subprogram</i>	is	<i>function-subprogram</i>
15			or	<i>subroutine-subprogram</i>
16			or	<i>separate-module-subprogram</i>

17 C1103 (R1104) If the *module-name* is specified in the *end-module-stmt*, it shall be identical to the
 18 *module-name* specified in the *module-stmt*.

19 C1104 (R1104) A module *specification-part* shall not contain a *stmt-function-stmt*, an *entry-stmt*, or a
 20 *format-stmt*.

21 C1105 (R1104) If an object that has a default-initialized direct component is declared in the *specification-*
 22 *part* of a module it shall have the ALLOCATABLE, POINTER, or SAVE attribute.

23 C1106 If an object that has a co-array ultimate component is declared in the *specification-part* of a
 24 module it shall have the SAVE attribute.

J3 internal note

Unresolved Technical Issue 010

This is better, but still inconsistent.

After 06-238r1, top-level allocatable co-arrays are allowed in a module without being **SAVED**.
 Allow that but requiring **SAVE** for things with co-array components does not make sense.

06-238r1 claims “save is required for a module variable if it has a co-array component. We wish
 to retain this since otherwise an unexpected implicit synchronization would be needed whenever
 the module ceases to be active.” Why would that be unexpected? *I* would expect that when the
 module goes out of scope, the allocatable (co-array thingo) either remains (as is allowed) or is
 deallocated (as is allowed), with only the latter case requiring (implicit) synchronisation.

NOTE 11.3

The module name is global to the program (16.2).

NOTE 11.4

Although statement function definitions, **ENTRY** statements, and **FORMAT** statements shall not
 appear in the specification part of a module, they may appear in the specification part of a module
 subprogram in the module.

NOTE 11.5

For a discussion of the impact of modules on dependent compilation, see subclause C.8.2.

NOTE 11.6

For examples of the use of modules, see subclause C.8.3.

1 If a procedure declared in the scoping unit of a module has an implicit interface, it shall be given the
 2 EXTERNAL attribute in that scoping unit; if it is a function, its type and type parameters shall be
 3 explicitly declared in a type declaration statement in that scoping unit.

4 If an intrinsic procedure is declared in the scoping unit of a module, it shall explicitly be given the
 5 INTRINSIC attribute in that scoping unit or be used as an intrinsic procedure in that scoping unit.

6 11.2.2 The USE statement and use association

7 The **USE statement** specifies use association. A USE statement is a **module reference** to the module
 8 it specifies. At the time a USE statement is processed, the public portions of the specified module shall
 9 be available. A module shall not reference itself, either directly or indirectly. A submodule shall not
 10 reference its ancestor module by use association, either directly or indirectly.

NOTE 11.7

It is possible for submodules with different ancestor modules to reference each others' ancestor modules by use association.

11 The **USE statement** provides the means by which a scoping unit accesses named data objects, derived
 12 types, interface blocks, procedures, abstract interfaces, module procedure interfaces, generic identifiers,
 13 macros, and namelist groups in a module. The entities in the scoping unit are **use associated** with the
 14 entities in the module. The accessed entities have the attributes specified in the module, except that an
 15 entity may have a different accessibility attribute or it may have the ASYNCHRONOUS or VOLATILE
 16 attribute in the local scoping unit even if the associated module entity does not. The entities made
 17 accessible are identified by the names or generic identifiers used to identify them in the module. By
 18 default, they are identified by the same identifiers in the scoping unit containing the USE statement,
 19 but it is possible to specify that different local identifiers are used.

NOTE 11.8

The accessibility of module entities may be controlled by accessibility attributes (4.5.2.2, 5.3.2), and the ONLY option of the USE statement. Definability of module entities can be controlled by the PROTECTED attribute (5.3.14).

20 R1109 *use-stmt* **is** USE [[, *module-nature*] ::] *module-name* [, *rename-list*]
 21 **or** USE [[, *module-nature*] ::] *module-name* , ■
 22 ■ **ONLY** : [*only-list*]
 23 R1110 *module-nature* **is** INTRINSIC
 24 **or** NON_INTRINSIC
 25 R1111 *rename* **is** *local-name* => *use-name*
 26 **or** OPERATOR (*local-defined-operator*) => ■
 27 ■ OPERATOR (*use-defined-operator*)
 28 R1112 *only* **is** *generic-spec*
 29 **or** *only-use-name*
 30 **or** *rename*
 31 R1113 *only-use-name* **is** *use-name*

32 C1107 (R1109) If *module-nature* is INTRINSIC, *module-name* shall be the name of an intrinsic module.

33 C1108 (R1109) If *module-nature* is NON_INTRINSIC, *module-name* shall be the name of a nonintrinsic
 34 module.

- 1 C1109 (R1109) A scoping unit shall not access an intrinsic module and a nonintrinsic module of the
2 same name.
- 3 C1110 (R1111) OPERATOR(*use-defined-operator*) shall not identify a *generic-binding*.
- 4 C1111 (R1112) The *generic-spec* shall not identify a *generic-binding*.

NOTE 11.9

The above two constraints do not prevent accessing a *generic-spec* that is declared by an interface block, even if a *generic-binding* has the same *generic-spec*.

- 5 C1112 (R1112) Each *generic-spec* shall be a public entity in the module.
- 6 C1113 (R1113) Each *use-name* shall be the name of a public entity in the module.
- 7 R1114 *local-defined-operator* **is** *defined-unary-op*
8 **or** *defined-binary-op*
- 9 R1115 *use-defined-operator* **is** *defined-unary-op*
10 **or** *defined-binary-op*
- 11 C1114 (R1115) Each *use-defined-operator* shall be a public entity in the module.
- 12 A *use-stmt* without a *module-nature* provides access either to an intrinsic or to a nonintrinsic module.
13 If the *module-name* is the name of both an intrinsic and a nonintrinsic module, the nonintrinsic module
14 is accessed.
- 15 The USE statement without the ONLY option provides access to all public entities in the specified
16 module.
- 17 A USE statement with the ONLY option provides access only to those entities that appear as *generic-*
18 *specs*, *use-names*, or *use-defined-operators* in the *only-list*.
- 19 More than one USE statement for a given module may appear in a scoping unit. If one of the USE
20 statements is without an ONLY option, all public entities in the module are accessible. If all the USE
21 statements have ONLY options, only those entities in one or more of the *only-lists* are accessible.
- 22 An accessible entity in the referenced module has one or more local identifiers. These identifiers are
- 23 (1) the identifier of the entity in the referenced module if that identifier appears as an *only-use-*
24 *name* or as the *defined-operator* of a *generic-spec* in any *only* for that module,
25 (2) each of the *local-names* or *local-defined-operators* that the entity is given in any *rename* for
26 that module, and
27 (3) the identifier of the entity in the referenced module if that identifier does not appear as a
28 *use-name* or *use-defined-operator* in any *rename* for that module.
- 29 Two or more accessible entities, other than generic interfaces or defined operators, may have the same
30 identifier only if the identifier is not used to refer to an entity in the scoping unit. Generic interfaces
31 and defined operators are handled as described in 12.4.3.3 and 12.4.3.3.4. Except for these cases, the
32 local identifier of any entity given accessibility by a USE statement shall differ from the local identifiers
33 of all other entities accessible to the scoping unit through USE statements and otherwise.

NOTE 11.10

There is no prohibition against a *use-name* or *use-defined-operator* appearing multiple times in one USE statement or in multiple USE statements involving the same module. As a result, it is possible for one use-associated entity to be accessible by more than one local identifier.

- 1 The local identifier of an entity made accessible by a USE statement shall not appear in any other
 2 nonexecutable statement that would cause any attribute (5.3) of the entity to be specified in the scoping
 3 unit that contains the USE statement, except that it may appear in a PUBLIC or PRIVATE statement
 4 in the scoping unit of a module and it may be given the ASYNCHRONOUS or VOLATILE attribute.
- 5 The appearance of such a local identifier in a PUBLIC statement in a module causes the entity accessible
 6 by the USE statement to be a public entity of that module. If the identifier appears in a PRIVATE
 7 statement in a module, the entity is not a public entity of that module. If the local identifier does not
 8 appear in either a PUBLIC or PRIVATE statement, it assumes the default accessibility attribute (5.4.1)
 9 of that scoping unit.

NOTE 11.11

The constraints in subclauses 5.7.1, 5.7.2, and 5.6 prohibit the *local-name* from appearing as a *common-block-object* in a COMMON statement, an *equivalence-object* in an EQUIVALENCE statement, or a *namelist-group-name* in a NAMELIST statement, respectively. There is no prohibition against the *local-name* appearing as a *common-block-name* or a *namelist-group-object*.

NOTE 11.12

For a discussion of the impact of the ONLY option and renaming on dependent compilation, see subclause C.8.2.1.

NOTE 11.13

Examples:

```
USE STATS_LIB
```

provides access to all public entities in the module STATS_LIB.

```
USE MATH_LIB; USE STATS_LIB, SPROD => PROD
```

makes all public entities in both MATH_LIB and STATS_LIB accessible. If MATH_LIB contains an entity called PROD, it is accessible by its own name while the entity PROD of STATS_LIB is accessible by the name SPROD.

```
USE STATS_LIB, ONLY: YPROD; USE STATS_LIB, ONLY : PROD
```

makes public entities YPROD and PROD in STATS_LIB accessible.

```
USE STATS_LIB, ONLY : YPROD; USE STATS_LIB
```

makes all public entities in STATS_LIB accessible.

10 11.2.3 Submodules

- 11 A **submodule** is a program unit that extends a module or another submodule. The program unit that
 12 it extends is its **parent**, and is specified by the *parent-identifier* in the *submodule-stmt*. A submodule
 13 is a **child** of its parent. An **ancestor** of a submodule is its parent or an ancestor of its parent. A
 14 **descendant** of a module or submodule is one of its children or a descendant of one of its children. The
 15 **submodule identifier** is the ordered pair whose first element is the ancestor module name and whose
 16 second element is the submodule name.

NOTE 11.14

A module and its submodules stand in a tree-like relationship one to another, with the module at the root. Therefore, a submodule has exactly one ancestor module and may optionally have one or more ancestor submodules.
--

1 A submodule accesses the scoping unit of its parent by host association (16.5.1.4).

2 A submodule may provide implementations for module procedures, each of which is declared by a module
3 procedure interface body (12.4.3.2) within that submodule or one of its ancestors, and declarations and
4 definitions of other entities that are accessible by host association in its descendants.

```
5 R1116  submodule                        is  submodule-stmt
6                                     [ specification-part ]
7                                     [ module-subprogram-part ]
8                                     end-submodule-stmt
```

```
9 R1117  submodule-stmt                    is  SUBMODULE ( parent-identifier ) submodule-name
```

```
10 R1118  parent-identifier                 is  ancestor-module-name [ : parent-submodule-name ]
```

```
11 R1119  end-submodule-stmt               is  END [ SUBMODULE [ submodule-name ] ]
```

12 C1115 (R1116) A submodule *specification-part* shall not contain a *format-stmt*, *entry-stmt*, or *stmt-*
13 *function-stmt*.

14 C1116 (R1116) An object with a default-initialized direct component that is declared in the *specification-*
15 *part* of a submodule shall have the ALLOCATABLE, POINTER, or SAVE attribute.

16 C1117 (R1118) The *ancestor-module-name* shall be the name of a nonintrinsic module; the *parent-*
17 *submodule-name* shall be the name of a descendant of that module.

18 C1118 (R1116) If a *submodule-name* appears in the *end-submodule-stmt*, it shall be identical to the one
19 in the *submodule-stmt*.

20 11.3 Block data program units

21 A **block data program unit** is used to provide initial values for data objects in named common blocks.

```
22 R1120  block-data                        is  block-data-stmt
23                                     [ specification-part ]
24                                     end-block-data-stmt
```

```
25 R1121  block-data-stmt                   is  BLOCK DATA [ block-data-name ]
```

```
26 R1122  end-block-data-stmt               is  END [ BLOCK DATA [ block-data-name ] ]
```

27 C1119 (R1120) The *block-data-name* shall be included in the *end-block-data-stmt* only if it was provided
28 in the *block-data-stmt* and, if included, shall be identical to the *block-data-name* in the *block-*
29 *data-stmt*.

30 C1120 (R1120) A *block-data specification-part* shall contain only derived-type definitions and ASYN-
31 CHRONOUS, BIND, COMMON, DATA, DIMENSION, EQUIVALENCE, IMPLICIT, INTRIN-
32 SIC, PARAMETER, POINTER, SAVE, TARGET, USE, VOLATILE, and type declaration
33 statements.

34 C1121 (R1120) A type declaration statement in a *block-data specification-part* shall not contain AL-
35 LOCATABLE, EXTERNAL, or BIND attribute specifiers.

NOTE 11.15

For explanatory information about the uses for the *block-data-name*, see subclause C.8.1.

- 1 If an object in a named common block is initially defined, all storage units in the common block storage
- 2 sequence shall be specified even if they are not all initially defined. More than one named common block
- 3 may have objects initially defined in a single block data program unit.

NOTE 11.16

In the example

```
BLOCK DATA INIT
  REAL A, B, C, D, E, F
  COMMON /BLOCK1/ A, B, C, D
  DATA A /1.2/, C /2.3/
  COMMON /BLOCK2/ E, F
  DATA F /6.5/
END BLOCK DATA INIT
```

common blocks BLOCK1 and BLOCK2 both have objects that are being initialized in a single block data program unit. B, D, and E are not initialized but they need to be specified as part of the common blocks.

- 4 Only an object in a named common block may be initially defined in a block data program unit.

NOTE 11.17

Objects associated with an object in a common block are considered to be in that common block.

- 5 The same named common block shall not be specified in more than one block data program unit in a
- 6 program.
- 7 There shall not be more than one unnamed block data program unit in a program.

NOTE 11.18

An example of a block data program unit is:

```
BLOCK DATA WORK
  COMMON /WRKCOM/ A, B, C (10, 10)
  REAL :: A, B, C
  DATA A /1.0/, B /2.0/, C /100 * 0.0/
END BLOCK DATA WORK
```


1 12 Procedures

2 12.1 Concepts

3 The concept of a procedure was introduced in 2.2.4. This clause contains a complete description of
4 procedures. The actions specified by a procedure are performed when the procedure is invoked by
5 execution of a reference to it.

6 The sequence of actions encapsulated by a procedure has access to entities in the invoking scoping
7 unit by way of argument association (12.5.2). A **dummy argument** is a name that appears in the
8 SUBROUTINE, FUNCTION, or ENTRY statement in the declaration of a procedure (R1235). Dummy
9 arguments are also specified for intrinsic procedures and procedures in intrinsic modules in Clauses 13,
10 14, and 15.

11 The entities in the invoking scoping unit are specified by actual arguments. An **actual argument** is an
12 entity that appears in a procedure reference (R1223).

13 A procedure may also have access to entities in other scoping units, not necessarily the invoking scoping
14 unit, by use association (16.5.1.3), host association (16.5.1.4), linkage association (16.5.1.5), storage
15 association (16.5.3), or by reference to external procedures (5.3.8).

16 12.2 Procedure classifications

17 12.2.1 Procedure classification by reference

18 The definition of a procedure specifies it to be a function or a subroutine. A reference to a function
19 either appears explicitly as a primary within an expression, or is implied by a defined operation (7.1.3)
20 within an expression. A reference to a subroutine is a CALL statement, a defined assignment statement
21 (7.4.1.4), the appearance of an object processed by user-defined derived-type input/output (9.5.4.7) in
22 an input/output list, or finalization (4.5.6).

23 A procedure is classified as **elemental** if it is a procedure that may be referenced elementally (12.8).

24 12.2.2 Procedure classification by means of definition

25 12.2.2.1 Intrinsic procedures

26 A procedure that is provided as an inherent part of the processor is an **intrinsic procedure**.

27 12.2.2.2 External, internal, and module procedures

28 An **external procedure** is a procedure that is defined by an external subprogram or by a means other
29 than Fortran.

30 An **internal procedure** is a procedure that is defined by an internal subprogram. Internal subprograms
31 may appear in the main program, in an external subprogram, or in a module subprogram. Internal
32 subprograms shall not appear in other internal subprograms. Internal subprograms are the same as
33 external subprograms except that the name of the internal procedure is not a global identifier, an
34 internal subprogram shall not contain an ENTRY statement, and the internal subprogram has access to
35 host entities by host association.

1 A **module procedure** is a procedure that is defined by a module subprogram.

2 A subprogram defines a procedure for the SUBROUTINE or FUNCTION statement. If the subprogram
3 has one or more ENTRY statements, it also defines a procedure for each of them.

4 **12.2.2.3 Dummy procedures**

5 A dummy argument that is specified to be a procedure or appears in a procedure reference is a **dummy**
6 **procedure**. A dummy procedure with the POINTER attribute is a dummy procedure pointer.

7 **12.2.2.4 Procedure pointers**

8 A procedure pointer is a procedure that has the EXTERNAL and POINTER attributes; it may be
9 pointer associated with an external procedure, a module procedure, an intrinsic procedure, or a dummy
10 procedure that is not a procedure pointer.

11 **12.2.2.5 Statement functions**

12 A function that is defined by a single statement is a **statement function** (12.6.4).

13 **12.3 Characteristics**

14 **12.3.1 Characteristics of procedures**

15 The **characteristics of a procedure** are the classification of the procedure as a function or subroutine,
16 whether it is pure, whether it is elemental, whether it has the BIND attribute, the characteristics of its
17 dummy arguments, and the characteristics of its result value if it is a function.

18 **12.3.2 Characteristics of dummy arguments**

19 **12.3.2.1 General**

20 Each dummy argument has the characteristic that it is a dummy data object, a dummy procedure, a
21 dummy procedure pointer, or an asterisk (alternate return indicator). A dummy argument other than an asterisk
22 may be specified to have the OPTIONAL attribute. This attribute means that the dummy argument
23 need not be associated with an actual argument for any particular reference to the procedure.

24 **12.3.2.2 Characteristics of dummy data objects**

25 The characteristics of a dummy data object are its type, its type parameters (if any), its shape, its intent
26 (5.3.9, 5.4.8), whether it is optional (5.3.11, 5.4.9), whether it is allocatable (5.3.7.4), whether it has the
27 ASYNCHRONOUS (5.3.4), CONTIGUOUS (5.3.6), VALUE (5.3.17), or VOLATILE (5.3.18) attributes,
28 whether it is polymorphic, and whether it is a pointer (5.3.13, 5.4.11) or a target (5.3.16, 5.4.14). If
29 a type parameter of an object or a bound of an array is not an initialization expression, the exact
30 dependence on the entities in the expression is a characteristic. If a shape, size, or type parameter is
31 assumed or deferred, it is a characteristic.

32 **12.3.2.3 Characteristics of dummy procedures and dummy procedure pointers**

33 The characteristics of a dummy procedure are the explicitness of its interface (12.4.2), its characteristics
34 as a procedure if the interface is explicit, whether it is a pointer, and whether it is optional (5.3.11, 5.4.9).

35 **12.3.2.4 Characteristics of asterisk dummy arguments**

36 An asterisk as a dummy argument has no characteristics.

1 12.3.3 Characteristics of function results

2 The characteristics of a function result are its type, type parameters (if any), rank, whether it is poly-
3 morphic, whether it is allocatable, whether it is a pointer, whether it has the CONTIGUOUS attribute,
4 and whether it is a procedure pointer. If a function result is an array that is not allocatable or a pointer,
5 its shape is a characteristic. If a type parameter of a function result or a bound of a function result
6 array is not an initialization expression, the exact dependence on the entities in the expression is a
7 characteristic. If type parameters of a function result are deferred, which parameters are deferred is a
8 characteristic. If the length of a character function result is assumed, this is a characteristic.

9 12.4 Procedure interface

10 12.4.1 General

11 The **interface** of a procedure determines the forms of reference through which it may be invoked.
12 The procedure's interface consists of its abstract interface, its name, its binding label if any, and the
13 procedure's generic identifiers, if any. The characteristics of a procedure are fixed, but the remainder
14 of the interface may differ in different scoping units, except that for a separate module procedure body
15 (12.6.2.4), the dummy argument names, binding label, and whether it is recursive shall be the same as
16 in its corresponding module procedure interface body (12.4.3.2).

17 An **abstract interface** consists of procedure characteristics and the names of dummy arguments.

NOTE 12.1

For more explanatory information on procedure interfaces, see C.9.3.

18 12.4.2 Implicit and explicit interfaces

19 If a procedure is accessible in a scoping unit, its interface is either **explicit** or **implicit** in that scoping
20 unit. The interface of an internal procedure, module procedure, or intrinsic procedure is always explicit
21 in such a scoping unit. The interface of a subroutine or a function with a separate result name is explicit
22 within the subprogram that defines it. The interface of a statement function is always implicit. The interface of
23 an external procedure or dummy procedure is explicit in a scoping unit other than its own if an interface
24 body (12.4.3.2) for the procedure is supplied or accessible, and implicit otherwise.

NOTE 12.2

For example, the subroutine LLS of C.8.3.5 has an explicit interface.

25 12.4.2.1 Explicit interface

26 A procedure other than a statement function shall have an explicit interface if it is referenced and

- 27 (1) a reference to the procedure appears
- 28 (a) with an argument keyword (12.5.2),
 - 29 (b) with an argument that is a co-indexed object (2.4.6),

J3 internal note**Unresolved Technical Issue 049**

This makes most old (F77-style) procedures useless for operating on co-indexed objects. I'm not convinced it is necessary.

I was told: "Surely it makes sense to prevent programming errors when possible."

However, this achieves it not.

Forcing the user to lie about his INTENT, use error-suppressing compiler switches, or forego the use of a third-party library which has unspecified intent arguments (in other words, the vast majority of software) does not prevent programming errors, it encourages them.

- 1 (c) in a context that requires it to be pure,
 2 (2) the procedure has a dummy argument that
 3 (a) has the ALLOCATABLE, ASYNCHRONOUS, OPTIONAL, POINTER, TARGET,
 4 VALUE, or VOLATILE attribute,
 5 (b) is an assumed-shape array,
 6 (c) is a co-array,
 7 (d) is of a parameterized derived type, or
 8 (e) is polymorphic,
 9 (3) the procedure has a result that
 10 (a) is an array,
 11 (b) is a pointer or is allocatable, or
 12 (c) has a nonassumed type parameter value that is not an initialization expression,
 13 (4) the procedure is elemental, or
 14 (5) the procedure has the BIND attribute.

12.4.3 Specification of the procedure interface**12.4.3.1 General**

17 The interface for an internal, external, module, or dummy procedure is specified by a FUNCTION,
 18 SUBROUTINE, or ENTRY statement and by specification statements for the dummy arguments and
 19 the result of a function. These statements may appear in the procedure definition, in an interface body,
 20 or both, except that the ENTRY statement shall not appear in an interface body.

NOTE 12.3

An interface body cannot be used to describe the interface of an internal procedure, a module procedure that is not a separate module procedure, or an intrinsic procedure because the interfaces of such procedures are already explicit. However, the name of a procedure may appear in a PROCEDURE statement in an interface block (12.4.3.2).

12.4.3.2 Interface block

- 22 R1201 *interface-block* **is** *interface-stmt*
 23 [*interface-specification*] ...
 24 *end-interface-stmt*
 25 R1202 *interface-specification* **is** *interface-body*
 26 **or** *procedure-stmt*
 27 R1203 *interface-stmt* **is** INTERFACE [*generic-spec*]
 28 **or** ABSTRACT INTERFACE
 29 R1204 *end-interface-stmt* **is** END INTERFACE [*generic-spec*]

- 1 R1205 *interface-body* is *function-stmt*
2 [*specification-part*]
3 *end-function-stmt*
4 or *subroutine-stmt*
5 [*specification-part*]
6 *end-subroutine-stmt*
- 7 R1206 *procedure-stmt* is [MODULE] PROCEDURE *procedure-name-list*
- 8 R1207 *generic-spec* is *generic-name*
9 or OPERATOR (*defined-operator*)
10 or ASSIGNMENT (=)
11 or *dtio-generic-spec*
- 12 R1208 *dtio-generic-spec* is READ (FORMATTED)
13 or READ (UNFORMATTED)
14 or WRITE (FORMATTED)
15 or WRITE (UNFORMATTED)
- 16 R1209 *import-stmt* is IMPORT [[::] *import-name-list*]
- 17 C1201 (R1201) An *interface-block* in a subprogram shall not contain an *interface-body* for a procedure
18 defined by that subprogram.
- 19 C1202 (R1201) The *generic-spec* shall be included in the *end-interface-stmt* only if it is provided in the
20 *interface-stmt*. If the *end-interface-stmt* includes *generic-name*, the *interface-stmt* shall specify
21 the same *generic-name*. If the *end-interface-stmt* includes ASSIGNMENT(=), the *interface-*
22 *stmt* shall specify ASSIGNMENT(=). If the *end-interface-stmt* includes *dtio-generic-spec*,
23 the *interface-stmt* shall specify the same *dtio-generic-spec*. If the *end-interface-stmt* includes
24 OPERATOR(*defined-operator*), the *interface-stmt* shall specify the same *defined-operator*. If
25 one *defined-operator* is .LT., .LE., .GT., .GE., .EQ., or .NE., the other is permitted to be the
26 corresponding operator <, <=, >, >=, ==, or /=.
- 27 C1203 (R1203) If the *interface-stmt* is ABSTRACT INTERFACE, then the *function-name* in the
28 *function-stmt* or the *subroutine-name* in the *subroutine-stmt* shall not be the same as a keyword
29 that specifies an intrinsic type.
- 30 C1204 (R1202) A *procedure-stmt* is allowed only in an interface block that has a *generic-spec*.
- 31 C1205 (R1205) An *interface-body* of a pure procedure shall specify the intents of all dummy arguments
32 except pointer, alternate return, and procedure arguments.
- 33 C1206 (R1205) An *interface-body* shall not contain an *entry-stmt*, *data-stmt*, *format-stmt*, or *stmt-*
34 *function-stmt*.
- 35 C1207 (R1206) A *procedure-name* shall have an explicit interface and shall refer to an accessible pro-
36 cedure pointer, external procedure, dummy procedure, or module procedure.
- 37 C1208 (R1206) If MODULE appears in a *procedure-stmt*, each *procedure-name* in that statement shall
38 be accessible in the current scope as a module procedure.
- 39 C1209 (R1206) A *procedure-name* shall not specify a procedure that is specified previously in any
40 *procedure-stmt* in any accessible interface with the same generic identifier.
- 41 C1210 (R1209) The IMPORT statement is allowed only in an *interface-body* that is not a module
42 procedure interface body.
- 43 C1211 (R1209) Each *import-name* shall be the name of an entity in the host scoping unit.
- 44 An external or module subprogram specifies a **specific interface** for the procedures defined in that
45 subprogram. Such a specific interface is explicit for module procedures and implicit for external proce-

1 dures.

2 An interface block introduced by ABSTRACT INTERFACE is an **abstract interface block**. An
3 interface body in an abstract interface block specifies an abstract interface. An interface block with a
4 generic specification is a **generic interface block**. An interface block with neither ABSTRACT nor a
5 generic specification is a **specific interface block**.

6 The name of the entity declared by an interface body is the *function-name* in the *function-stmt* or the
7 *subroutine-name* in the *subroutine-stmt* that begins the interface body.

8 A **module procedure interface body** is an interface body whose initial statement contains the
9 keyword MODULE. It defines the **module procedure interface** for a separate module procedure
10 (12.6.2.4). A separate module procedure is accessible by use association if and only if its interface body
11 is declared in the specification part of a module and is public. If a corresponding (12.6.2.4) separate
12 module procedure is not defined, the interface may be used to specify an explicit specific interface but
13 the procedure shall not be used in any other way.

14 C1212 (R1205) A module procedure interface body shall not appear in an abstract interface block.

15 An interface body in a generic or specific interface block specifies the EXTERNAL attribute and an
16 explicit specific interface for an external procedure or a dummy procedure. If the name of the declared
17 procedure is that of a dummy argument in the subprogram containing the interface body, the procedure
18 is a dummy procedure; otherwise, it is an external procedure.

19 An interface body specifies all of the characteristics of the explicit specific interface or abstract interface.
20 The specification part of an interface body may specify attributes or define values for data entities that
21 do not determine characteristics of the procedure. Such specifications have no effect.

22 If an explicit specific interface is specified by an interface body or a procedure declaration statement
23 (12.4.3.5) for an external procedure, the characteristics shall be consistent with those specified in the
24 procedure definition, except that the interface may specify a procedure that is not pure if the procedure
25 is defined to be pure. An interface for a procedure named by an ENTRY statement may be specified by
26 using the entry name as the procedure name in the interface body. If an external procedure does not
27 exist in the program, an interface body for it may be used to specify an explicit specific interface but
28 the procedure shall not be used in any other way. An explicit specific interface may be specified by an
29 interface body for an external procedure that does not exist in the program if the procedure is never
30 used in any way. A procedure shall not have more than one explicit specific interface in a given scoping
31 unit.

NOTE 12.4

The dummy argument names in an interface body may be different from the corresponding dummy argument names in the procedure definition because the name of a dummy argument is not a characteristic.

32 The IMPORT statement specifies that the named entities from the host scoping unit are accessible in
33 the interface body by host association. An entity that is imported in this manner and is defined in the
34 host scoping unit shall be explicitly declared prior to the interface body. The name of an entity made
35 accessible by an IMPORT statement shall not appear in any of the contexts described in 16.5.1.4 that
36 cause the host entity of that name to be inaccessible.

37 Within an interface body, if an IMPORT statement with no *import-name-list* appears, each host entity
38 not named in an IMPORT statement also is made accessible by host association if its name does not
39 appear in any of the contexts described in 16.5.1.4 that cause the host entity of that name to be
40 inaccessible. If an entity that is made accessible by this means is accessed by host association and is
41 defined in the host scoping unit, it shall be explicitly declared prior to the interface body.

NOTE 12.5

An example of an interface block without a generic specification is:

```

INTERFACE
  SUBROUTINE EXT1 (X, Y, Z)
    REAL, DIMENSION (100, 100) :: X, Y, Z
  END SUBROUTINE EXT1
  SUBROUTINE EXT2 (X, Z)
    REAL X
    COMPLEX (KIND = 4) Z (2000)
  END SUBROUTINE EXT2
  FUNCTION EXT3 (P, Q)
    LOGICAL EXT3
    INTEGER P (1000)
    LOGICAL Q (1000)
  END FUNCTION EXT3
END INTERFACE

```

This interface block specifies explicit interfaces for the three external procedures EXT1, EXT2, and EXT3. Invocations of these procedures may use argument keywords (12.5.2); for example:

```
EXT3 (Q = P_MASK (N+1 : N+1000), P = ACTUAL_P)
```

NOTE 12.6

The IMPORT statement can be used to allow module procedures to have dummy arguments that are procedures with assumed-shape arguments of an opaque type. For example:

```

MODULE M
  TYPE T
    PRIVATE    ! T is an opaque type
    ...
  END TYPE
CONTAINS
  SUBROUTINE PROCESS(X,Y,RESULT,MONITOR)
    TYPE(T),INTENT(IN) :: X(:,,:),Y(:,,:)
    TYPE(T),INTENT(OUT) :: RESULT(:,,:)
    INTERFACE
      SUBROUTINE MONITOR(ITERATION_NUMBER,CURRENT_ESTIMATE)
        IMPORT T
        INTEGER,INTENT(IN) :: ITERATION_NUMBER
        TYPE(T),INTENT(IN) :: CURRENT_ESTIMATE(:,,:)
      END SUBROUTINE
    END INTERFACE
    ...
  END SUBROUTINE
END MODULE

```

The MONITOR dummy procedure requires an explicit interface because it has an assumed-shape array argument, but TYPE(T) would not be available inside the interface body without the IMPORT statement.

- 1 A generic interface block specifies a **generic interface** for each of the procedures in the interface
 2 block. The PROCEDURE statement lists procedure pointers, external procedures, dummy procedures,
 3 or module procedures that have this generic interface. A generic interface is always explicit.
- 4 Any procedure may be referenced via its specific interface if the specific interface is accessible. It also
 5 may be referenced via a generic interface. The *generic-spec* in an *interface-stmt* is a **generic identifier**
 6 for all the procedures in the interface block. The rules specifying how any two procedures with the same
 7 generic identifier shall differ are given in 12.4.3.3.4. They ensure that any generic invocation applies to
 8 at most one specific procedure.
- 9 A **generic name** specifies a single name to reference all of the procedure names in the interface block.
 10 A generic name may be the same as any one of the procedure names in the interface block, or the same
 11 as any accessible generic name.
- 12 A generic name may be the same as a derived-type name, in which case all of the procedures in the
 13 interface block shall be functions.

NOTE 12.7

An example of a generic procedure interface is:

```
INTERFACE SWITCH
  SUBROUTINE INT_SWITCH (X, Y)
    INTEGER, INTENT (INOUT) :: X, Y
  END SUBROUTINE INT_SWITCH
  SUBROUTINE REAL_SWITCH (X, Y)
    REAL, INTENT (INOUT) :: X, Y
  END SUBROUTINE REAL_SWITCH
  SUBROUTINE COMPLEX_SWITCH (X, Y)
    COMPLEX, INTENT (INOUT) :: X, Y
  END SUBROUTINE COMPLEX_SWITCH
END INTERFACE SWITCH
```

Any of these three subroutines (INT_SWITCH, REAL_SWITCH, COMPLEX_SWITCH) may be referenced with the generic name SWITCH, as well as by its specific name. For example, a reference to INT_SWITCH could take the form:

```
CALL SWITCH (MAX_VAL, LOC_VAL) ! MAX_VAL and LOC_VAL are of type INTEGER
```

- 14 An *interface-stmt* having a *dtio-generic-spec* is an interface for a user-defined derived-type input/output
 15 procedure (9.5.4.7)

16 12.4.3.3.1 Defined operations

- 17 If OPERATOR is specified in a generic specification, all of the procedures specified in the generic interface
 18 shall be functions that may be referenced as defined operations (7.1.3, 7.1.8.9, 7.2, 12.5). In the case of
 19 functions of two arguments, infix binary operator notation is implied. In the case of functions of one
 20 argument, prefix operator notation is implied. OPERATOR shall not be specified for functions with no
 21 arguments or for functions with more than two arguments. The dummy arguments shall be nonoptional
 22 dummy data objects and shall be specified with INTENT (IN). The function result shall not have assumed
 23 character length. If the operator is an *intrinsic-operator* (R310), the number of function arguments shall
 24 be consistent with the intrinsic uses of that operator, and the types, kind type parameters, or ranks of
 25 the dummy arguments shall differ from those required for the intrinsic operation (7.1.2).

- 26 A defined operation is treated as a reference to the function. For a unary defined operation, the operand

- 1 corresponds to the function's dummy argument; for a binary operation, the left-hand operand corre-
 2 sponds to the first dummy argument of the function and the right-hand operand corresponds to the
 3 second dummy argument.

NOTE 12.8

An example of the use of the OPERATOR generic specification is:

```
INTERFACE OPERATOR ( * )
  FUNCTION BOOLEAN_AND (B1, B2)
    LOGICAL, INTENT (IN) :: B1 (:), B2 (SIZE (B1))
    LOGICAL :: BOOLEAN_AND (SIZE (B1))
  END FUNCTION BOOLEAN_AND
END INTERFACE OPERATOR ( * )
```

This allows, for example

```
SENSOR (1:N) * ACTION (1:N)
```

as an alternative to the function call

```
BOOLEAN_AND (SENSOR (1:N), ACTION (1:N))    ! SENSOR and ACTION are
                                              ! of type LOGICAL
```

- 4 A given defined operator may, as with generic names, apply to more than one function, in which case
 5 it is generic in exact analogy to generic procedure names. For intrinsic operator symbols, the generic
 6 properties include the intrinsic operations they represent. Because both forms of each relational operator
 7 have the same interpretation (7.2), extending one form (such as <=) has the effect of defining both forms
 8 (<= and .LE.).

NOTE 12.9

In Fortran 90 and Fortran 95, it was not possible to define operations on pointers because pointer dummy arguments were disallowed from having an INTENT attribute. The restriction against INTENT for pointer dummy arguments is now lifted, so defined operations on pointers are now possible.

However, the POINTER attribute cannot be used to resolve generic procedures (12.4.3.3.4), so it is not possible to define a generic operator that has one procedure for pointers and another procedure for nonpointers.

9 12.4.3.3.2 Defined assignments

- 10 If ASSIGNMENT (=) is specified in a generic specification, all the procedures in the generic interface
 11 shall be subroutines that may be referenced as defined assignments (7.4.1.4). Defined assignment may,
 12 as with generic names, apply to more than one subroutine, in which case it is generic in exact analogy
 13 to generic procedure names.

- 14 Each of these subroutines shall have exactly two dummy arguments. The dummy arguments shall be
 15 nonoptional dummy data objects. The first argument shall have INTENT (OUT) or INTENT (INOUT)
 16 and the second argument shall have INTENT (IN). Either the second argument shall be an array whose
 17 rank differs from that of the first argument, the declared types and kind type parameters of the arguments
 18 shall not conform as specified in Table 7.11, or the first argument shall be of derived type. A defined
 19 assignment is treated as a reference to the subroutine, with the left-hand side as the first argument
 20 and the right-hand side enclosed in parentheses as the second argument. The ASSIGNMENT generic

- 1 specification specifies that assignment is extended or redefined.

NOTE 12.10

An example of the use of the ASSIGNMENT generic specification is:

```
INTERFACE ASSIGNMENT ( = )

  SUBROUTINE LOGICAL_TO_NUMERIC (N, B)
    INTEGER, INTENT (OUT) :: N
    LOGICAL, INTENT (IN)  :: B
  END SUBROUTINE LOGICAL_TO_NUMERIC
  SUBROUTINE CHAR_TO_STRING (S, C)
    USE STRING_MODULE      ! Contains definition of type STRING
    TYPE (STRING), INTENT (OUT) :: S ! A variable-length string
    CHARACTER (*), INTENT (IN)  :: C
  END SUBROUTINE CHAR_TO_STRING
END INTERFACE ASSIGNMENT ( = )
```

Example assignments are:

```
KOUNT = SENSOR (J) ! CALL LOGICAL_TO_NUMERIC (KOUNT, (SENSOR (J)))
NOTE  = '89AB'     ! CALL CHAR_TO_STRING (NOTE, ('89AB'))
```

2 **12.4.3.3.3 User-defined derived-type input/output procedure interfaces**

- 3 All of the procedures specified in an interface block for a user-defined derived-type input/output procedure shall be subroutines that have interfaces as described in 9.5.4.7.2.

5 **12.4.3.3.4 Restrictions on generic declarations**

- 6 This subclause contains the rules that shall be satisfied by every pair of specific procedures that have
7 the same generic identifier within a scoping unit. If a generic procedure is accessed from a module, the
8 rules apply to all the specific versions even if some of them are inaccessible by their specific names.

NOTE 12.11

In most scoping units, the possible sources of procedures with a particular generic identifier are the accessible interface blocks and the generic bindings other than names for the accessible objects in that scoping unit. In a type definition, they are the generic bindings, including those from a parent type.

- 9 A dummy argument is type, kind, and rank compatible, or **TKR compatible**, with another dummy
10 argument if both have the same rank, and either the first is type compatible with the second and the
11 kind type parameters of the first have the same values as the corresponding kind type parameters of the
12 second, or the two are bits compatible.

- 13 Two dummy arguments are **distinguishable** if

- 14 • one is a procedure and the other is a data object,
- 15 • they are both data objects or known to be functions, and neither is TKR compatible with the
16 other,
- 17 • one has the ALLOCATABLE attribute and the other has the POINTER attribute, or
- 18 • one is a function with nonzero rank and the other is not known to be a function.

1 Within a scoping unit, if two procedures have the same generic operator and the same number of
 2 arguments or both define assignment, one shall have a dummy argument that corresponds by position
 3 in the argument list to a dummy argument of the other that is distinguishable with it.

4 Within a scoping unit, if two procedures have the same *dtio-generic-spec* (12.4.3.2), their *dtv* arguments
 5 shall be type incompatible or have different kind type parameters.

6 Within a scoping unit, two procedures that have the same generic name shall both be subroutines or
 7 both be functions, and

- 8 (1) there is a non-passed-object dummy data object in one or the other of them such that
 - 9 (a) the number of dummy data objects in one that are nonoptional, are not passed-object,
 10 and with which that dummy data object is TKR compatible, possibly including that
 11 dummy data object itself,
 12 exceeds
 - 13 (b) the number of non-passed-object dummy data objects, both optional and nonoptional,
 14 in the other that are not distinguishable with that dummy data object,
- 15 (2) both have passed-object dummy arguments and the passed-object dummy arguments are
 16 distinguishable, or
- 17 (3) at least one of them shall have both
 - 18 (a) a nonoptional non-passed-object dummy argument at an effective position such that
 19 either the other procedure has no dummy argument at that effective position or the
 20 dummy argument at that position is distinguishable with it, and
 - 21 (b) a nonoptional non-passed-object dummy argument whose name is such that either the
 22 other procedure has no dummy argument with that name or the dummy argument
 23 with that name is distinguishable with it.

24 and the dummy argument that disambiguates by position shall either be the same as or
 25 occur earlier in the argument list than the one that disambiguates by name.

26 The **effective position** of a dummy argument is its position in the argument list after any passed-object
 27 dummy argument has been removed.

28 Within a scoping unit, if a generic name is the same as the name of a generic intrinsic procedure, the
 29 generic intrinsic procedure is not accessible if the procedures in the interface and the intrinsic procedure
 30 are not all functions or not all subroutines. If a generic invocation applies to both a specific procedure
 31 from an interface and an accessible generic intrinsic procedure, it is the specific procedure from the
 32 interface that is referenced.

NOTE 12.12

An extensive explanation of the application of these rules is in C.12.2.

33 12.4.3.4 EXTERNAL statement

34 An **EXTERNAL statement** specifies the EXTERNAL attribute (5.3.8) for a list of names.

35 R1210 *external-stmt* **is** EXTERNAL [::] *external-name-list*

36 The appearance of the name of a block data program unit in an EXTERNAL statement confirms that
 37 the block data program unit is a part of the program.

NOTE 12.13

For explanatory information on potential portability problems with external procedures, see sub-
 clause C.9.1.

NOTE 12.14

An example of an EXTERNAL statement is:

```
EXTERNAL FOCUS
```

1 12.4.3.5 Procedure declaration statement

2 A procedure declaration statement declares procedure pointers, dummy procedures, and external pro-
3 cedures. It specifies the EXTERNAL attribute (5.3.8) for all procedure entities in the *proc-decl-list*.

```
4 R1211 procedure-declaration-stmt is PROCEDURE ( [ proc-interface ] ) ■
5                                     ■ [ [ , proc-attr-spec ] ... :: ] proc-decl-list
6 R1212 proc-interface             is interface-name
7                                     or declaration-type-spec
8 R1213 proc-attr-spec             is access-spec
9                                     or proc-language-binding-spec
10                                    or INTENT ( intent-spec )
11                                    or OPTIONAL
12                                    or POINTER
13                                    or SAVE
14 R1214 proc-decl                 is procedure-entity-name[ => proc-pointer-init ]
15 R1215 interface-name            is name
16 R1216 proc-pointer-init         is null-init
17                                     or initial-proc-target
18 R1217 initial-proc-target        is procedure-name
```

19

20 C1213 (R1215) The *name* shall be the name of an abstract interface or of a procedure that has an
21 explicit interface. If *name* is declared by a *procedure-declaration-stmt* it shall be previously
22 declared. If *name* denotes an intrinsic procedure it shall be one that is listed in 13.6 and not
23 marked with a bullet (●).

24 C1214 (R1215) The *name* shall not be the same as a keyword that specifies an intrinsic type.

25 C1215 If a procedure entity has the INTENT attribute or SAVE attribute, it shall also have the
26 POINTER attribute.

27 C1216 (R1211) If a *proc-interface* describes an elemental procedure, each *procedure-entity-name* shall
28 specify an external procedure.

29 C1217 (R1214) If => appears in *proc-decl*, the procedure entity shall have the POINTER attribute.

30 C1218 (R1217) The *procedure-name* shall be the name of an initialization target.

31 C1219 (R1211) If *proc-language-binding-spec* with a NAME= is specified, then *proc-decl-list* shall con-
32 tain exactly one *proc-decl*, which shall neither have the POINTER attribute nor be a dummy
33 procedure.

34 C1220 (R1211) If *proc-language-binding-spec* is specified, the *proc-interface* shall appear, it shall be an
35 *interface-name*, and *interface-name* shall be declared with a *proc-language-binding-spec*.

36 If *proc-interface* appears and consists of *interface-name*, it specifies an explicit specific interface (12.4.3.2)
37 for the declared procedures or procedure pointers. The abstract interface (12.4) is that specified by the
38 interface named by *interface-name*.

39 If *proc-interface* appears and consists of *declaration-type-spec*, it specifies that the declared procedures

- 1 or procedure pointers are functions having implicit interfaces and the specified result type. If a type is
 2 specified for an external function, its function definition (12.6.2.1) shall specify the same result type and
 3 type parameters.
- 4 If *proc-interface* does not appear, the procedure declaration statement does not specify whether the
 5 declared procedures or procedure pointers are subroutines or functions.
- 6 If a *proc-attr-spec* other than a *proc-language-binding-spec* appears, it specifies that the declared proce-
 7 dures or procedure pointers have that attribute. These attributes are described in 5.3. If a *proc-language-*
 8 *binding-spec* with NAME= appears, it specifies a binding label or its absence, as described in 15.5.2.
 9 A *proc-language-binding-spec* without a NAME= is allowed, but is redundant with the *proc-interface*
 10 required by C1220.
- 11 A procedure is an **initialization target** if it is a nonelemental external or module procedure, or a
 12 specific intrinsic function listed in 13.6 and not marked with a bullet (●).
- 13 If => appears in a *proc-decl* in a *procedure-declaration-stmt* it specifies the initial association status
 14 of the corresponding procedure entity, and implies the SAVE attribute, which may be confirmed by
 15 explicit specification. If => *null-init* appears, the procedure entity is initially disassociated. If =>
 16 *initial-proc-target* appears, the procedure entity is initially associated with the target.
- 17 If *proc-entity-name* has an explicit interface, its characteristics shall be the same as *initial-proc-target*
 18 except that *initial-proc-target* may be pure even if *proc-entity-name* is not pure and *initial-proc-target*
 19 may be an elemental intrinsic procedure.
- 20 If the characteristics of *proc-entity-name* or *initial-proc-target* are such that an explicit interface is
 21 required, both *proc-entity-name* and *initial-proc-target* shall have an explicit interface.
- 22 If *proc-entity-name* has an implicit interface and is explicitly typed or referenced as a function, *initial-*
 23 *proc-target* shall be a function. If *proc-entity-name* has an implicit interface and is referenced as a
 24 subroutine, *initial-proc-target* shall be a subroutine.
- 25 If *initial-proc-target* and *proc-entity-name* are functions, they shall have the same type; corresponding
 26 type parameters shall either both be deferred or both have the same value.

NOTE 12.15

In contrast to the EXTERNAL statement, it is not possible to use the PROCEDURE statement to identify a BLOCK DATA subprogram.

NOTE 12.16

The following code illustrates PROCEDURE statements. Note 7.48 illustrates the use of the P and BESSEL defined by this code.

```

ABSTRACT INTERFACE
  FUNCTION REAL_FUNC (X)
    REAL, INTENT (IN) :: X
    REAL :: REAL_FUNC
  END FUNCTION REAL_FUNC
END INTERFACE

INTERFACE
  SUBROUTINE SUB (X)
    REAL, INTENT (IN) :: X
  END SUBROUTINE SUB
END INTERFACE

```

NOTE 12.16 (cont.)

```

!-- Some external or dummy procedures with explicit interface.
PROCEDURE (REAL_FUNC) :: BESSEL, GFUN
PROCEDURE (SUB) :: PRINT_REAL
!-- Some procedure pointers with explicit interface,
!-- one initialized to NULL().
PROCEDURE (REAL_FUNC), POINTER :: P, R => NULL()
PROCEDURE (REAL_FUNC), POINTER :: PTR_TO_GFUN
!-- A derived type with a procedure pointer component ...
TYPE STRUCT_TYPE
  PROCEDURE (REAL_FUNC), POINTER :: COMPONENT
END TYPE STRUCT_TYPE
!-- ... and a variable of that type.
TYPE(STRUCT_TYPE) :: STRUCT
!-- An external or dummy function with implicit interface
PROCEDURE (REAL) :: PSI

```

1 12.4.3.6 INTRINSIC statement

2 An **INTRINSIC statement** specifies a list of names that have the INTRINSIC attribute (5.3.10).

3 R1218 *intrinsic-stmt* **is** INTRINSIC [::] *intrinsic-procedure-name-list*

4 C1221 (R1218) Each *intrinsic-procedure-name* shall be the name of an intrinsic procedure.

NOTE 12.17

A name shall not be explicitly specified to have both the EXTERNAL and INTRINSIC attributes in the same scoping unit.

5 12.4.3.7 Implicit interface specification

6 In a scoping unit where the interface of a function is implicit, the type and type parameters of the
7 function result are specified by an implicit or explicit type specification of the function name. The type,
8 type parameters, and shape of dummy arguments of a procedure invoked from a scoping unit where the
9 interface of the procedure is implicit shall be such that the actual arguments are consistent with the
10 characteristics of the dummy arguments.

11 12.5 Procedure reference**12 12.5.1 Syntax**

13 The form of a procedure reference is dependent on the interface of the procedure or procedure pointer,
14 but is independent of the means by which the procedure is defined. The forms of procedure references
15 are as follows.

16 R1219 *function-reference* **is** *procedure-designator* ([*actual-arg-spec-list*])

17 C1222 (R1219) The *procedure-designator* shall designate a function.

18 C1223 (R1219) The *actual-arg-spec-list* shall not contain an *alt-return-spec*.

19 R1220 *call-stmt* **is** CALL *procedure-designator* [([*actual-arg-spec-list*])]

- 1 C1224 (R1220) The *procedure-designator* shall designate a subroutine.
- 2 R1221 *procedure-designator* **is** *procedure-name*
 3 **or** *proc-component-ref*
 4 **or** *data-ref* % *binding-name*
- 5 C1225 (R1221) A *procedure-name* shall be the name of a procedure or procedure pointer.
- 6 C1226 (R1221) A *binding-name* shall be a binding name (4.5.5) of the declared type of *data-ref*.
- 7 C1227 (R1221) If *data-ref* is an array, the referenced type-bound procedure shall have the PASS at-
 8 tribute.
- 9 Resolving references to type-bound procedures is described in 12.5.6.
- 10 A function may also be referenced as a defined operation (12.4.3.3.1). A subroutine may also be ref-
 11 erenced as a defined assignment (12.4.3.3.2), by user-defined derived-type input/output (9.5.4.7), or by
 12 finalization (4.5.6).

NOTE 12.18

If image *I* executes a procedure reference in which the *variable* of a *proc-component-ref* specifies a procedure pointer on image *J*, the procedure pointer association is fetched from image *J* but the invocation of the associated procedure occurs on image *I*.

- 13 R1222 *actual-arg-spec* **is** [*keyword* =] *actual-arg*
 14 R1223 *actual-arg* **is** *expr*
 15 **or** *variable*
 16 **or** *procedure-name*
 17 **or** *proc-component-ref*
 18 **or** *alt-return-spec*
 19 R1224 *alt-return-spec* **is** * *label*
- 20 C1228 (R1222) The *keyword* = shall not appear if the interface of the procedure is implicit in the
 21 scoping unit.
- 22 C1229 (R1222) The *keyword* = shall not be omitted from an *actual-arg-spec* unless it has been omitted
 23 from each preceding *actual-arg-spec* in the argument list.
- 24 C1230 (R1222) Each *keyword* shall be the name of a dummy argument in the explicit interface of the
 25 procedure.
- 26 C1231 (R1223) A nonintrinsic elemental procedure shall not be used as an actual argument.
- 27 C1232 (R1223) A *procedure-name* shall be the name of an external, internal, module, or dummy proce-
 28 dure, a specific intrinsic function listed in 13.6 and not marked with a bullet (●), or a procedure
 29 pointer.
- 30 C1233 (R1224) The *label* shall be the statement label of a branch target statement that appears in the same scoping
 31 unit as the *call-stmt*.

NOTE 12.19

Successive commas shall not be used to omit optional arguments.

NOTE 12.20

Examples of procedure reference using procedure pointers:

```
P => BESSEL
WRITE (*, *) P(2.5)      !-- BESSEL(2.5)

S => PRINT_REAL
CALL S(3.14)
```

1 12.5.2 Actual arguments, dummy arguments, and argument association

2 12.5.2.1 Argument correspondence

3 In either a subroutine reference or a function reference, the actual argument list identifies the corre-
 4 spondence between the actual arguments supplied and the dummy arguments of the procedure. This
 5 correspondence may be established either by keyword or by position. If an argument keyword appears,
 6 the actual argument corresponds to the dummy argument whose name is the same as the argument
 7 keyword (using the dummy argument names from the interface accessible in the scoping unit containing
 8 the procedure reference). In the absence of an argument keyword, an actual argument corresponds to the
 9 dummy argument occupying the corresponding position in the reduced dummy argument list; that is,
 10 the first actual argument corresponds to the first dummy argument in the reduced list, the second actual
 11 argument corresponds to the second dummy argument in the reduced list, etc. The reduced dummy
 12 argument list is either the full dummy argument list or, if there is a passed-object dummy argument
 13 (4.5.4.4), the dummy argument list with the passed object dummy argument omitted. Exactly one
 14 actual argument shall correspond to each nonoptional dummy argument. At most one actual argument
 15 shall correspond to each optional dummy argument. Each actual argument shall correspond to a dummy
 16 argument.

NOTE 12.21

For example, the procedure defined by

```
SUBROUTINE SOLVE (FUNCT, SOLUTION, METHOD, STRATEGY, PRINT)
  INTERFACE
    FUNCTION FUNCT (X)
      REAL FUNCT, X
    END FUNCTION FUNCT
  END INTERFACE
  REAL SOLUTION
  INTEGER, OPTIONAL :: METHOD, STRATEGY, PRINT
  ...
```

may be invoked with

```
CALL SOLVE (FUN, SOL, PRINT = 6)
```

provided its interface is explicit; if the interface is specified by an interface block, the name of the last argument shall be PRINT.

17 12.5.2.2 The passed-object dummy argument and argument correspondence

18 In a reference to a type-bound procedure, or a procedure pointer component, that has a passed-object
 19 dummy argument (4.5.4.4), the *data-ref* of the *function-reference* or *call-stmt* corresponds, as an actual
 20 argument, with the passed-object dummy argument.

1 12.5.2.3 Argument association

2 Except in references to intrinsic inquiry functions, if a nonoptional nonpointer dummy argument corre-
3 sponds to a pointer actual argument, the actual argument shall be pointer associated with a target and
4 the dummy argument becomes argument associated with that target. If an optional nonpointer dummy
5 argument corresponds to a pointer actual argument that is pointer associated with a target the dummy
6 argument becomes argument associated with that target. A present nonpointer dummy argument that
7 corresponds to a nonpointer actual argument becomes argument associated with that actual argument.

8 A present pointer dummy argument that corresponds to a pointer actual argument becomes argument
9 associated with that actual argument. A present pointer dummy argument that does not correspond to
10 a pointer actual argument is not argument associated.

11 The entity that is argument-associated with a dummy argument is called its **effective argument**.

12 12.5.2.4 Compatibility with bits objects

13 An entity is **bits compatible** with another entity if and only if one is of type bits, the other is of type
14 bits, integer, real, complex, or logical, and scalar entities of these types have the same size expressed in
15 bits.

J3 internal note

Unresolved Technical Issue 016

Can “bits compatible” be folded into a new concept of “type and kind compatible”?

16 12.5.2.5 Ordinary dummy variables

17 The requirements in this subclause apply to actual arguments that correspond to nonallocatable non-
18 pointer dummy data objects.

19 The dummy argument shall be type compatible with the actual argument or it shall be of type bits and
20 bits compatible with the actual argument.

21 Unless the actual argument and the corresponding dummy argument are bits compatible, the type
22 parameter values of the actual argument shall agree with the corresponding ones of the dummy argument
23 that are not assumed, except for the case of the character length parameter of an actual argument of
24 type default character associated with a dummy argument that is not assumed shape.

25 If a scalar dummy argument is of type default character, the length *len* of the dummy argument shall
26 be less than or equal to the length of the actual argument. The dummy argument becomes associated
27 with the leftmost *len* characters of the actual argument. If an array dummy argument is of type default
28 character and is not assumed shape, it becomes associated with the leftmost characters of the actual
29 argument element sequence (12.5.2.12) and it shall not extend beyond the end of that sequence.

30 The values of assumed type parameters of a dummy argument are assumed from the corresponding type
31 parameters of the actual argument.

32 If the actual argument is a co-indexed object, the dummy argument shall not be a co-array and shall
33 have the INTENT(IN) or the VALUE attribute.

J3 internal note

Unresolved Technical Issue 050

It doesn't seem like forcing the user to write

```
localvar = x[i]
CALL sub(localvar)
x[i] = localvar
```

is a substantial improvement in, well, anything. Not to mention that the user cannot do that if `x[i]` is undefined or only partially defined before the call to `sub`. It doesn't seem to be any more work for the processor than what it already does for

```
CALL sub(a(1:n:2))
```

for explicit-shape dummies. And, *MUCH WORSE*, this makes defined assignment unusable with co-arrays. Can we say "not even integrated with Fortran 90"?

Subgroup said: "this requires more complications to the memory consistency rules since the copy out may overwrite other changes to the variable."

The editor says: That is just so untrue it is not funny.

Is everyone *REALLY* that ignorant of our dummy argument aliasing rules?

We've not repealed them, and there are no edits to subvert them for co-arrays. If they need to be subverted (something I doubt, but maybe the go-faster department want to go slower), something needs to be done.

Either way this is a serious technical flaw.

NOTE 12.22

If the actual argument is a co-indexed object and the corresponding dummy argument is not a co-array, it is likely that a processor will make a copy on the executing image of the actual argument, including copies of any allocated allocatable subcomponents, before argument association occurs.

J3 internal note

Unresolved Technical Issue 051

That doesn't seem likely for the vast majority of Fortran processing systems, which are going to be shared-memory small-cpu-count. They are, presumably, going to do the "right thing" here.

That does not involve making any copies.

Anyway, what's this "and the ... dummy is not a co-array." guff. Last time I looked, a co-indexed object was not a valid actual argument to a co-array dummy.

- 1 Except in references to intrinsic inquiry functions, if the dummy argument is nonoptional and the actual
- 2 argument is allocatable, the corresponding actual argument shall be allocated.
- 3 If the dummy argument has the VALUE attribute it becomes associated with a definable anonymous
- 4 data object whose initial value is that of the actual argument. Subsequent changes to the value or
- 5 definition status of the dummy argument do not affect the actual argument.

NOTE 12.23

Fortran argument association is usually similar to call by reference and call by value-result. If the VALUE attribute is specified, the effect is as if the actual argument is assigned to a temporary, and the temporary is then argument associated with the dummy argument. The actual mechanism by which this happens is determined by the processor.

- 6 If the dummy argument does not have the TARGET attribute, or is not type-compatible with the
- 7 actual argument, any pointers associated with the effective argument do not become associated with the

1 corresponding dummy argument on invocation of the procedure. If such a dummy argument is used as
2 an actual argument that corresponds to a dummy argument with the TARGET attribute, whether any
3 pointers associated with the original effective argument become associated with the dummy argument
4 with the TARGET attribute is processor dependent.

5 If the dummy argument has the TARGET attribute, is type-compatible with the actual argument, does
6 not have the VALUE attribute, and is either a scalar or an assumed-shape array that does not have the
7 CONTIGUOUS attribute, and the effective argument has the TARGET attribute but is not a co-indexed
8 object or an array section with a vector subscript then

- 9 (1) any pointers associated with the effective argument become associated with the correspond-
10 ing dummy argument on invocation of the procedure, and
- 11 (2) when execution of the procedure completes, any pointers that do not become undefined
12 (16.5.2.2.3) and are associated with the dummy argument remain associated with the effec-
13 tive argument.

14 If the dummy argument has the TARGET attribute and is an explicit-shape array, an assumed-shape
15 array with the CONTIGUOUS attribute, or an assumed-size array, and the effective argument has the
16 TARGET attribute but is not an array section with a vector subscript then

- 17 (1) on invocation of the procedure, whether any pointers associated with the effective argument
18 become associated with the corresponding dummy argument is processor dependent, and
- 19 (2) when execution of the procedure completes, the pointer association status of any pointer
20 that is pointer associated with the dummy argument is processor dependent.

21 If the dummy argument has the TARGET attribute and the effective argument does not have the
22 TARGET attribute or is an array section with a vector subscript, or the dummy argument is not
23 type-compatible with the actual argument, any pointers associated with the dummy argument become
24 undefined when execution of the procedure completes.

25 If the dummy argument has the TARGET attribute and the VALUE attribute, any pointers associated
26 with the dummy argument become undefined when execution of the procedure completes.

27 If the actual argument is scalar, the corresponding dummy argument shall be scalar unless the actual
28 argument is of type default character, of type character with the C character kind (15.2), or is an element
29 or substring of an element of an array that is not an assumed-shape, pointer, or polymorphic array. If
30 the procedure is nonelemental and is referenced by a generic name or as a defined operator or defined
31 assignment, the ranks of the actual arguments and corresponding dummy arguments shall agree.

32 If a dummy argument is an assumed-shape array, the rank of the actual argument shall be the same as
33 the rank of the dummy argument; the actual argument shall not be an assumed-size array (including an
34 array element designator or an array element substring designator).

35 Except when a procedure reference is elemental (12.8), each element of an array actual argument or of
36 a sequence in a sequence association (12.5.2.12) is associated with the element of the dummy array that
37 has the same position in array element order (6.2.2.2).

NOTE 12.24

For type default character sequence associations, the interpretation of element is provided in 12.5.2.12.

38 A scalar dummy argument of a nonelemental procedure shall be associated only with a scalar actual
39 argument.

40 If a dummy argument has INTENT (OUT) or INTENT (INOUT), the actual argument shall be definable.

41 If a dummy argument has INTENT (OUT), the actual argument becomes undefined at the time the
42 association is established, except for direct components of an object of derived type for which default

- 1 initialization has been specified. If the dummy argument is not polymorphic and the type of the effective
 2 argument is an extension type of the type of the dummy argument, only the part of the effective argument
 3 that is of the same type as the dummy argument becomes undefined.
- 4 If the actual argument is an array section having a vector subscript, the dummy argument is not defin-
 5 able and shall not have the ASYNCHRONOUS, INTENT (OUT), INTENT (INOUT), or VOLATILE
 6 attributes.

NOTE 12.25

Argument intent specifications serve several purposes. See Note 5.17.

NOTE 12.26

For more explanatory information on argument association and evaluation, see subclause C.9.5.
 For more explanatory information on targets as dummy arguments, see subclause C.9.6.

J3 internal note

Unresolved Technical Issue 055

Probably need this constraint:

C1234 An actual argument that is a co-indexed object shall not be associated with a dummy argument that has the ASYNCHRONOUS attribute.

It is true that allowing VALUE+ASYNCHRONOUS makes a mockery of the restrictions on ASYNCHRONOUS dummies and note (04-007) 12.26 about their purpose. But should co-arrays try to be consistent or should they copy this inconsistency?

- 7 C1235 (R1223) If an actual argument is an array section or an assumed-shape array, and the corre-
 8 sponding dummy argument has either the VOLATILE or ASYNCHRONOUS attribute, that
 9 dummy argument shall be an assumed-shape array.
- 10 C1236 (R1223) If an actual argument is a pointer array, and the corresponding dummy argument
 11 has either the VOLATILE or ASYNCHRONOUS attribute, that dummy argument shall be an
 12 assumed-shape array that does not have the CONTIGUOUS attribute or a pointer array.

NOTE 12.27

The constraints on actual arguments that correspond to a dummy argument with either the ASYNCHRONOUS or VOLATILE attribute are designed to avoid forcing a processor to use the so-called copy-in/copy-out argument passing mechanism. Making a copy of actual arguments whose values are likely to change due to an asynchronous I/O operation completing or in some unpredictable manner will cause those new values to be lost when a called procedure returns and the copy-out overwrites the actual argument.

13 12.5.2.6 Allocatable and pointer dummy variables

14 The requirements in this subclause apply to actual arguments that correspond to either allocatable or
 15 pointer dummy data objects.

16 The actual argument shall be polymorphic if and only if the associated dummy argument is polymorphic,
 17 and either both the actual and dummy arguments shall be unlimited polymorphic, or the declared type
 18 of the actual argument shall be the same as the declared type of the dummy argument.

NOTE 12.28

The dynamic type of a polymorphic allocatable or pointer dummy argument may change as a result of execution of an allocate statement or pointer assignment in the subprogram. Because of this the corresponding actual argument needs to be polymorphic and have a declared type that

NOTE 12.28 (cont.)

is the same as the declared type of the dummy argument or an extension of that type. However, type compatibility requires that the declared type of the dummy argument be the same as, or an extension of, the type of the actual argument. Therefore, the dummy and actual arguments need to have the same declared type.

Dynamic type information is not maintained for a nonpolymorphic allocatable or pointer dummy argument. However, allocating or pointer assigning such a dummy argument would require maintenance of this information if the corresponding actual argument is polymorphic. Therefore, the corresponding actual argument needs to be nonpolymorphic.

1 The rank of the actual argument shall be the same as that of the dummy argument. The type parameter
2 values of the actual argument shall agree with the corresponding ones of the dummy argument that are
3 not assumed or deferred.

4 The values of assumed type parameters of a dummy argument are assumed from the corresponding type
5 parameters of the associated actual argument.

6 The actual argument shall have deferred the same type parameters as the dummy argument.

7 **12.5.2.7 Allocatable dummy variables**

8 The requirements in this subclause apply to actual arguments that correspond to allocatable dummy
9 data objects.

10 The actual argument shall be allocatable. It is permissible for the actual argument to have an allocation
11 status of unallocated.

12 If the dummy argument does not have the TARGET attribute, any pointers associated with the ac-
13 tual argument do not become associated with the corresponding dummy argument on invocation of the
14 procedure. If such a dummy argument is used as an actual argument that is associated with a dummy ar-
15 gument with the TARGET attribute, whether any pointers associated with the original actual argument
16 become associated with the dummy argument with the TARGET attribute is processor dependent.

17 If the dummy argument has the TARGET attribute, does not have the INTENT(OUT) or VALUE
18 attribute, and the corresponding actual argument has the TARGET attribute then

19 (1) any pointers associated with the actual argument become associated with the corresponding
20 dummy argument on invocation of the procedure, and

21 (2) when execution of the procedure completes, any pointers that do not become undefined
22 (16.5.2.2.3) and are associated with the dummy argument remain associated with the actual
23 argument.

24 If a dummy argument has INTENT (OUT) or INTENT (INOUT), the actual argument shall be definable.

25 If a dummy argument has INTENT (OUT), an allocated actual argument is deallocated on procedure
26 invocation (6.3.3.1).

27 **12.5.2.8 Pointer dummy variables**

28 The requirements in this subclause apply to actual arguments that correspond to dummy data pointers.

29 If the dummy argument does not have the INTENT(IN) attribute, the actual argument shall be a
30 pointer. Otherwise, the actual argument shall be a pointer or a valid target for the dummy pointer
31 in an assignment statement. If the actual argument is not a pointer, the dummy pointer becomes
32 pointer-associated with the actual argument.

- 1 The nondeferred type parameters and ranks shall agree.
- 2 If the dummy pointer has the CONTIGUOUS attribute and the actual argument does not have the
3 CONTIGUOUS attribute, the actual argument shall satisfy the following conditions.
- 4 • Its base object shall either have the CONTIGUOUS attribute or shall not be a pointer or assumed-
5 shape array.
 - 6 • It shall not be the real or imaginary part of an array of type complex.
 - 7 • Its designator shall not contain a *substring-range*.
 - 8 • It shall not have a vector subscript.
 - 9 • Only its final *part-ref* shall have nonzero rank.
 - 10 • If that *part-ref* has a *section-subscript-list*, the *section-subscript-list* shall satisfy these conditions:
 - 11 – no *stride* shall appear;
 - 12 – if any *section-subscript* is a subscript, it shall not be followed by a *subscript-triplet*;
 - 13 – all but the last *subscript-triplet* shall consist of a single colon with no *subscript*.

NOTE 12.29

This requires that the actual argument be obviously contiguous at compile time. Columns, planes, cubes and hypercubes of contiguous base objects can be passed, for example:

```

ARRAY1 (10:20, 3) ! passes part of the third column of ARRAY1.
X3D (:, i:j, 2)   ! passes part of the second plane of X3D (or the whole
                  ! plane if i==LBOUND(X3D,2) and j==UBOUND(X3D,2)).
Y5D (:, :, :, :, 7) ! passes the seventh hypercube of Y5D.

```

- 14 If the dummy argument has INTENT(OUT), the pointer association status of the actual argument
15 becomes undefined on invocation of the procedure.
- 16 If the dummy argument is nonoptional and the actual argument is allocatable, the actual argument shall
17 be allocated.

NOTE 12.30

For more explanatory information on pointers as dummy arguments, see subclause C.9.6.

18 **12.5.2.9 Co-array arguments**

- 19 The actual argument shall be an allocatable co-array if and only if the dummy argument is an allocatable
20 co-array with the same rank and co-rank.
- 21 If a dummy argument is a nonallocatable co-array, the co-rank and co-bounds are specified by the local
22 declaration and are independent of those of the actual argument. The actual argument shall be a co-array
23 or a subobject of a co-array without any co-subscripts, vector-valued subscripts, allocatable component
24 selection, or pointer component selection.

J3 internal note

Unresolved Technical Issue 054
Broken. I assume you want to be able to pass

`VARIABLE%CO_ARRAY_COMPONENT`

to a dummy co-array, no?

Subgroup says yes. It's not entirely trivial to fix the wording though. (It probably needs to be something like "no *part-ref* after the *part-ref* with co-rank non-zero shall have the pointer or allocatable attribute, or have a vector-valued subscript".)

NOTE 12.31

The co-array on the executing image is specified as the actual argument and this is associated with the dummy co-array argument on the same image. There are corresponding co-arrays on the other images. For example:

```
INTERFACE
  SUBROUTINE SUB(X, Y)
    REAL :: X(:)[*], Y(:)[*]
  END SUBROUTINE SUB
END INTERFACE
...
REAL, ALLOCATABLE :: A(:)[:], B(:,:)[:]
...
CALL SUB(A(:), B(1,:))
```

NOTE 12.32

The bounds and co-bounds of a dummy co-array may differ on each image executing a reference to the procedure. A co-indexed *data-ref* always uses the bounds and co-bounds on the executing image.

- 1 If an assumed-shape co-array that does not have the CONTIGUOUS attribute or a subobject of an
- 2 assumed-shape co-array that does not have the CONTIGUOUS attribute is an actual argument associ-
- 3 ated with a dummy co-array, the dummy co-array shall be of assumed shape.
- 4 If an array section is an actual argument associated with a dummy co-array that is not of assumed shape
- 5 or has the CONTIGUOUS attribute, the section shall be contiguous (5.3.6).

NOTE 12.33

The requirements on an actual argument that corresponds to a dummy co-array that is not of assumed-shape or has the CONTIGUOUS attribute are designed to avoid forcing a processor to use the so-called copy-in/copy-out argument passing mechanism.

J3 internal note

Unresolved Technical Issue 079

This constraint on dummy co-arrays and their arguments appears to be unnecessary without some relaxation in the rules in 12.5.2.14 (the "high performance" rules for Fortran dummy arguments).

6 12.5.2.10 Actual arguments associated with dummy procedure entities

- 1 If the actual argument is the name of an internal subprogram, the host instance of the dummy argument
2 is the innermost currently executing instance of the host of that internal subprogram. If the actual
3 argument has a host instance the host instance of the dummy argument is that instance. Otherwise the
4 dummy argument has no host instance.
- 5 If a dummy argument is a procedure pointer, the associated actual argument shall be a procedure pointer,
6 a reference to a function that returns a procedure pointer, or a reference to the NULL intrinsic function.
- 7 If a dummy argument is a dummy procedure without the POINTER attribute, the associated actual
8 argument shall be an external, internal, module, or dummy procedure, or a specific intrinsic procedure
9 listed in 13.6 and not marked with a bullet (●). If the specific name is also a generic name, only the
10 specific procedure is associated with the dummy argument.
- 11 If an external procedure name or a dummy procedure name is used as an actual argument, its interface
12 shall be explicit or it shall be explicitly declared to have the EXTERNAL attribute.
- 13 If the interface of a dummy procedure is explicit, the characteristics listed in 12.3.1 shall be the same
14 for the associated actual argument and the corresponding dummy argument, except that a pure actual
15 argument may be associated with a dummy argument that is not pure and an elemental intrinsic actual
16 procedure may be associated with a dummy procedure (which is prohibited from being elemental).
- 17 If the interface of a dummy procedure is implicit and either the dummy argument is explicitly typed
18 or referenced as a function, it shall not be referenced as a subroutine and any corresponding actual
19 argument shall be a function, function procedure pointer, or dummy procedure.
- 20 If the interface of a dummy procedure is implicit and a reference to it appears as a subroutine reference,
21 any corresponding actual argument shall be a subroutine, subroutine procedure pointer, or dummy
22 procedure.

23 **12.5.2.11 Actual arguments associated with alternate return indicators**

- 24 If a dummy argument is an asterisk (12.6.2.2), the associated actual argument shall be an alternate return specifier (12.5).

25 **12.5.2.12 Sequence association**

- 26 An actual argument represents an **element sequence** if it is an array expression, an array element
27 designator, a scalar of type default character, or a scalar of type character with the C character kind
28 (15.2.2). If the actual argument is an array expression, the element sequence consists of the elements
29 in array element order. If the actual argument is an array element designator, the element sequence
30 consists of that array element and each element that follows it in array element order.
- 31 If the actual argument is of type default character or of type character with the C character kind, and is
32 an array expression, array element, or array element substring designator, the element sequence consists
33 of the storage units beginning with the first storage unit of the actual argument and continuing to the
34 end of the array. The storage units of an array element substring designator are viewed as array elements
35 consisting of consecutive groups of storage units having the character length of the dummy array.
- 36 If the actual argument is of type default character or of type character with the C character kind, and
37 is a scalar that is not an array element or array element substring designator, the element sequence
38 consists of the storage units of the actual argument.

NOTE 12.34

Some of the elements in the element sequence may consist of storage units from different elements of the original array.
--

- 39 An actual argument that represents an element sequence and corresponds to a dummy argument that is

1 an array is sequence associated with the dummy argument if the dummy argument is an explicit-shape
2 or assumed-size array. The rank and shape of the actual argument need not agree with the rank and
3 shape of the dummy argument, but the number of elements in the dummy argument shall not exceed
4 the number of elements in the element sequence of the actual argument. If the dummy argument is
5 assumed-size, the number of elements in the dummy argument is exactly the number of elements in the
6 element sequence.

7 **12.5.2.13 Argument presence and restrictions on arguments not present**

8 A dummy argument or an entity that is host associated with a dummy argument is not **present** if the
9 dummy argument

- 10 (1) does not correspond to an actual argument,
- 11 (2) corresponds to an actual argument that is not present, or
- 12 (3) does not have the ALLOCATABLE or POINTER attribute, and corresponds to an actual
13 argument that
 - 14 (a) has the ALLOCATABLE attribute and is not allocated, or
 - 15 (b) has the POINTER attribute and is disassociated.

16 Otherwise, it is present. A nonoptional dummy argument shall be present. If an optional nonpointer
17 dummy argument corresponds to a pointer actual argument, the pointer association status of the actual
18 argument shall not be undefined.

19 An optional dummy argument that is not present is subject to the following restrictions.

- 20 (1) If it is a data object, it shall not be referenced or be defined. If it is of a type for which
21 default initialization is specified for a direct component, the initialization has no effect.
- 22 (2) It shall not be used as the *data-target* or *proc-target* of a pointer assignment.
- 23 (3) If it is a procedure or procedure pointer, it shall not be invoked.
- 24 (4) It shall not be supplied as an actual argument corresponding to a nonoptional dummy
25 argument other than as the argument of the PRESENT intrinsic function or as an argument
26 of a function reference that meets the requirements of (6) or (8) in 7.1.7.
- 27 (5) A designator with it as the base object and with at least one subobject selector shall not
28 be supplied as an actual argument.
- 29 (6) If it is an array, it shall not be supplied as an actual argument to an elemental procedure
30 unless an array of the same rank is supplied as an actual argument corresponding to a
31 nonoptional dummy argument of that elemental procedure.
- 32 (7) If it is a pointer, it shall not be allocated, deallocated, nullified, pointer-assigned, or supplied
33 as an actual argument corresponding to an optional nonpointer dummy argument.
- 34 (8) If it is allocatable, it shall not be allocated, deallocated, or supplied as an actual argument
35 corresponding to an optional nonallocatable dummy argument.
- 36 (9) If it has length type parameters, they shall not be the subject of an inquiry.
- 37 (10) It shall not be used as the *selector* in a SELECT TYPE or ASSOCIATE construct.

38 Except as noted in the list above, it may be supplied as an actual argument corresponding to an optional
39 dummy argument, which is then also considered not to be present.

40 **12.5.2.14 Restrictions on entities associated with dummy arguments**

41 While an entity is associated with a dummy argument, the following restrictions hold.

- 42 (1) Action that affects the allocation status of the entity or a subobject thereof shall be taken
43 through the dummy argument. Action that affects the value of the entity or any subobject
44 of it shall be taken only through the dummy argument unless

- 1 (a) the dummy argument has the POINTER attribute or
 2 (b) the dummy argument has the TARGET attribute, the dummy argument does not
 3 have INTENT (IN), the dummy argument is a scalar object or an assumed-shape
 4 array, and the actual argument is a target other than an array section with a vector
 5 subscript.

NOTE 12.35

In

```

SUBROUTINE OUTER
  REAL, POINTER :: A (:)
  ...
  ALLOCATE (A (1:N))
  ...
  CALL INNER (A)
  ...
CONTAINS
  SUBROUTINE INNER (B)
    REAL :: B (:)
    ...
  END SUBROUTINE INNER
  SUBROUTINE SET (C, D)
    REAL, INTENT (OUT) :: C
    REAL, INTENT (IN) :: D
    C = D
  END SUBROUTINE SET
END SUBROUTINE OUTER

```

an assignment statement such as

```
A (1) = 1.0
```

would not be permitted during the execution of INNER because this would be changing A without using B, but statements such as

```
B (1) = 1.0
```

or

```
CALL SET (B (1), 1.0)
```

would be allowed. Similarly,

```
DEALLOCATE (A)
```

would not be allowed because this affects the allocation of B without using B. In this case,

```
DEALLOCATE (B)
```

also would not be permitted. If B were declared with the POINTER attribute, either of the statements

```
DEALLOCATE (A)
```

NOTE 12.35 (cont.)

and

DEALLOCATE (B)

would be permitted, but not both.

NOTE 12.36

If there is a partial or complete overlap between the actual arguments associated with two different dummy arguments of the same procedure and the dummy arguments have neither the POINTER nor TARGET attribute, the overlapped portions shall not be defined, redefined, or become undefined during the execution of the procedure. For example, in

```
CALL SUB (A (1:5), A (3:9))
```

A (3:5) shall not be defined, redefined, or become undefined through the first dummy argument because it is part of the argument associated with the second dummy argument and shall not be defined, redefined, or become undefined through the second dummy argument because it is part of the argument associated with the first dummy argument. A (1:2) remains definable through the first dummy argument and A (6:9) remains definable through the second dummy argument.

NOTE 12.37

This restriction applies equally to pointer targets. In

```
REAL, DIMENSION (10), TARGET :: A
REAL, DIMENSION (:), POINTER :: B, C
B => A (1:5)
C => A (3:9)
```

```
CALL SUB (B, C) ! The dummy arguments of SUB are neither pointers nor targets.
```

B (3:5) cannot be defined because it is part of the argument associated with the second dummy argument. C (1:3) cannot be defined because it is part of the argument associated with the first dummy argument. A (1:2) [which is B (1:2)] remains definable through the first dummy argument and A (6:9) [which is C (4:7)] remains definable through the second dummy argument.

NOTE 12.38

Because a nonpointer dummy argument declared with INTENT(IN) shall not be used to change the associated actual argument, the associated actual argument remains constant throughout the execution of the procedure.

- 1 (2) If the allocation status of the entity or a subobject thereof is affected through the dummy
 2 argument, then at any time during the execution of the procedure, either before or after the
 3 allocation or deallocation, it may be referenced only through the dummy argument. If the
 4 value of the entity or any subobject of it is affected through the dummy argument, then at
 5 any time during the execution of the procedure, either before or after the definition, it may
 6 be referenced only through that dummy argument unless
 7 (a) the dummy argument has the POINTER attribute or
 8 (b) the dummy argument has the TARGET attribute, the dummy argument does not
 9 have INTENT (IN), the dummy argument is a scalar object or an assumed-shape
 10 array, and the actual argument is a target other than an array section with a vector
 11 subscript.

NOTE 12.39

In

```

MODULE DATA
  REAL :: W, X, Y, Z
END MODULE DATA

PROGRAM MAIN
  USE DATA
  ...
  CALL INIT (X)
  ...
END PROGRAM MAIN
SUBROUTINE INIT (V)
  USE DATA
  ...
  READ (*, *) V
  ...
END SUBROUTINE INIT

```

variable X shall not be directly referenced at any time during the execution of INIT because it is being defined through the dummy argument V. X may be (indirectly) referenced through V. W, Y, and Z may be directly referenced. X may, of course, be directly referenced once execution of INIT is complete.

NOTE 12.40

The restrictions on entities associated with dummy arguments are intended to facilitate a variety of optimizations in the translation of the subprogram, including implementations of argument association in which the value of an actual argument that is neither a pointer nor a target is maintained in a register or in local storage.

1 12.5.3 Function reference

2 A function is invoked during expression evaluation by a *function-reference* or by a defined operation
 3 (7.1.3). When it is invoked, all actual argument expressions are evaluated, then the arguments are
 4 associated, and then the function is executed. When execution of the function is complete, the value
 5 of the function result is available for use in the expression that caused the function to be invoked. The
 6 characteristics of the function result (12.3.3) are determined by the interface of the function. A reference
 7 to an elemental function (12.8) is an elemental reference if one or more actual arguments are arrays and
 8 all array arguments have the same shape.

9 12.5.4 Subroutine reference

10 A subroutine is invoked by execution of a CALL statement, execution of a defined assignment statement
 11 (7.4.1.4), user-defined derived-type input/output(9.5.4.7.1), or finalization(4.5.6). When a subroutine is
 12 invoked, all actual argument expressions are evaluated, then the arguments are associated, and then the
 13 subroutine is executed. When the actions specified by the subroutine are completed, the execution of the
 14 CALL statement, the execution of the defined assignment statement, the processing of an input or output
 15 list item, or finalization of an object is also completed. If a CALL statement includes one or more alternate
 16 return specifiers among its arguments, control may be transferred to one of the statements indicated, depending on the
 17 action specified by the subroutine. A reference to an elemental subroutine (12.8) is an elemental reference if
 18 there is at least one actual argument corresponding to an INTENT(OUT) or INTENT(INOUT) dummy

1 argument, all such actual arguments are arrays, and all actual arguments are conformable.

2 **12.5.5 Resolving named procedure references**

3 **12.5.5.1 Establishment of procedure names**

4 The rules for interpreting a procedure reference depend on whether the procedure name in the reference
5 is established by the available declarations and specifications to be generic in the scoping unit containing
6 the reference, is established to be only specific in the scoping unit containing the reference, or is not
7 established.

8 A procedure name is established to be generic in a scoping unit

- 9 (1) if that scoping unit contains an interface block with that name;
- 10 (2) if that scoping unit contains an INTRINSIC attribute specification for that name and it is
11 the name of a generic intrinsic procedure;
- 12 (3) if that scoping unit contains a USE statement that makes that procedure name accessible
13 and the corresponding name in the module is established to be generic; or
- 14 (4) if that scoping unit contains no declarations of that name, that scoping unit has a host
15 scoping unit, and that name is established to be generic in the host scoping unit.

16 A procedure name is established to be only specific in a scoping unit if it is established to be specific
17 and not established to be generic. It is established to be specific

- 18 (1) if that scoping unit contains a module subprogram, internal subprogram, or statement function
19 that defines a procedure with that name;
- 20 (2) if that scoping unit contains an INTRINSIC attribute specification for that name and it is
21 the name of a specific intrinsic procedure;
- 22 (3) if that scoping unit contains an explicit EXTERNAL attribute specification (5.3.8) for that
23 name;
- 24 (4) if that scoping unit contains a USE statement that makes that procedure name accessible
25 and the corresponding name in the module is established to be specific; or
- 26 (5) if that scoping unit contains no declarations of that name, that scoping unit has a host
27 scoping unit, and that name is established to be specific in the host scoping unit.

28 A procedure name is not established in a scoping unit if it is neither established to be generic nor
29 established to be specific.

30 **12.5.5.2 Resolving procedure references to names established to be generic**

31 If the reference is consistent with a nonelemental reference to one of the specific interfaces of a generic
32 interface that has that name and either is in the scoping unit in which the reference appears or is made
33 accessible by a USE statement in the scoping unit, the reference is to the specific procedure in the
34 interface block that provides that interface. The rules in 12.4.3.3.4 ensure that there can be at most one
35 such specific procedure.

36 Otherwise, if the reference is consistent with an elemental reference to one of the specific interfaces of
37 a generic interface that has that name and either is in the scoping unit in which the reference appears
38 or is made accessible by a USE statement in the scoping unit, the reference is to the specific elemental
39 procedure in the interface block that provides that interface. The rules in 12.4.3.3.4 ensure that there
40 can be at most one such specific elemental procedure.

41 Otherwise, if the scoping unit contains either an INTRINSIC attribute specification for that name or
42 a USE statement that makes that name accessible from a module in which the corresponding name is
43 specified to have the INTRINSIC attribute, and if the reference is consistent with the interface of that
intrinsic procedure, the reference is to that intrinsic procedure.

1

2 Otherwise, if the scoping unit has a host scoping unit, the name is established to be generic in that host
 3 scoping unit, and there is agreement between the scoping unit and the host scoping unit as to whether
 4 the name is a function name or a subroutine name, the name is resolved by applying the rules in this
 5 subclause to the host scoping unit.

6 Otherwise, if the name is that of an intrinsic procedure and the reference is consistent with that intrinsic
 7 procedure, the reference is to that intrinsic procedure.

NOTE 12.41

These rules allow particular specific procedures of a generic procedure to be used for particular array ranks and a general elemental version to be used for other ranks. For example, given an interface block such as:

```
INTERFACE RANF
  ELEMENTAL FUNCTION SCALAR_RANF(X)
    REAL, INTENT(IN) :: X
  END FUNCTION SCALAR_RANF
  FUNCTION VECTOR_RANDOM(X)
    REAL X(:)
    REAL VECTOR_RANDOM(SIZE(X))
  END FUNCTION VECTOR_RANDOM
END INTERFACE RANF
```

and a declaration such as:

```
REAL A(10,10), AA(10,10)
```

then the statement

```
A = RANF(AA)
```

is an elemental reference to SCALAR_RANF. The statement

```
A = RANF(AA(6:10,2))
```

is a nonelemental reference to VECTOR_RANDOM.

NOTE 12.42

In the USE statement case, it is possible, because of the renaming facility, for the name in the reference to be different from the name of the intrinsic procedure.

8 **12.5.5.3 Resolving procedure references to names established to be only specific**

9 If the scoping unit contains an interface body or EXTERNAL attribute specification for the name and
 10 the name is the name of a dummy argument of the scoping unit, the dummy argument is a dummy
 11 procedure and the reference is to that dummy procedure. That is, the procedure invoked by executing
 12 that reference is the procedure supplied as the actual argument corresponding to that dummy procedure.

13 If the scoping unit contains an interface body or EXTERNAL attribute specification for the name and
 14 the name is not the name of a dummy argument of the scoping unit, the reference is to an external
 15 procedure with that name.

- 1 If the scoping unit contains a module subprogram, internal subprogram, or statement function that defines
2 a procedure with the name, the reference is to the procedure so defined.
- 3 If the scoping unit contains an INTRINSIC attribute specification for the name, the reference is to the
4 intrinsic with that name.
- 5 If the scoping unit contains a USE statement that makes a procedure accessible by the name, the
6 reference is to that procedure.

NOTE 12.43

Because of the renaming facility of the USE statement, the name in the reference may be different from the original name of the procedure.

- 7 If none of the above apply, the scoping unit shall have a host scoping unit, and the reference is resolved
8 by applying the rules in this subclause to the host scoping unit.

9 12.5.5.4 Resolving procedure references to names not established

- 10 If the name is the name of a dummy argument of the scoping unit, the dummy argument is a dummy
11 procedure and the reference is to that dummy procedure. That is, the procedure invoked by executing
12 that reference is the procedure supplied as the actual argument corresponding to that dummy procedure.
- 13 Otherwise, if the name is the name of an intrinsic procedure, and if there is agreement between the
14 reference and the status of the intrinsic procedure as being a function or subroutine, the reference is to
15 that intrinsic procedure.
- 16 Otherwise, the reference is to an external procedure with that name.

17 12.5.6 Resolving type-bound procedure references

- 18 If the *binding-name* in a *procedure-designator* (R1221) is that of a specific type-bound procedure, the
19 procedure referenced is the one bound to that name in the dynamic type of the *data-ref*.
- 20 If the *binding-name* in a *procedure-designator* is that of a generic type bound procedure, the generic
21 binding with that name in the declared type of the *data-ref* is used to select a specific binding using the
22 following criteria.
- 23 (1) If the reference is consistent with one of the specific bindings of that generic binding, that
24 specific binding is selected.
 - 25 (2) Otherwise, the reference shall be consistent with an elemental reference to one of the specific
26 bindings of that generic binding; that specific binding is selected.

- 27 The reference is to the procedure bound to the same name as the selected specific binding in the dynamic
28 type of the *data-ref*.

29 12.6 Procedure definition**30 12.6.1 Intrinsic procedure definition**

- 31 Intrinsic procedures are defined as an inherent part of the processor. A standard-conforming processor
32 shall include the intrinsic procedures described in Clause 13, but may include others. However, a
33 standard-conforming program shall not make use of intrinsic procedures other than those described in
34 Clause 13.

1 12.6.2 Procedures defined by subprograms

2 When a procedure defined by a subprogram is invoked, an instance (12.6.2.3) of the subprogram is
3 created and executed.

4 12.6.2.1 Function subprogram

5 A **function subprogram** is a subprogram that has a FUNCTION statement as its first statement.

```
6 R1225 function-subprogram      is function-stmt
7                               [ specification-part ]
8                               [ execution-part ]
9                               [ internal-subprogram-part ]
10                              end-function-stmt
11 R1226 function-stmt           is [ prefix ] FUNCTION function-name ■
12                               ■ ( [ dummy-arg-name-list ] ) [ suffix ]
```

13 C1237 (R1226) If RESULT appears, *result-name* shall not be the same as *function-name* and shall not
14 be the same as the *entry-name* in any ENTRY statement in the subprogram.

15 C1238 (R1226) If RESULT appears, the *function-name* shall not appear in any specification statement
16 in the scoping unit of the function subprogram.

```
17 R1227 proc-language-binding-spec is language-binding-spec
```

18 C1239 (R1227) A *proc-language-binding-spec* with a NAME= specifier shall not be specified in the
19 *function-stmt* or *subroutine-stmt* of an interface body for an abstract interface or a dummy
20 procedure.

21 C1240 (R1227) A *proc-language-binding-spec* shall not be specified for an internal procedure.

22 C1241 (R1227) If *proc-language-binding-spec* is specified for a procedure, each of the procedure's dummy
23 arguments shall be a nonoptional interoperable variable (15.3.5, 15.3.6) or a nonoptional interoper-
24 able procedure (15.3.7). If *proc-language-binding-spec* is specified for a function, the function
25 result shall be an interoperable scalar variable.

```
26 R1228 dummy-arg-name          is name
```

27 C1242 (R1228) A *dummy-arg-name* shall be the name of a dummy argument.

```
28 R1229 prefix                  is prefix-spec [ prefix-spec ] ...
```

```
29 R1230 prefix-spec             is declaration-type-spec
30                               or ELEMENTAL
31                               or IMPURE
32                               or MODULE
33                               or PURE
34                               or RECURSIVE
```

35 C1243 (R1229) A *prefix* shall contain at most one of each *prefix-spec*.

36 C1244 (R1229) A *prefix* shall not specify both PURE and IMPURE.

37 C1245 (R1229) A *prefix* shall not specify both ELEMENTAL and RECURSIVE.

38 C1246 (R1229) A *prefix* shall not specify ELEMENTAL if *proc-language-binding-spec* appears in the
39 *function-stmt* or *subroutine-stmt*.

40 C1247 (R1229) MODULE shall appear only within the *function-stmt* or *subroutine-stmt* of a module

- 1 subprogram or of an interface body that is declared in the scoping unit of a module or submodule.
- 2 C1248 (R1229) If MODULE appears within the *prefix* in a module subprogram, an accessible module
3 procedure interface having the same name as the subprogram shall be declared in the module
4 or submodule in which the subprogram is defined, or shall be declared in an ancestor of that
5 program unit.
- 6 C1249 (R1229) If MODULE appears within the *prefix* in a module subprogram, the subprogram shall
7 specify the same characteristics and dummy argument names as its corresponding (12.6.2.4)
8 module procedure interface body.
- 9 C1250 (R1229) If MODULE appears within the *prefix* in a module subprogram and a binding label
10 is specified, it shall be the same as the binding label specified in the corresponding module
11 procedure interface body.
- 12 C1251 (R1229) If MODULE appears within the *prefix* in a module subprogram, RECURSIVE shall
13 appear if and only if RECURSIVE appears in the *prefix* in the corresponding module procedure
14 interface body.
- 15 R1231 *suffix* **is** *proc-language-binding-spec* [RESULT (*result-name*)]
16 **or** RESULT (*result-name*) [*proc-language-binding-spec*]
- 17 R1232 *end-function-stmt* **is** END [FUNCTION [*function-name*]]
- 18 C1252 (R1225) An internal function subprogram shall not contain an ENTRY statement.
- 19 C1253 (R1225) An internal function subprogram shall not contain an *internal-subprogram-part*.
- 20 C1254 (R1232) If a *function-name* appears in the *end-function-stmt*, it shall be identical to the *function-*
21 *name* specified in the *function-stmt*.
- 22 The name of the function is *function-name*.
- 23 The type and type parameters (if any) of the result of the function defined by a function subprogram
24 may be specified by a type specification in the FUNCTION statement or by the name of the result
25 variable appearing in a type declaration statement in the declaration part of the function subprogram.
26 They shall not be specified both ways. If they are not specified either way, they are determined by
27 the implicit typing rules in force within the function subprogram. If the function result is an array,
28 allocatable, or a pointer, this shall be specified by specifications of the name of the result variable within
29 the function body. The specifications of the function result attributes, the specification of dummy
30 argument attributes, and the information in the procedure heading collectively define the characteristics
31 of the function (12.3.1).
- 32 The RECURSIVE *prefix-spec* shall appear if any procedure defined by the subprogram directly or indi-
33 rectly invokes itself or any other procedure defined by the subprogram.

NOTE 12.40a

Each ENTRY statement in the subprogram defines an additional function.

- 34 If RESULT appears, the name of the result variable of the function is *result-name* and all occurrences
35 of the function name in *execution-part* statements in the scoping unit refer to the function itself. If
36 RESULT does not appear, the result variable is *function-name* and all occurrences of the function name
37 in *execution-part* statements in the scoping unit are references to the result variable. The characteristics
38 (12.3.3) of the function result are those of the result variable. On completion of execution of the function,
39 the value returned is that of its result variable. If the function result is a pointer, the shape of the value
40 returned by the function is determined by the shape of the result variable when the execution of the

- 1 function is completed. If the result variable is not a pointer, its value shall be defined by the function.
- 2 If the function result is a pointer, on return the pointer association status of the result variable shall not
- 3 be undefined.

NOTE 12.44

The result variable is similar to any other variable local to a function subprogram. Its existence begins when execution of the function is initiated and ends when execution of the function is terminated. However, because the final value of this variable is used subsequently in the evaluation of the expression that invoked the function, an implementation may wish to defer releasing the storage occupied by that variable until after its value has been used in expression evaluation.

- 4 If the *prefix-spec* PURE appears, or the *prefix-spec* ELEMENTAL appears and IMPURE does not appear,
- 5 the subprogram is a pure subprogram and shall meet the additional constraints of 12.7.
- 6 If the *prefix-spec* ELEMENTAL appears, the subprogram is an elemental subprogram and shall meet
- 7 the additional constraints of 12.8.1.

NOTE 12.45

An example of a recursive function is:

```

RECURSIVE FUNCTION CUMM_SUM (ARRAY) RESULT (C_SUM)
  REAL, INTENT (IN), DIMENSION (:): ARRAY
  REAL, DIMENSION (SIZE (ARRAY)): C_SUM
  INTEGER N
  N = SIZE (ARRAY)
  IF (N <= 1) THEN
    C_SUM = ARRAY
  ELSE
    N = N / 2
    C_SUM (:N) = CUMM_SUM (ARRAY (:N))
    C_SUM (N+1:) = C_SUM (N) + CUMM_SUM (ARRAY (N+1:))
  END IF
END FUNCTION CUMM_SUM

```

NOTE 12.46

The following is an example of the declaration of an interface body with the BIND attribute, and a reference to the procedure declared.

```

USE, INTRINSIC :: ISO_C_BINDING

INTERFACE
  FUNCTION JOE (I, J, R) BIND(C,NAME="FrEd")
    USE, INTRINSIC :: ISO_C_BINDING
    INTEGER(C_INT) :: JOE
    INTEGER(C_INT), VALUE :: I, J
    REAL(C_FLOAT), VALUE :: R
  END FUNCTION JOE
END INTERFACE

INT = JOE(1_C_INT, 3_C_INT, 4.0_C_FLOAT)
END PROGRAM

```

NOTE 12.46 (cont.)

The invocation of the function JOE results in a reference to a function with a binding label "FrEd". FrEd may be a C function described by the C prototype

```
int FrEd(int n, int m, float x);
```

1 **12.6.2.2 Subroutine subprogram**

2 A **subroutine subprogram** is a subprogram that has a SUBROUTINE statement as its first statement.

3 R1233 *subroutine-subprogram* **is** *subroutine-stmt*
 4 [*specification-part*]
 5 [*execution-part*]
 6 [*internal-subprogram-part*]
 7 *end-subroutine-stmt*
 8 R1234 *subroutine-stmt* **is** [*prefix*] SUBROUTINE *subroutine-name* ■
 9 ■ [([*dummy-arg-list*]) [*proc-language-binding-spec*]]

10 C1255 (R1234) The *prefix* of a *subroutine-stmt* shall not contain a *declaration-type-spec*.

11 R1235 *dummy-arg* **is** *dummy-arg-name*
 12 **or** *

13 R1236 *end-subroutine-stmt* **is** END [SUBROUTINE [*subroutine-name*]]

14 C1256 (R1233) An internal subroutine subprogram shall not contain an ENTRY statement.

15 C1257 (R1233) An internal subroutine subprogram shall not contain an *internal-subprogram-part*.

16 C1258 (R1236) If a *subroutine-name* appears in the *end-subroutine-stmt*, it shall be identical to the
 17 *subroutine-name* specified in the *subroutine-stmt*.

18 The name of the subroutine is *subroutine-name*.

19 The RECURSIVE *prefix-spec* shall appear if any procedure defined by the subprogram directly or indi-
 20 rectly invokes itself or any other procedure defined by the subprogram.

NOTE 12.43a

Each ENTRY statement in the subprogram defines an additional subroutine.

21 If the *prefix-spec* PURE appears, or the *prefix-spec* ELEMENTAL appears and IMPURE does not appear,
 22 the subprogram is a pure subprogram and shall meet the additional constraints of 12.7.

23 If the *prefix-spec* ELEMENTAL appears, the subprogram is an elemental subprogram and shall meet
 24 the additional constraints of 12.8.1.

25 **12.6.2.3 Instances of a subprogram**

26 When a procedure defined by a subprogram is invoked, an **instance** of that subprogram is created.
 27 Execution begins with the first executable construct following the FUNCTION, SUBROUTINE, or
 28 ENTRY statement specifying the name of the procedure invoked.

29 When a statement function is invoked, an instance of that statement function is created.

30 When execution of an instance completes it ceases to exist.

1 Each instance has an independent sequence of execution and an independent set of dummy arguments
 2 and local unsaved data objects. If an internal procedure or statement function in the subprogram is invoked
 3 by name from an instance of the subprogram or from an internal subprogram or statement function that
 4 has access to the entities of that instance, the created instance of the internal subprogram or statement
 5 function also has access to the entities of that instance of the host subprogram. If an internal procedure
 6 is invoked via a dummy procedure or procedure pointer, the internal procedure has access to the entities
 7 of the host instance of that dummy procedure or procedure pointer.

8 All other entities are shared by all instances of the subprogram.

NOTE 12.47

The value of a saved data object appearing in one instance may have been defined in a previous instance or by initialization in a DATA statement or type declaration statement.

9 **12.6.2.4 Separate module procedures**

10 A **separate module procedure** is a module procedure defined by a *separate-module-subprogram*,
 11 by a *function-subprogram* whose initial statement contains the keyword MODULE, or by a *subroutine-*
 12 *subprogram* whose initial statement contains the keyword MODULE. Its interface is declared by a module
 13 procedure interface body (12.4.3.2) in the *specification-part* of the module or submodule in which the
 14 procedure is defined, or in an ancestor module or submodule.

15 R1237 *separate-module-subprogram* **is** *mp-subprogram-stmt*
 16 [*specification-part*]
 17 [*execution-part*]
 18 [*internal-subprogram-part*]
 19 *end-mp-subprogram-stmt*

20 R1238 *mp-subprogram-stmt* **is** MODULE PROCEDURE *procedure-name*
 21

22 R1239 *end-mp-subprogram-stmt* **is** END [PROCEDURE [*procedure-name*]]

23 C1259 (R1237) The *procedure-name* shall be the same as the name of an accessible module procedure
 24 interface that is declared in the module or submodule in which the *separate-module-subprogram*
 25 is defined, or is declared in an ancestor of that program unit.

26 C1260 (R1239) If a *procedure-name* appears in the *end-mp-subprogram-stmt*, it shall be identical to the
 27 *procedure-name* in the MODULE PROCEDURE statement.

28 A module procedure interface body and a subprogram that defines a separate module procedure **corre-**
 29 **spond** if they have the same name, and the module procedure interface is declared in the same program
 30 unit as the subprogram or is declared in an ancestor of the program unit in which the procedure is
 31 defined and is accessible by host association from that ancestor. A module procedure interface body
 32 shall not correspond to more than one subprogram that defines a separate module procedure.

NOTE 12.48

A separate module procedure can be accessed by use association only if its interface body is declared in the specification part of a module and is public.

33 If a procedure is defined by a *separate-module-subprogram*, its characteristics are specified by the corre-
 34 sponding module procedure interface body.

35 If a separate module procedure is a function defined by a *separate-module-subprogram*, the result variable
 36 name is determined by the FUNCTION statement in the module procedure interface body. Otherwise,

1 the result variable name is determined by the FUNCTION statement in the module subprogram.

2 12.6.2.5 ENTRY statement

3 An **ENTRY statement** permits a procedure reference to begin with a particular executable statement
4 within the function or subroutine subprogram in which the ENTRY statement appears.

5 R1240 *entry-stmt* **is** ENTRY *entry-name* [([*dummy-arg-list*]) [*suffix*]]

6 C1261 (R1240) If RESULT appears, the *entry-name* shall not appear in any specification or type-
7 declaration statement in the scoping unit of the function program.

8 C1262 (R1240) An *entry-stmt* shall appear only in an *external-subprogram* or a *module-subprogram*
9 that does not define a separate module procedure. An *entry-stmt* shall not appear within an
10 *executable-construct*.

11 C1263 (R1240) RESULT shall appear only if the *entry-stmt* is in a function subprogram.

12 C1264 (R1240) Within the subprogram containing the *entry-stmt*, the *entry-name* shall not appear
13 as a dummy argument in the FUNCTION or SUBROUTINE statement or in another ENTRY
14 statement nor shall it appear in an EXTERNAL, INTRINSIC, or PROCEDURE statement.

15 C1265 (R1240) A *dummy-arg* shall not be an alternate return indicator if the ENTRY statement is in a function
16 subprogram.

17 C1266 (R1240) If RESULT appears, *result-name* shall not be the same as the *function-name* in the
18 FUNCTION statement and shall not be the same as the *entry-name* in any ENTRY statement
19 in the subprogram.

20 Optionally, a subprogram may have one or more ENTRY statements.

21 If the ENTRY statement is in a function subprogram, an additional function is defined by that subpro-
22 gram. The name of the function is *entry-name* and the name of its result variable is *result-name* or
23 is *entry-name* if no *result-name* is provided. The characteristics of the function result are specified by
24 specifications of the result variable. The dummy arguments of the function are those specified in the
25 ENTRY statement. If the characteristics of the result of the function named in the ENTRY statement
26 are the same as the characteristics of the result of the function named in the FUNCTION statement,
27 their result variables identify the same variable, although their names need not be the same. Otherwise,
28 they are storage associated and shall all be nonpointer, nonallocatable scalars of type default integer,
29 default real, double precision real, default complex, or default logical.

30 If the ENTRY statement is in a subroutine subprogram, an additional subroutine is defined by that
31 subprogram. The name of the subroutine is *entry-name*. The dummy arguments of the subroutine are
32 those specified in the ENTRY statement.

33 The order, number, types, kind type parameters, and names of the dummy arguments in an ENTRY
34 statement may differ from the order, number, types, kind type parameters, and names of the dummy
35 arguments in the FUNCTION or SUBROUTINE statement in the containing subprogram.

36 Because an ENTRY statement defines an additional function or an additional subroutine, it is referenced
37 in the same manner as any other function or subroutine (12.5).

38 In a subprogram, a name that appears as a dummy argument in an ENTRY statement shall not appear
39 in an executable statement preceding that ENTRY statement, unless it also appears in a FUNCTION,
40 SUBROUTINE, or ENTRY statement that precedes the executable statement.

J3 internal note

Unresolved Technical Issue 083

The above makes it invalid to use it as an *ac-implicit-do* variable (in an executable statement) ahead of the ENTRY, even though it has been established (by interp) that the use as an *ac-implicit-do* variable does not create a local entity.

Similarly to have it used as an *associate-name* in an ASSOCIATE or SELECT TYPE construct, or to declare a construct entity with that name in a BLOCK construct.

The latter should *obviously* be allowed, as should be *associate-name*. Maybe there should be an interp request for F2003, but in any case we should get this right for F2008.

I'll note that the following paragraph gets it right for statement function dummy arguments; if we can do that for an obsolescent feature, we can do it for new ones too.

The paragraph after that is ok, because it talks about the dummy argument appearing, whereas the above paragraph talks about the name.

1 In a subprogram, a name that appears as a dummy argument in an ENTRY statement shall not appear in the expression
2 of a statement function unless the name is also a dummy argument of the statement function, appears in a FUNCTION
3 or SUBROUTINE statement, or appears in an ENTRY statement that precedes the statement function statement.

4 If a dummy argument appears in an executable statement, the execution of the executable statement is
5 permitted during the execution of a reference to the function or subroutine only if the dummy argument
6 appears in the dummy argument list of the procedure name referenced.

7 If a dummy argument is used in a specification expression to specify an array bound or character length
8 of an object, the appearance of the object in a statement that is executed during a procedure reference
9 is permitted only if the dummy argument appears in the dummy argument list of the procedure name
10 referenced and it is present (12.5.2.13).

11 A scoping unit containing a reference to a procedure defined by an ENTRY statement may have access to
12 an interface body for the procedure. The procedure header for the interface body shall be a FUNCTION
13 statement for an entry in a function subprogram and shall be a SUBROUTINE statement for an entry
14 in a subroutine subprogram.

15 The keyword RECURSIVE is not used in an ENTRY statement. Instead, the presence or absence of
16 RECURSIVE in the initial SUBROUTINE or FUNCTION statement controls whether the procedure
17 defined by an ENTRY statement is permitted to reference itself or another procedure defined by the
18 subprogram.

19 The keywords PURE and IMPURE are not used in an ENTRY statement. Instead, the procedure
20 defined by an ENTRY statement is pure if and only if the subprogram is a pure subprogram.

21 The keyword ELEMENTAL is not used in an ENTRY statement. Instead, the procedure defined by
22 an ENTRY statement is elemental if and only if ELEMENTAL is specified in the SUBROUTINE or
23 FUNCTION statement.

24 12.6.2.6 RETURN statement

25 R1241 *return-stmt* **is** RETURN [*scalar-int-expr*]

26 C1267 (R1241) The *return-stmt* shall be in the scoping unit of a function or subroutine subprogram.

27 C1268 (R1241) The *scalar-int-expr* is allowed only in the scoping unit of a subroutine subprogram.

28 Execution of the RETURN statement completes execution of the instance of the subprogram in which
29 it appears. If the expression appears and has a value *n* between 1 and the number of asterisks in the dummy argument
30 list, the CALL statement that invoked the subroutine transfers control to the statement identified by the *n*th alternate
31 return specifier in the actual argument list of the referenced procedure. If the expression is omitted or has a value outside

1 the required range, there is no transfer of control to an alternate return.

2 Execution of an *end-function-stmt* or *end-subroutine-stmt* is equivalent to executing a RETURN state-
3 ment with no expression.

4 **12.6.2.7 CONTAINS statement**

5 R1242 *contains-stmt* is CONTAINS

6 The CONTAINS statement separates the body of a main program, module, submodule, or subprogram
7 from any internal or module subprograms it may contain, or it introduces the type-bound procedure
8 part of a derived-type definition (4.5.2). The CONTAINS statement is not executable.

9 **12.6.3 Definition and invocation of procedures by means other than Fortran**

10 A procedure may be defined by means other than Fortran. The interface of a procedure defined by
11 means other than Fortran may be specified by an interface body or procedure declaration statement. If
12 the interface of such a procedure does not have a *proc-language-binding-spec*, the means by which the
13 procedure is defined are processor dependent. A reference to such a procedure is made as though it were
14 defined by an external subprogram.

15 If the interface of a procedure has a *proc-language-binding-spec*, the procedure is interoperable (15.5).

16 Interoperation with C functions is described in 15.5.

NOTE 12.49

For explanatory information on definition of procedures by means other than Fortran, see subclause C.9.2.

17 **12.6.4 Statement function**

18 A statement function is a function defined by a single statement.

19 R1243 *stmt-function-stmt* is *function-name* ([*dummy-arg-name-list*]) = *scalar-expr*

20 C1269 (R1243) The *primaries* of the *scalar-expr* shall be constants (literal and named), references to variables, references
21 to functions and function dummy procedures, and intrinsic operations. If *scalar-expr* contains a reference to a
22 function or a function dummy procedure, the reference shall not require an explicit interface, the function shall
23 not require an explicit interface unless it is an intrinsic function, the function shall not be a transformational
24 intrinsic, and the result shall be scalar. If an argument to a function or a function dummy procedure is an array,
25 it shall be an array name. If a reference to a statement function appears in *scalar-expr*, its definition shall have
26 been provided earlier in the scoping unit and shall not be the name of the statement function being defined.

27 C1270 (R1243) Named constants in *scalar-expr* shall have been declared earlier in the scoping unit or made accessible
28 by use or host association. If array elements appear in *scalar-expr*, the array shall have been declared as an array
29 earlier in the scoping unit or made accessible by use or host association.

30 C1271 (R1243) If a *dummy-arg-name*, variable, function reference, or dummy function reference is typed by the implicit
31 typing rules, its appearance in any subsequent type declaration statement shall confirm this implied type and
32 the values of any implied type parameters.

33 C1272 (R1243) The *function-name* and each *dummy-arg-name* shall be specified, explicitly or implicitly, to be scalar.

34 C1273 (R1243) A given *dummy-arg-name* shall not appear more than once in any *dummy-arg-name-list*.

35 C1274 (R1243) Each variable reference in *scalar-expr* may be either a reference to a dummy argument of the statement
function or a reference to a variable accessible in the same scoping unit as the statement function statement.

- 1 The definition of a statement function with the same name as an accessible entity from the host shall be preceded by the
2 declaration of its type in a type declaration statement.
- 3 The dummy arguments have a scope of the statement function statement. Each dummy argument has the same type and
4 type parameters as the entity of the same name in the scoping unit containing the statement function.
- 5 A statement function shall not be supplied as a procedure argument.
- 6 The value of a statement function reference is obtained by evaluating the expression using the values of the actual arguments
7 for the values of the corresponding dummy arguments and, if necessary, converting the result to the declared type and
8 type parameters of the function.
- 9 A function reference in the scalar expression shall not cause a dummy argument of the statement function to become
10 redefined or undefined.

11 12.7 Pure procedures

12 A pure procedure is

- 13 (1) a pure intrinsic procedure (13.1),
14 (2) defined by a pure subprogram, or
15 (3) a statement function that references only pure functions.

16 A pure subprogram is a subprogram that has the *prefix-spec* PURE or that has the *prefix-spec* ELE-
17 MENTAL and does not have the *prefix-spec* IMPURE. The following additional constraints apply to
18 pure subprograms.

- 19 C1275 The *specification-part* of a pure function subprogram shall specify that all its nonpointer dummy
20 data objects have INTENT(IN).
- 21 C1276 The *specification-part* of a pure subroutine subprogram shall specify the intents of all its non-
22 pointer dummy data objects.
- 23 C1277 A local variable declared in the *specification-part* of a pure subprogram, or within the *specification-*
24 *part* of a BLOCK construct within a pure subprogram, shall not have the SAVE attribute.

NOTE 12.50

Variable initialization in a *type-declaration-stmt* or a *data-stmt* implies the SAVE attribute; therefore, such initialization is also disallowed.

- 25 C1278 The *specification-part* of a pure subprogram shall specify that all its dummy procedures are
26 pure.
- 27 C1279 If a procedure that is neither an intrinsic procedure nor a statement function is used in a context
28 that requires it to be pure, then its interface shall be explicit in the scope of that use. The
29 interface shall specify that the procedure is pure.
- 30 C1280 All internal subprograms in a pure subprogram shall be pure.
- 31 C1281 In a pure subprogram any designator with a base object that is in common or accessed by
32 host or use association, is a dummy argument of a pure function, is a dummy argument with
33 INTENT (IN) of a pure subroutine, or an object that is storage associated with any such variable,
34 shall not be used
- 35 (1) in a variable definition context(16.6.7),
36 (2) as the *data-target* in a *pointer-assignment-stmt*,

- 1 (3) as the *expr* corresponding to a component with the POINTER attribute in a *structure-*
 2 *constructor*,
 3 (4) as the *expr* of an intrinsic assignment statement in which the variable is of a derived type
 4 if the derived type has a pointer component at any level of component selection, or
 5 (5) as an actual argument associated with a dummy argument with INTENT (OUT) or IN-
 6 TENT (INOUT) or with the POINTER attribute.

NOTE 12.51

Item 3 requires that processors be able to determine if entities with the PRIVATE attribute or with private components have a pointer component.

- 7 C1282 Any procedure referenced in a pure subprogram, including one referenced via a defined operation,
 8 defined assignment, user-defined derived-type input/output, or finalization, shall be pure.
- 9 C1283 A pure subprogram shall not contain a *print-stmt*, *open-stmt*, *close-stmt*, *backspace-stmt*, *endfile-*
 10 *stmt*, *rewind-stmt*, *flush-stmt*, *wait-stmt*, or *inquire-stmt*.
- 11 C1284 A pure subprogram shall not contain a *read-stmt* or *write-stmt* whose *io-unit* is a *file-unit-*
 12 *number* or *.
- 13 C1285 A pure subprogram shall not contain a *stop-stmt*.
- 14 C1286 A co-indexed object shall not appear in a variable definition context in a pure subprogram.
- 15 C1287 A pure subprogram shall not contain an image control statement (8.5.1).

NOTE 12.52

The above constraints are designed to guarantee that a pure procedure is free from side effects (modifications of data visible outside the procedure), which means that it is safe to reference it in constructs such as a FORALL *assignment-stmt* or a DO CONCURRENT construct, where there is no explicit order of evaluation.

The constraints on pure subprograms may appear complicated, but it is not necessary for a programmer to be intimately familiar with them. From the programmer's point of view, these constraints can be summarized as follows: a pure subprogram shall not contain any operation that could conceivably result in an assignment or pointer assignment to a common variable, a variable accessed by use or host association, or an INTENT (IN) dummy argument; nor shall a pure subprogram contain any operation that could conceivably perform any external file input/output or STOP operation. Note the use of the word conceivably; it is not sufficient for a pure subprogram merely to be side-effect free in practice. For example, a function that contains an assignment to a global variable but in a block that is not executed in any invocation of the function is nevertheless not a pure function. The exclusion of functions of this nature is required if strict compile-time checking is to be used.

It is expected that most library procedures will conform to the constraints required of pure procedures, and so can be declared pure and referenced in FORALL statements and constructs, DO CONCURRENT constructs, and within user-defined pure procedures.

NOTE 12.53

Pure subroutines are included to allow subroutine calls from pure procedures in a safe way, and to allow *forall-assignment-stmts* to be defined assignments. The constraints for pure subroutines are based on the same principles as for pure functions, except that side effects to INTENT (OUT), INTENT (INOUT), and pointer dummy arguments are permitted.

1 12.8 Elemental procedures

2 12.8.1 Elemental procedure declaration and interface

3 An **elemental procedure** is an elemental intrinsic procedure or a procedure that is defined by an
4 elemental subprogram.

5 An elemental subprogram has the *prefix-spec* ELEMENTAL. An elemental subprogram is a pure sub-
6 program unless it has the *prefix-spec* IMPURE. The following additional constraints apply to elemental
7 subprograms.

8 C1288 All dummy arguments of an elemental procedure shall be scalar dummy data objects and shall
9 not have the POINTER or ALLOCATABLE attribute.

10 C1289 The result variable of an elemental function shall be scalar and shall not have the POINTER or
11 ALLOCATABLE attribute.

12 C1290 In the scoping unit of an elemental subprogram, an object designator with a dummy argument
13 as the base object shall not appear in a *specification-expr* except as the argument to one of the
14 intrinsic functions BIT_SIZE, KIND, LEN, or the numeric inquiry functions (13.5.6).

NOTE 12.54

The restriction on dummy arguments in specification expressions is imposed primarily to facilitate optimization. An example of usage that is not permitted is

```
ELEMENTAL REAL FUNCTION F (A, N)
  REAL, INTENT (IN) :: A
  INTEGER, INTENT (IN) :: N
  REAL :: WORK_ARRAY(N) ! Invalid
  ...
END FUNCTION F
```

An example of usage that is permitted is

```
ELEMENTAL REAL FUNCTION F (A)
  REAL, INTENT (IN) :: A
  REAL (SELECTED_REAL_KIND (PRECISION (A)*2)) :: WORK
  ...
END FUNCTION F
```

15 12.8.2 Elemental function actual arguments and results

16 If a generic name or a specific name is used to reference an elemental function, the shape of the result is
17 the same as the shape of the actual argument with the greatest rank. If there are no actual arguments
18 or the actual arguments are all scalar, the result is scalar. For those elemental functions that have more
19 than one argument, all actual arguments shall be conformable. In the array case, the values of the
20 elements, if any, of the result are the same as would have been obtained if the scalar function had been
21 applied separately, in array element order, to corresponding elements of each array actual argument.

NOTE 12.55

An example of an elemental reference to the intrinsic function MAX:

if X and Y are arrays of shape (M, N),

NOTE 12.55 (cont.)

```
MAX (X, 0.0, Y)
```

is an array expression of shape (M, N) whose elements have values

```
MAX (X(I, J), 0.0, Y(I, J)), I = 1, 2, ..., M, J = 1, 2, ..., N
```

1 12.8.3 Elemental subroutine actual arguments

- 2 An elemental subroutine is one that has only scalar dummy arguments, but may have array actual
3 arguments. In a reference to an elemental subroutine, either all actual arguments shall be scalar, or
4 all actual arguments associated with INTENT (OUT) and INTENT (INOUT) dummy arguments shall
5 be arrays of the same shape and the remaining actual arguments shall be conformable with them. In
6 the case that the actual arguments associated with INTENT (OUT) and INTENT (INOUT) dummy
7 arguments are arrays, the values of the elements, if any, of the results are the same as would be obtained
8 if the subroutine had been applied separately, in array element order, to corresponding elements of each
9 array actual argument.
- 10 In a reference to the intrinsic subroutine MVBITS, the actual arguments corresponding to the TO and
11 FROM dummy arguments may be the same variable and may be associated scalar variables or associated
12 array variables all of whose corresponding elements are associated. Apart from this, the actual arguments
13 in a reference to an elemental subroutine must satisfy the restrictions of 12.5.2.14.

1 13 Intrinsic procedures and modules

2 13.1 Classes of intrinsic procedures

3 There are four classes of intrinsic procedures: inquiry functions, elemental functions, transformational
4 functions, and subroutines. Some intrinsic subroutines are elemental. Some intrinsic subroutines are
5 collective.

6 An **inquiry function** is one whose result depends on the properties of one or more of its arguments
7 instead of their values; in fact, these argument values may be undefined. Unless the description of
8 an inquiry function states otherwise, these arguments are permitted to be unallocated allocatables or
9 pointers that are not associated. An **elemental intrinsic function** is one that is specified for scalar
10 arguments, but may be applied to array arguments as described in 12.8. All other intrinsic functions
11 are **transformational functions**; they almost all have one or more array arguments or an array result.
12 All standard intrinsic functions are pure.

13 The subroutine MOVE_ALLOC and the elemental subroutine MVBITS are pure. No other standard
14 intrinsic subroutine is pure.

15 A **collective subroutine** is one of the intrinsic subroutines named in 13.5.15. If it is invoked by one
16 image of a team, it shall be invoked by the same statement on all images of the team. There is an implicit
17 team synchronization (8.5.3) at the beginning and end of the execution of a collective subroutine.

NOTE 13.1

As with user-written elemental subroutines, an elemental intrinsic subroutine is pure. The effects of MOVE_ALLOC are limited to its arguments. The remaining nonelemental intrinsic subroutines all have side effects (or reflect system side effects) and thus are not pure.

18 **Generic names** of standard intrinsic procedures are listed in 13.5. In most cases, generic functions
19 accept arguments of more than one type and the type of the result is the same as the type of the
20 arguments. **Specific names** of standard intrinsic functions with corresponding generic names are listed
21 in 13.6.

22 If an intrinsic procedure is used as an actual argument to a procedure, its specific name shall be used
23 and it may be referenced in the called procedure only with scalar arguments. If an intrinsic procedure
24 does not have a specific name, it shall not be used as an actual argument (12.5.2.10).

25 Elemental intrinsic procedures behave as described in 12.8.

26 13.2 Arguments to intrinsic procedures

27 13.2.1 General rules

28 All intrinsic procedures may be invoked with either positional arguments or argument keywords (12.5).
29 The descriptions in 13.5 through 13.7 give the argument keyword names and positional sequence for
30 standard intrinsic procedures.

31 Many of the intrinsic procedures have optional arguments. These arguments are identified by the notation
32 “optional” in the argument descriptions. In addition, the names of the optional arguments are enclosed
33 in square brackets in description headings and in lists of procedures. The valid forms of reference for

1 procedures with optional arguments are described in 12.5.2.

NOTE 13.2

The text CMPLX (X [, Y, KIND]) indicates that Y and KIND are both optional arguments. Valid reference forms include CMPLX(*x*), CMPLX(*x*, *y*), CMPLX(*x*, KIND=*kind*), CMPLX(*x*, *y*, *kind*), and CMPLX(KIND=*kind*, X=*x*, Y=*y*).

NOTE 13.3

Some intrinsic procedures impose additional requirements on their optional arguments. For example, SELECTED_REAL_KIND requires that at least one of its optional arguments be present, and RANDOM_SEED requires that at most one of its optional arguments be present.

2 The dummy arguments of the specific intrinsic procedures in 13.6 have INTENT(IN). The dummy
3 arguments of the generic intrinsic procedures in 13.7 have INTENT(IN) if the intent is not stated
4 explicitly.

5 The actual argument associated with an intrinsic function dummy argument named KIND shall be a
6 scalar integer initialization expression and its value shall specify a representation method for the function
7 result that exists on the processor.

8 Intrinsic subroutines that assign values to arguments of type character do so in accordance with the
9 rules of intrinsic assignment (7.4.1.3).

10 13.2.2 The shape of array arguments

11 Unless otherwise specified, the inquiry intrinsic functions accept array arguments for which the shape
12 need not be defined. The shape of array arguments to transformational and elemental intrinsic functions
13 shall be defined.

14 13.2.3 Mask arguments

15 Some array intrinsic functions have an optional MASK argument of type logical that is used by the
16 function to select the elements of one or more arguments to be operated on by the function. Any
17 element not selected by the mask need not be defined at the time the function is invoked.

18 The MASK affects only the value of the function, and does not affect the evaluation, prior to invoking
19 the function, of arguments that are array expressions.

20 13.2.4 Arguments to collective subroutines

21 Each argument to a collective subroutine shall have the same shape on all images of the team. A co-array
22 argument shall have the same bounds on all images of the team. An INTENT(IN) argument of type
23 IMAGE_TEAM shall have a value that was constructed by corresponding invocations of FORM_TEAM
24 for the team.

25 13.3 Bit model

26 The bit manipulation procedures and inquiry functions are described in terms of a model for the repre-
27 sentation and behaviour of bits on a processor.

28 For an integer, a bit is defined to be a binary digit *w* located at position *k* of a nonnegative integer scalar
29 object based on a model nonnegative integer defined by

$$j = \sum_{k=0}^{z-1} w_k \times 2^k$$

1 and for which w_k may have the value 0 or 1. An example of a model number compatible with the
2 examples used in 13.4 would have $z = 32$, thereby defining a 32-bit integer.

3 Using the notation of the formula above, the value of an object of type bits and kind z is represented as
4 the ordered sequence of bits with w_k the bit at position k . The rightmost bit is w_0 and the leftmost bit
5 is w_{z-1} . Such a bits object can be interpreted as a nonnegative integer with the value j .

6 An inquiry function BIT_SIZE is available to determine the parameter z of the model.

7 Effectively, this model defines an integer object to consist of z bits in sequence numbered from right
8 to left from 0 to $z - 1$. This model is valid only in the context of the use of such an object as the
9 argument or result of one of the bit manipulation procedures. In all other contexts, the model defined
10 for an integer in 13.4 applies. In particular, whereas the models are identical for $r = 2$ and $w_{z-1} = 0$,
11 they do not correspond for $r \neq 2$ or $w_{z-1} = 1$ and the interpretation of bits in such objects is processor
12 dependent.

13 13.4 Numeric models

14 The numeric manipulation and inquiry functions are described in terms of a model for the representation
15 and behavior of numbers on a processor. The model has parameters that are determined so as to make
16 the model best fit the machine on which the program is executed.

17 The model set for integer i is defined by

$$i = s \times \sum_{k=0}^{q-1} w_k \times r^k$$

18 where r is an integer exceeding one, q is a positive integer, each w_k is a nonnegative integer less than r ,
19 and s is +1 or -1.

20 The model set for real x is defined by

$$x = \begin{cases} 0 \text{ or} \\ s \times b^e \times \sum_{k=1}^p f_k \times b^{-k} \end{cases},$$

21 where b and p are integers exceeding one; each f_k is a nonnegative integer less than b , with f_1 nonzero; s
22 is +1 or -1; and e is an integer that lies between some integer maximum e_{\max} and some integer minimum
23 e_{\min} inclusively. For $x = 0$, its exponent e and digits f_k are defined to be zero. The integer parameters
24 r and q determine the set of model integers and the integer parameters b , p , e_{\min} , and e_{\max} determine
25 the set of model floating point numbers. The parameters of the integer and real models are available
26 for each representation method of the integer and real types. The parameters characterize the set of
27 available numbers in the definition of the model. The floating-point manipulation functions (13.5.10)
28 and numeric inquiry functions (13.5.6) provide values of some parameters and other values related to
29 the models.

NOTE 13.4

Examples of these functions in 13.7 use the models

$$i = s \times \sum_{k=0}^{30} w_k \times 2^k$$

and

$$x = 0 \text{ or } s \times 2^e \times \left(\frac{1}{2} + \sum_{k=2}^{24} f_k \times 2^{-k} \right), \quad -126 \leq e \leq 127$$

1 13.5 Standard generic intrinsic procedures

- 2 For all of the standard intrinsic procedures, the arguments shown are the names that shall be used for
3 argument keywords if the keyword form is used for actual arguments.

NOTE 13.5

For example, a reference to CMPLX may be written in the form CMPLX (A, B, M) or in the form CMPLX (Y = B, KIND = M, X = A).

NOTE 13.6

Many of the argument keywords have names that are indicative of their usage. For example:

KIND	Describes the kind type parameter of the result
STRING, STRING_A	An arbitrary character string
BACK	Controls the direction of string scan (forward or backward)
MASK	A mask that may be applied to the arguments
DIM	A selected dimension of an array argument

4 13.5.1 Numeric functions

5	ABS (A)	Absolute value
6	AIMAG (Z)	Imaginary part of a complex number
7	AINIT (A [, KIND])	Truncation to whole number
8	ANINT (A [, KIND])	Nearest whole number
9	CEILING (A [, KIND])	Least integer greater than or equal to number
10	CMPLX (X [, Y, KIND])	Conversion to complex type
11	CONJG (Z)	Conjugate of a complex number
12	DBLE (A)	Conversion to double precision real type
13	DIM (X, Y)	Positive difference
14	DPROD (X, Y)	Double precision real product
15	FLOOR (A [, KIND])	Greatest integer less than or equal to number
16	INT (A [, KIND])	Conversion to integer type
17	MAX (A1, A2 [, A3,...])	Maximum value
18	MIN (A1, A2 [, A3,...])	Minimum value
19	MOD (A, P)	Remainder function
20	MODULO (A, P)	Modulo function

1	NINT (A [, KIND])	Nearest integer
2	REAL (A [, KIND])	Conversion to real type
3	SIGN (A, B)	Transfer of sign

4 13.5.2 Mathematical functions

5	ACOS (X)	Arccosine
6	ACOSH (X)	Hyperbolic arccosine
7	ASIN (X)	Arcsine
8	ASINH (X)	Hyperbolic arcsine
9	ATAN (X) or ATAN (Y, X)	Arctangent
10	ATAN2 (Y, X)	Arctangent
11	ATANH (X)	Hyperbolic arctangent
12	BESSEL_J0 (X)	Bessel function of the first kind of order zero
13	BESSEL_J1 (X)	Bessel function of the first kind of order one
14	BESSEL_JN (N,X)	Bessel function of the first kind of order N
15	BESSEL_Y0 (X)	Bessel function of the second kind of order zero
16	BESSEL_Y1 (X)	Bessel function of the second kind of order one
17	BESSEL_YN (N, X)	Bessel function of the second kind of order N
18	COS (X)	Cosine
19	COSH (X)	Hyperbolic cosine
20	ERF (X)	Error function
21	ERFC (X)	Complementary error function
22	ERFC_SCALED (X)	Exponentially-scaled complementary error function
23		
24	EXP (X)	Exponential
25	GAMMA (X)	Gamma function
26	HYPOT (X, Y)	Euclidean distance function
27	LOG (X)	Natural logarithm
28	LOG.GAMMA (X)	Logarithm of absolute value of gamma function
29	LOG10 (X)	Common logarithm (base 10)
30	SIN (X)	Sine
31	SINH (X)	Hyperbolic sine
32	SQRT (X)	Square root
33	TAN (X)	Tangent
34	TANH (X)	Hyperbolic tangent

35 13.5.3 Character functions

36	ACHAR (I [, KIND])	Character in given position in ASCII collating sequence
37		
38	ADJUSTL (STRING)	Adjust left
39	ADJUSTR (STRING)	Adjust right
40	CHAR (I [, KIND])	Character in given position in processor collating sequence
41		
42	IACHAR (C [, KIND])	Position of a character in ASCII collating sequence
43		
44	ICHAR (C [, KIND])	Position of a character in processor collating sequence
45		
46	INDEX (STRING, SUBSTRING [, BACK, KIND])	Starting position of a substring
47	LEN_TRIM (STRING [, KIND])	Length without trailing blank characters
48	LGE (STRING_A, STRING_B)	Lexically greater than or equal
49	LGT (STRING_A, STRING_B)	Lexically greater than

1	LLE (STRING_A, STRING_B)	Lexically less than or equal
2	LLT (STRING_A, STRING_B)	Lexically less than
3	MAX (A1, A2 [, A3,...])	Maximum value
4	MIN (A1, A2 [, A3,...])	Minimum value
5	REPEAT (STRING, NCOPIES)	Repeated concatenation
6	SCAN (STRING, SET [, BACK, KIND])	Scan a string for a character in a set
7	TRIM (STRING)	Remove trailing blank characters
8	VERIFY (STRING, SET [, BACK, KIND])	Verify the set of characters in a string
9	13.5.4 Kind functions	
10	BITS_KIND (X)	Bits kind type parameter value, compatible with the argument
11	KIND (X)	Kind type parameter value
12	SELECTED_BITS_KIND (N)	Bits kind type parameter value, given number of bits
13	SELECTED_CHAR_KIND (NAME)	Character kind type parameter value, given character set name
14	SELECTED_CHAR_KIND (NAME)	Character kind type parameter value, given character set name
15	SELECTED_CHAR_KIND (NAME)	Character kind type parameter value, given character set name
16	SELECTED_CHAR_KIND (NAME)	Character kind type parameter value, given character set name
17	SELECTED_INT_KIND (R)	Integer kind type parameter value, given range
18	SELECTED_REAL_KIND ([P, R, RADIX])	Real kind type parameter value, given precision, range, and radix
19		
20	13.5.5 Miscellaneous type conversion functions	
21	BITS (A [, KIND])	Convert to bits type
22	LOGICAL (L [, KIND])	Convert between objects of type logical with different kind type parameters
23		
24	TRANSFER (SOURCE, MOLD [, SIZE])	Treat first argument as if of type of second argument
25		
26	13.5.6 Numeric inquiry functions	
27	DIGITS (X)	Number of significant digits of the model
28	EPSILON (X)	Number that is almost negligible compared to one
29		
30	HUGE (X)	Largest number of the model
31	MAXEXPONENT (X)	Maximum exponent of the model
32	MINEXPONENT (X)	Minimum exponent of the model
33	PRECISION (X)	Decimal precision
34	RADIX (X)	Base of the model
35	RANGE (X)	Decimal exponent range
36	TINY (X)	Smallest positive number of the model
37	13.5.7 Array inquiry functions	
38	CO_LBOUND (CO_ARRAY [, DIM, KIND])	Lower co-bounds of a co-array
39	CO_UBOUND (CO_ARRAY [, DIM, KIND])	Upper co-bounds of a co-array
40	LBOUND (ARRAY [, DIM, KIND])	Lower dimension bounds of an array
41	SHAPE (SOURCE [, KIND])	Shape of an array or scalar
42	SIZE (ARRAY [, DIM, KIND])	Total number of elements in an array
43	UBOUND (ARRAY [, DIM, KIND])	Upper dimension bounds of an array

1 13.5.8 Other inquiry functions

2	ALLOCATED (ARRAY) or	Allocation status
3	ALLOCATED (SCALAR)	
4	ASSOCIATED (POINTER [, TARGET])	Association status inquiry or comparison
5	BIT_SIZE (I)	Number of bits of the model
6	EXTENDS_TYPE_OF (A, MOLD)	Same dynamic type or an extension
7	IS_CONTIGUOUS	Contiguity inquiry
8	LEN (STRING [, KIND])	Length of a character entity
9	NEW_LINE (A)	Newline character
10	PRESENT (A)	Argument presence
11	SAME_TYPE_AS (A, B)	Same dynamic type
12	STORAGE_SIZE (A [, KIND])	Size in bits of an array element

12 13.5.9 Bit manipulation procedures

13	BTEST (I, POS)	Bit testing
14	DSHIFTL (I, J, SHIFT)	Combined left shift
15	DSHIFTR (I, J, SHIFT)	Combined right shift
16	IAND (I, J)	Bitwise AND
17	IBCLR (I, POS)	Clear bit
18	IBITS (I, POS, LEN)	Bit extraction
19	IBSET (I, POS)	Set bit
20	IEOR (I, J)	Exclusive OR
21	IOR (I, J)	Inclusive OR
22	ISHFT (I, SHIFT)	Logical shift
23	ISHFTC (I, SHIFT [, SIZE])	Circular shift
24	LEADZ (I)	Number of leading zero bits
25	MASKL (I [, KIND])	Left justified bit mask
26	MASKR (I [, KIND])	Right justified bit mask
27	MERGE_BITS (I, J, MASK)	Merge bits under mask
28	MVBITS (FROM, FROMPOS, LEN, TO, TOPOS)	Copies bits from one integer to another
29	NOT (I)	Bitwise complement
30	POPCNT (I)	Number of one bits
31	POPPAR (I)	Parity of one bits
32	SHIFTA (I, SHIFT)	Right shift with fill
33	SHIFTL (I, SHIFT)	Left shift
34	SHIFTR (I, SHIFT)	Right shift
35	TRAILZ (I)	Number of trailing zero bits

36 13.5.10 Floating-point manipulation functions

37	EXPONENT (X)	Exponent part of a model number
38	FRACTION (X)	Fractional part of a number
39	NEAREST (X, S)	Nearest different processor number in given direction
40		
41	RRSPACING (X)	Reciprocal of the relative spacing of model numbers near given number
42		
43	SCALE (X, I)	Multiply a real by its base to an integer power
44	SET_EXPONENT (X, I)	Set exponent part of a number
45	SPACING (X)	Absolute spacing of model numbers near given number
46		

1 13.5.11 Vector and matrix multiply functions

2	DOT_PRODUCT (VECTOR_A, VECTOR_B)	Dot product of two rank-one arrays
3	MATMUL (MATRIX_A, MATRIX_B)	Matrix multiplication

4 13.5.12 Array reduction functions

5	ALL (MASK [, DIM])	True if all values are true
6	ANY (MASK [, DIM])	True if any value is true
7	COUNT (MASK [, DIM, KIND])	Number of true elements in an array
8	IALL (ARRAY, DIM [, MASK]) or IALL (ARRAY [, MASK])	Bitwise AND of array elements
9	IANY (ARRAY, DIM [, MASK]) or IANY (ARRAY [, MASK])	Bitwise OR of array elements
10	IPARITY (ARRAY, DIM [, MASK]) or IPARITY (ARRAY [, MASK])	Bitwise exclusive OR of array elements
11	MAXVAL (ARRAY, DIM [, MASK]) or MAXVAL (ARRAY [, MASK])	Maximum value in an array
12	MINVAL (ARRAY, DIM [, MASK]) or MINVAL (ARRAY [, MASK])	Minimum value in an array
13	NORM2 (X)	L_2 norm of an array
14	PARITY (MASK [, DIM])	True if an odd number of values is true
15	PRODUCT (ARRAY, DIM [, MASK]) or PRODUCT (ARRAY [, MASK])	Product of array elements
16	SUM (ARRAY, DIM [, MASK]) or SUM (ARRAY [, MASK])	Sum of array elements

17 13.5.13 Array construction functions

18	CSHIFT (ARRAY, SHIFT [, DIM])	Circular shift
19	EOSHIFT (ARRAY, SHIFT [, BOUNDARY, DIM])	End-off shift
20	MERGE (TSOURCE, FSOURCE, MASK)	Merge under mask
21	PACK (ARRAY, MASK [, VECTOR])	Pack an array into an array of rank one under a mask
22	RESHAPE (SOURCE, SHAPE[, PAD, ORDER])	Reshape an array
23	SPREAD (SOURCE, DIM, NCOPIES)	Replicates array by adding a dimension
24	TRANSPOSE (MATRIX)	Transpose of an array of rank two
25	UNPACK (VECTOR, MASK, FIELD)	Unpack an array of rank one into an array under a mask

28 13.5.14 Array location functions

29	FINDLOC (ARRAY, VALUE, DIM, [, MASK, KIND, BACK]) or FINDLOC (ARRAY, VALUE, [, MASK, KIND, BACK])	Location of value in an array
30	MAXLOC (ARRAY, DIM [, MASK, KIND, BACK]) or MAXLOC (ARRAY [, MASK, KIND, BACK])	Location of a maximum value in an array
31	MINLOC (ARRAY, DIM [, MASK, KIND, BACK]) or MINLOC (ARRAY [, MASK, KIND, BACK])	Location of a minimum value in an array

1 13.5.15 Collective subroutines

2	CO_ALL (MASK, RESULT [, TEAM])	Whether all values are true
3	CO_ANY (MASK, RESULT [, TEAM])	Whether any value is true
4	CO_COUNT (MASK, RESULT [, TEAM])	Number of true elements in a co-array
5	CO_MAXLOC (CO_ARRAY, RESULT [, TEAM])	Image indices of maximum values
6	CO_MAXVAL (CO_ARRAY, RESULT [, TEAM])	Maximum values
7	CO_MINLOC (CO_ARRAY, RESULT [, TEAM])	Image indices of minimum values
8	CO_MINVAL (CO_ARRAY, RESULT [, TEAM])	Minimum values
9	CO_PRODUCT (CO_ARRAY, RESULT [, TEAM])	Products of elements
10	CO_SUM (CO_ARRAY, RESULT [, TEAM])	Sums of elements
11	FORM_TEAM (TEAM, IMAGES)	Forms a team of images

12 13.5.16 Null function

13	NULL ([MOLD])	Returns disassociated or unallocated result
----	---------------	---

14 13.5.17 Allocation transfer procedure

15	MOVE_ALLOC (FROM, TO)	Moves an allocation from one allocatable object to another
----	-----------------------	--

17 13.5.18 Random number subroutines

18	RANDOM_NUMBER (HARVEST)	Returns pseudorandom number
19	RANDOM_SEED ([SIZE, PUT, GET])	Initializes or restarts the pseudorandom number generator

21 13.5.19 System environment procedures

22	COMMAND_ARGUMENT_COUNT ()	Number of command arguments
23	CPU_TIME (TIME)	Obtain processor time
24	DATE_AND_TIME ([DATE, TIME, ZONE, VALUES])	Obtain date and time
25	EXECUTE_COMMAND_LINE (COMMAND, [WAIT, EXITSTAT, CMDSTAT, CMDMSG])	Execute a command line
26	GET_COMMAND ([COMMAND, LENGTH, STATUS])	Returns entire command
27	GET_COMMAND_ARGUMENT (NUMBER [, VALUE, LENGTH, STATUS])	Returns a command argument
28	GET_ENVIRONMENT_VARIABLE (NAME [, VALUE, LENGTH, STATUS, TRIM_NAME])	Obtain the value of an environment variable
29	IMAGE_INDEX (CO_ARRAY, SUB)	Index of an image
30	IS_IOSTAT_END (I)	Test for end-of-file value
31	IS_IOSTAT_EOR (I)	Test for end-of-record value
32	NUM_IMAGES ()	Number of images
	SYSTEM_CLOCK ([COUNT, COUNT_RATE, COUNT_MAX])	Obtain data from the system clock

1		
2	TEAM_IMAGES (TEAM)	Indices of the images in a team
3	THIS_IMAGE ()	Index of the invoking image
4	THIS_IMAGE (CO_ARRAY [, DIM])	Co-subscript list

5 13.6 Specific names for standard intrinsic functions

6 Except for AMAX0, AMIN0, MAX1, and MIN1, the result type of the specific function is the same as the
 7 result type of the corresponding generic function would be if it were invoked with the same arguments
 8 as the specific function.

Specific Name	Generic Name	Argument Type
ABS	ABS	default real
ACOS	ACOS	default real
AIMAG	AIMAG	default complex
AINT	AINT	default real
ALOG	LOG	default real
ALOG10	LOG10	default real
• AMAX0 (...)	REAL (MAX (...))	default integer
• AMAX1	MAX	default real
• AMIN0 (...)	REAL (MIN (...))	default integer
• AMIN1	MIN	default real
AMOD	MOD	default real
ANINT	ANINT	default real
ASIN	ASIN	default real
ATAN (X)	ATAN	default real
ATAN2	ATAN2	default real
CABS	ABS	default complex
CCOS	COS	default complex
CEXP	EXP	default complex
• CHAR	CHAR	default integer
CLOG	LOG	default complex
CONJG	CONJG	default complex
COS	COS	default real
COSH	COSH	default real
CSIN	SIN	default complex
CSQRT	SQRT	default complex
DABS	ABS	double precision real
DACOS	ACOS	double precision real
DASIN	ASIN	double precision real
DATAN	ATAN	double precision real
DATAN2	ATAN2	double precision real
DCOS	COS	double precision real
DCOSH	COSH	double precision real
DDIM	DIM	double precision real
DEXP	EXP	double precision real
DIM	DIM	default real
DINT	AINT	double precision real
DLOG	LOG	double precision real
DLOG10	LOG10	double precision real
• DMAX1	MAX	double precision real
• DMIN1	MIN	double precision real
DMOD	MOD	double precision real

Specific Name	Generic Name	Argument Type
DNINT	ANINT	double precision real
DPROD	DPROD	default real
DSIGN	SIGN	double precision real
DSIN	SIN	double precision real
DSINH	SINH	double precision real
DSQRT	SQRT	double precision real
DTAN	TAN	double precision real
DTANH	TANH	double precision real
EXP	EXP	default real
• FLOAT	REAL	default integer
IABS	ABS	default integer
• ICHAR	ICHAR	default character
IDIM	DIM	default integer
• IDINT	INT	double precision real
IDNINT	NINT	double precision real
• IFIX	INT	default real
INDEX	INDEX	default character
• INT	INT	default real
ISIGN	SIGN	default integer
LEN	LEN	default character
• LGE	LGE	default character
• LGT	LGT	default character
• LLE	LLE	default character
• LLT	LLT	default character
• MAX0	MAX	default integer
• MAX1 (...)	INT (MAX (...))	default real
• MIN0	MIN	default integer
• MIN1 (...)	INT (MIN (...))	default real
MOD	MOD	default integer
NINT	NINT	default real
• REAL	REAL	default integer
SIGN	SIGN	default real
SIN	SIN	default real
SINH	SINH	default real
• SNGL	REAL	double precision real
SQRT	SQRT	default real
TAN	TAN	default real
TANH	TANH	default real

1 A specific intrinsic function marked with a bullet (•) shall not be used as an actual argument or as a
2 target in a procedure pointer assignment statement.

3 13.7 Specifications of the standard intrinsic procedures

4 Detailed specifications of the standard generic intrinsic procedures are provided here in alphabetical
5 order.

6 The types and type parameters of standard intrinsic procedure arguments and function results are de-
7 termined by these specifications. The “Argument(s)” paragraphs specify requirements on the actual
8 arguments of the procedures. The result characteristics are sometimes specified in terms of the char-
9 acteristics of dummy arguments. A program is prohibited from invoking an intrinsic procedure under

1 circumstances where a value to be returned in a subroutine argument or function result is outside the
2 range of values representable by objects of the specified type and type parameters, unless the intrinsic
3 module IEEE_ARITHMETIC (clause 14) is accessible and there is support for an infinite or a NaN
4 result, as appropriate. If an infinite result is returned, the flag IEEE_OVERFLOW or IEEE_DIVIDE_
5 BY_ZERO shall signal; if a NaN result is returned, the flag IEEE_INVALID shall signal. If all results
6 are normal, these flags shall have the same status as when the intrinsic procedure was invoked.

7 13.7.1 ABS (A)

8 **Description.** Absolute value.

9 **Class.** Elemental function.

10 **Argument.** A shall be of type integer, real, or complex.

11 **Result Characteristics.** The same as A except that if A is complex, the result is real.

12 **Result Value.** If A is of type integer or real, the value of the result is $|A|$; if A is complex with
13 value (x, y) , the result is equal to a processor-dependent approximation to $\sqrt{x^2 + y^2}$ computed
14 without undue overflow or underflow.

15 **Example.** ABS ((3.0, 4.0)) has the value 5.0 (approximately).

16 13.7.2 ACHAR (I [, KIND])

17 **Description.** Returns the character in a specified position of the ASCII collating sequence. It is
18 the inverse of the IACHAR function.

19 **Class.** Elemental function.

20 **Arguments.**

21 I shall be of type integer or bits. If it is of type bits, it is interpreted as a
nonnegative integer as described in 13.3.

22 KIND (optional) shall be a scalar integer initialization expression.

23 **Result Characteristics.** Character of length one. If KIND is present, the kind type parameter
24 is that specified by the value of KIND; otherwise, the kind type parameter is that of default
25 character type.

26 **Result Value.** If I has a value in the range $0 \leq I \leq 127$, the result is the character in
27 position I of the ASCII collating sequence, provided the processor is capable of representing
28 that character in the character type of the result; otherwise, the result is processor dependent.
29 ACHAR (IACHAR (C)) shall have the value C for any character C capable of representation in
30 the default character type.

31 **Examples.** ACHAR (88) has the value 'X'. ACHAR (Z'41') has the value 'A'.

32 13.7.3 ACOS (X)

33 **Description.** Arccosine (inverse cosine) function.

34 **Class.** Elemental function.

35 **Argument.** X shall be of type real with a value that satisfies the inequality $|X| \leq 1$, or of type
36 complex.

1 **Result Characteristics.** Same as X.

2 **Result Value.** The result has a value equal to a processor-dependent approximation to $\arccos(X)$.
 3 If it is real it is expressed in radians and lies in the range $0 \leq \text{ACOS}(X) \leq \pi$. If it is complex the
 4 real part is expressed in radians and lies in the range $0 \leq \text{REAL}(\text{ACOS}(X)) \leq \pi$.

5 **Example.** $\text{ACOS}(0.54030231)$ has the value 1.0 (approximately).

6 **13.7.4 ADJUSTL (STRING)**

7 **Description.** Adjust to the left, removing leading blanks and inserting trailing blanks.

8 **Class.** Elemental function.

9 **Argument.** STRING shall be of type character.

10 **Result Characteristics.** Character of the same length and kind type parameter as STRING.

11 **Result Value.** The value of the result is the same as STRING except that any leading blanks
 12 have been deleted and the same number of trailing blanks have been inserted.

13 **Example.** $\text{ADJUSTL}(' \text{WORD}')$ has the value 'WORD '.

14 **13.7.5 ACOSH (X)**

15 **Description.** Inverse hyperbolic cosine function.

16 **Class.** Elemental function.

17 **Argument.** X shall be of type real or complex.

18 **Result Characteristics.** Same as X.

19 **Result Value.** The result has a value equal to a processor-dependent approximation to the
 20 inverse hyperbolic cosine function of X. If the result is complex the imaginary part is expressed
 21 in radians and lies in the range $0 \leq \text{AIMAG}(\text{ACOSH}(X)) \leq \pi$

22 **Example.** $\text{ACOSH}(1.5430806)$ has the value 1.0 (approximately).

23 **13.7.6 ADJUSTR (STRING)**

24 **Description.** Adjust to the right, removing trailing blanks and inserting leading blanks.

25 **Class.** Elemental function.

26 **Argument.** STRING shall be of type character.

27 **Result Characteristics.** Character of the same length and kind type parameter as STRING.

28 **Result Value.** The value of the result is the same as STRING except that any trailing blanks
 29 have been deleted and the same number of leading blanks have been inserted.

30 **Example.** $\text{ADJUSTR}(' \text{WORD}')$ has the value ' WORD'.

31 **13.7.7 AIMAG (Z)**

32 **Description.** Imaginary part of a complex number.

33 **Class.** Elemental function.

1 **Argument.** Z shall be of type complex.

2 **Result Characteristics.** Real with the same kind type parameter as Z.

3 **Result Value.** If Z has the value (x, y) , the result has the value y .

4 **Example.** AIMAG ((2.0, 3.0)) has the value 3.0.

5 13.7.8 AINT (A [, KIND])

6 **Description.** Truncation to a whole number.

7 **Class.** Elemental function.

8 **Arguments.**

9 A shall be of type real.

10 KIND (optional) shall be a scalar integer initialization expression.

11 **Result Characteristics.** The result is of type real. If KIND is present, the kind type parameter is that specified by the value of KIND; otherwise, the kind type parameter is that of A.

12 **Result Value.** If $|A| < 1$, AINT (A) has the value 0; if $|A| \geq 1$, AINT (A) has a value equal to the integer whose magnitude is the largest integer that does not exceed the magnitude of A and whose sign is the same as the sign of A.

13 **Examples.** AINT (2.783) has the value 2.0. AINT (-2.783) has the value -2.0.

14 13.7.9 ALL (MASK [, DIM])

15 **Description.** Determine whether all values are true in MASK along dimension DIM.

16 **Class.** Transformational function.

17 **Arguments.**

18 MASK shall be of type logical. It shall be an array.

19 DIM (optional) shall be scalar and of type integer with value in the range $1 \leq \text{DIM} \leq n$, where n is the rank of MASK. The corresponding actual argument shall not be an optional dummy argument.

20 **Result Characteristics.** The result is of type logical with the same kind type parameter as MASK. It is scalar if DIM is absent; otherwise, the result has rank $n - 1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of MASK.

21 **Result Value.**

22 *Case (i):* The result of ALL (MASK) has the value true if all elements of MASK are true or if MASK has size zero, and the result has value false if any element of MASK is false.

23 *Case (ii):* If MASK has rank one, ALL(MASK,DIM) is equal to ALL(MASK). Otherwise, the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$ of ALL (MASK, DIM) is equal to ALL (MASK $(s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)$).

24 **Examples.**

25 *Case (i):* The value of ALL ((/ .TRUE., .FALSE., .TRUE. /)) is false.

1 *Case (ii):* If B is the array $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$ and C is the array $\begin{bmatrix} 0 & 3 & 5 \\ 7 & 4 & 8 \end{bmatrix}$ then ALL (B /= C,
2 DIM = 1) is [true, false, false] and ALL (B /= C, DIM = 2) is [false, false].

3 13.7.10 ALLOCATED (ARRAY) or ALLOCATED (SCALAR)

4 **Description.** Indicate whether an allocatable variable is allocated.

5 **Class.** Inquiry function.

6 **Arguments.**

7 ARRAY shall be an allocatable array.

8 SCALAR shall be an allocatable scalar.

9 **Result Characteristics.** Default logical scalar.

10 **Result Value.** The result has the value true if the argument (ARRAY or SCALAR) is allocated
11 and has the value false if the argument is unallocated.

12 13.7.11 ANINT (A [, KIND])

13 **Description.** Nearest whole number.

14 **Class.** Elemental function.

15 **Arguments.**

16 A shall be of type real.

17 KIND (optional) shall be a scalar integer initialization expression.

18 **Result Characteristics.** The result is of type real. If KIND is present, the kind type parameter
19 is that specified by the value of KIND; otherwise, the kind type parameter is that of A.

20 **Result Value.** The result is the integer nearest A, or if there are two integers equally near A,
21 the result is whichever such integer has the greater magnitude.

22 **Examples.** ANINT (2.783) has the value 3.0. ANINT (-2.783) has the value -3.0.

23 13.7.12 ANY (MASK [, DIM])

24 **Description.** Determine whether any value is true in MASK along dimension DIM.

25 **Class.** Transformational function.

26 **Arguments.**

27 MASK shall be of type logical. It shall be an array.

28 DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
where n is the rank of MASK. The corresponding actual argument shall not be
an optional dummy argument.

29 **Result Characteristics.** The result is of type logical with the same kind type parameter as
30 MASK. It is scalar if DIM is absent; otherwise, the result has rank $n - 1$ and shape $(d_1, d_2,$
31 $\dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of MASK.

1 **Result Value.**

2 *Case (i):* The result of ANY (MASK) has the value true if any element of MASK is true and
 3 has the value false if no elements are true or if MASK has size zero.

4 *Case (ii):* If MASK has rank one, ANY(MASK,DIM) is equal to ANY(MASK). Otherwise,
 5 the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$ of ANY(MASK, DIM)
 6 is equal to ANY(MASK $(s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)$).

7 **Examples.**

8 *Case (i):* The value of ANY ((/ .TRUE., .FALSE., .TRUE. /)) is true.

9 *Case (ii):* If B is the array $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$ and C is the array $\begin{bmatrix} 0 & 3 & 5 \\ 7 & 4 & 8 \end{bmatrix}$ then ANY(B /= C,
 10 DIM = 1) is [true, false, true] and ANY(B /= C, DIM = 2) is [true, true].

11 **13.7.13 ASIN (X)**

12 **Description.** Arcsine (inverse sine) function.

13 **Class.** Elemental function.

14 **Argument.** X shall be of type real with a value that satisfies the inequality $|X| \leq 1$, or of type
 15 complex.

16 **Result Characteristics.** Same as X.

17 **Result Value.** The result has a value equal to a processor-dependent approximation to arcsin(X).
 18 If it is real it is expressed in radians and lies in the range $-\pi/2 \leq \text{ASIN}(X) \leq \pi/2$. If it is
 19 complex the real part is expressed in radians and lies in the range $-\pi/2 \leq \text{REAL}(\text{ASIN}(X)) \leq \pi/2$.

20 **Example.** ASIN (0.84147098) has the value 1.0 (approximately).

21 **13.7.14 ASINH (X)**

22 **Description.** Inverse hyperbolic sine function.

23 **Class.** Elemental function.

24 **Argument.** X shall be of type real or complex.

25 **Result Characteristics.** Same as X.

26 **Result Value.** The result has a value equal to a processor-dependent approximation to the
 27 inverse hyperbolic sine function of X. If the result is complex the imaginary part is expressed in
 28 radians and lies in the range $-\pi/2 \leq \text{AIMAG}(\text{ASINH}(X)) \leq \pi/2$.

29 **Example.** ASINH(1.1752012) has the value 1.0 (approximately).

30 **13.7.15 ASSOCIATED (POINTER [, TARGET])**

31 **Description.** Returns the association status of its pointer argument or indicates whether the
 32 pointer is associated with the target.

33 **Class.** Inquiry function.

34 **Arguments.**

POINTER shall be a pointer. It may be of any type or may be a procedure pointer. Its
 pointer association status shall not be undefined.

TARGET shall be allowable as the *data-target* or *proc-target* in a pointer assignment statement (7.4.2) in which POINTER is *data-pointer-object* or *proc-pointer-object*. If TARGET is a pointer then its pointer association status shall not be undefined.

Result Characteristics. Default logical scalar.

Result Value.

Case (i): If TARGET is absent, the result is true if POINTER is associated with a target and false if it is not.

Case (ii): If TARGET is present and is a procedure, the result is true if POINTER is associated with TARGET.

Case (iii): If TARGET is present and is a procedure pointer, the result is true if POINTER and TARGET are associated with the same procedure. If either POINTER or TARGET is disassociated, the result is false.

Case (iv): If TARGET is present and is a scalar target, the result is true if TARGET is not a zero-sized storage sequence and the target associated with POINTER occupies the same storage units as TARGET. Otherwise, the result is false. If the POINTER is disassociated, the result is false.

Case (v): If TARGET is present and is an array target, the result is true if the target associated with POINTER and TARGET have the same shape, are neither of size zero nor arrays whose elements are zero-sized storage sequences, and occupy the same storage units in array element order. Otherwise, the result is false. If POINTER is disassociated, the result is false.

Case (vi): If TARGET is present and is a scalar pointer, the result is true if the target associated with POINTER and the target associated with TARGET are not zero-sized storage sequences and they occupy the same storage units. Otherwise, the result is false. If either POINTER or TARGET is disassociated, the result is false.

Case (vii): If TARGET is present and is an array pointer, the result is true if the target associated with POINTER and the target associated with TARGET have the same shape, are neither of size zero nor arrays whose elements are zero-sized storage sequences, and occupy the same storage units in array element order. Otherwise, the result is false. If either POINTER or TARGET is disassociated, the result is false.

Examples. ASSOCIATED (CURRENT, HEAD) is true if CURRENT is associated with the target HEAD. After the execution of

```
A_PART => A (:N)
```

ASSOCIATED (A_PART, A) is true if N is equal to UBOUND (A, DIM = 1). After the execution of

```
NULLIFY (CUR); NULLIFY (TOP)
```

ASSOCIATED (CUR, TOP) is false.

13.7.16 ATAN (X) or ATAN (Y, X)

Description. Arctangent (inverse tangent) function.

Class. Elemental function.

Arguments.

1 Y shall be of type real.
 X If Y is present, X shall be of type real with the same kind type parameter as
 2 Y. If Y has the value zero, X shall not have the value zero. If Y is absent, X
 shall be of type real or complex.

3 **Result Characteristics.** Same as X.

4 **Result Value.** If Y is present, the result is the same as the result of ATAN2(Y,X). If Y is absent,
 5 the result has a value equal to a processor-dependent approximation to $\arctan(X)$ whose real part
 6 is expressed in radians and lies in the range $-\pi/2 \leq \text{ATAN}(X) \leq \pi/2$.

7 **Example.** ATAN (1.5574077) has the value 1.0 (approximately).

8 13.7.17 ATAN2 (Y, X)

9 **Description.** Arctangent (inverse tangent) function. The result is the principal value of the
 10 argument of the nonzero complex number (X, Y).

11 **Class.** Elemental function.

12 **Arguments.**

13 Y shall be of type real.

14 X shall be of the same type and kind type parameter as Y. If Y has the value
 zero, X shall not have the value zero.

15 **Result Characteristics.** Same as X.

16 **Result Value.** The result has a value equal to a processor-dependent approximation to the
 17 principal value of the argument of the complex number (X, Y), expressed in radians. It lies in the
 18 range $-\pi \leq \text{ATAN2}(Y,X) \leq \pi$ and is equal to a processor-dependent approximation to a value of
 19 $\arctan(Y/X)$ if $X \neq 0$. If $Y > 0$, the result is positive. If $Y = 0$ and $X > 0$, the result is Y. If $Y = 0$
 20 and $X < 0$, then the result is approximately π if Y is positive real zero or the processor cannot
 21 distinguish between positive and negative real zero, and approximately $-\pi$ if Y is negative real
 22 zero. If $Y < 0$, the result is negative. If $X = 0$, the absolute value of the result is approximately
 23 $\pi/2$.

24 **Examples.** ATAN2 (1.5574077, 1.0) has the value 1.0 (approximately). If Y has the value
 25 $\begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}$ and X has the value $\begin{bmatrix} -1 & 1 \\ -1 & 1 \end{bmatrix}$, the value of ATAN2 (Y, X) is approximately
 26 $\begin{bmatrix} \frac{3\pi}{4} & \frac{\pi}{4} \\ -\frac{3\pi}{4} & -\frac{\pi}{4} \end{bmatrix}$.

27 13.7.18 ATANH (X)

28 **Description.** Inverse hyperbolic tangent function.

29 **Class.** Elemental function.

30 **Argument.** X shall be of type real or complex.

31 **Result Characteristics.** Same as X.

32 **Result Value.** The result has a value equal to a processor-dependent approximation to the
 33 inverse hyperbolic tangent function of X. If the result is complex the imaginary part is expressed
 34 in radians and lies in the range $-\pi/2 \leq \text{AIMAG}(\text{ATANH}(X)) \leq \pi/2$.

1 **Example.** ATANH(0.76159416) has the value 1.0 (approximately).

2 **13.7.19 BESSEL_J0 (X)**

3 **Description.** Bessel function of the first kind of order zero.

4 **Class.** Elemental function.

5 **Argument.** X shall be of type real.

6 **Result Characteristics.** Same as X.

7 **Result Value.** The result has a value equal to a processor-dependent approximation to the Bessel
8 function of the first kind of order zero of X.

9 **Example.** BESSEL_J0 (1.0) has the value 0.765 (approximately).

10 **13.7.20 BESSEL_J1 (X)**

11 **Description.** Bessel function of the first kind of order one.

12 **Class.** Elemental function.

13 **Argument.** X shall be of type real.

14 **Result Characteristics.** Same as X.

15 **Result Value.** The result has a value equal to a processor-dependent approximation to the Bessel
16 function of the first kind of order one of X.

17 **Example.** BESSEL_J1 (1.0) has the value 0.440 (approximately).

18 **13.7.21 BESSEL_JN (N, X)**

19 **Description.** Bessel function of the first kind of order N.

20 **Class.** Elemental function.

21 **Arguments.**

22 N shall be of type integer and nonnegative.

23 X shall be of type real.

24 **Result Characteristics.** Same as X.

25 **Result Value.** The result has a value equal to a processor-dependent approximation to the Bessel
26 function of the first kind of order N of X.

27 **Example.** BESSEL_JN (2, 1.0) has the value 0.115 (approximately).

28 **13.7.22 BESSEL_Y0 (X)**

29 **Description.** Bessel function of the second kind of order zero.

30 **Class.** Elemental function.

31 **Argument.** X shall be of type real. Its value shall be greater than zero.

32 **Result Characteristics.** Same as X.

1 **Result Value.** The result has a value equal to a processor-dependent approximation to the Bessel
2 function of the second kind of order zero of X.

3 **Example.** BESSEL_Y0(1.0) has the value 0.088 (approximately).

4 **13.7.23 BESSEL_Y1 (X)**

5 **Description.** Bessel function of the second kind of order one.

6 **Class.** Elemental function.

7 **Argument.** X shall be of type real. Its value shall be greater than zero.

8 **Result Characteristics.** Same as X.

9 **Result Value.** The result has a value equal to a processor-dependent approximation to the Bessel
10 function of the second kind of order one of X.

11 **Example.** BESSEL_Y1 (1.0) has the value -0.781 (approximately).

12 **13.7.24 BESSEL_YN (N, X)**

13 **Description.** Bessel function of the second kind of order N.

14 **Class.** Elemental function.

15 **Arguments.**

16 N shall be of type integer and nonnegative.

17 X shall be of type real. Its value shall be greater than zero.

18 **Result Characteristics.** Same as X.

19 **Result Value.** The result has a value equal to a processor-dependent approximation to the Bessel
20 function of the second kind of order N of X.

21 **Example.** BESSEL_YN (2, 1.0) has the value -1.651 (approximately).

22 **13.7.25 BIT_SIZE (I)**

23 **Description.** Returns the number of bits z defined by the model of 13.3.

24 **Class.** Inquiry function.

25 **Argument.** I shall be of type integer or bits. It may be a scalar or an array.

26 **Result Characteristics.** Scalar integer. If I is of type integer the kind type parameter is that
27 of I, otherwise it is that of default integer type.

28 **Result Value.** If I is of type integer, the result has the value of the number of bits z of the model
29 integer defined for bit manipulation contexts in 13.3. If I is of type bits, the result has the value
30 KIND(I).

31 **Examples.** BIT_SIZE (1) has the value 32 if z of the model is 32. BIT_SIZE(Z'00FF') has the
32 value 16.

33 **13.7.26 BITS (A [, KIND])**

1 **Description.** Convert to bits type.

2 **Class.** Elemental function.

3 **Arguments.**

4 A shall be of type integer, real, complex, logical, or bits.

5 KIND (optional) shall be a scalar integer initialization expression.

6 **Result Characteristics.** Bits. If KIND is present, the kind type parameter is that specified by
7 the value of KIND; otherwise, the kind type parameter is BITS_KIND(A).

8 **Result Value.** If A is of type bits and the kind type parameter of A is the same as that of the
9 result, the result value is the value of A.

10 If A is of type bits with a smaller kind type parameter value than that of the result, the rightmost
11 KIND(A) bits of the result value are the same as those of A and the remaining bits of the result
12 are 0.

13 If A is of type bits with a larger kind type parameter value than that of the result, the result
14 value consists of the rightmost KIND(*result*) bits of A.

15 If A is not of type bits and BITS_KIND(A) is greater than or equal to the kind type param-
16 eter of the result, the result value consists of the rightmost KIND(*result*) bits of the physical
17 representation of A.

18 If A is not of type bits and BITS_KIND(A) is less than or equal to the kind type parameter of
19 the result, the rightmost bits of the result are those of the physical representation of A and the
20 remaining bits of the result are 0.

21 **Examples.** BITS (Z'AB', 16) has the value Z'00AB'. BITS (-1) has the value Z'FFFFFFFF'
22 for a processor whose default integer representation is 32-bit two's-complement. BITS (.TRUE.,
23 5) has the value B'00001' for a processor that represents the logical value true by setting the low
24 order bit of the internal value and clearing the other bits. BITS (X) has the value Z'7F800000' if
25 the value of X is an IEEE single precision positive infinity.

26 13.7.27 BITS_KIND (X)

27 **Description.** Return bits kind compatible with the argument.

28 **Class.** Inquiry function.

29 **Arguments.**

30 X shall be of type bits, integer, real, complex, or logical.

31 **Result Characteristics.** Default integer scalar.

32 **Result Value.** If X is of type bits, the result has the value KIND (X). If X is of type default
33 integer, default real, or default logical, the result has the value NUMERIC_STORAGE_SIZE
34 (13.8.3.14). If X is of type double precision or default complex, the result has the value 2 ×
35 NUMERIC_STORAGE_SIZE. If X is of type complex with the same kind type parameter as that
36 of double precision, the result has the value 4 × NUMERIC_STORAGE_SIZE. If X is of type
37 non-default integer, the result has the value BIT_SIZE (X). If X is of a non-default logical or
38 non-default non-integer numeric type, the result value is the number of bits of storage used by
39 the processor to represent scalar objects of that type and kind.

J3 internal note

Unresolved Technical Issue 061

COMPLEX(KIND(0d0)) is neither required nor guaranteed to take double the space of DOUBLE PRECISION; it takes a single (unique) storage unit.

BIT_SIZE(X) does not return the number of bits taken for the storage of X, it returns the number of bits available for bit operations. If you want to use BITS type to hold the whole physical representation (including any functionally useless bits or padding bits), this means that BIT_SIZE and BITS_KIND are not (necessarily) going to return the same value.

Indeed, the number of bits for the “storage” of X is not actually a completely well-defined concept except in the context of storage association, and even there there is some leeway. Different entities with the same type and type parameters may use different amounts of storage (so long as argument association etc. gives the right answer) or even different amounts of storage at different times.

Indeed, BIT_SIZE(default integer) is neither required nor guaranteed to be the same as NUMERIC_STORAGE_SIZE. Nor does this requirement appear to be necessary or even all that useful for the BITS feature.

If (and only if) we want to impose this unnecessary requirement on all future processors, it should be done properly as a straight-forward requirement (that BIT_SIZE(17)==NUMERIC_STORAGE_SIZE), not obscurely by the BITS_KIND specification.

Not imposing this requirement means just deleting the 3rd and 4th sentences and simplifying the wording of the final sentence to be an “Otherwise...”.

1 **Example.** The value of BITS_KIND (0) is 32 if the size of a numeric storage unit is 32 bits.

2 **13.7.28 BTEST (I, POS)**

3 **Description.** Tests a bit of an integer or bits value.

4 **Class.** Elemental function.

5 **Arguments.**

6 I shall be of type integer or bits.

7 POS shall be of type integer. It shall be nonnegative and be less than BIT_SIZE (I).

8 **Result Characteristics.** Default logical.

9 **Result Value.** The result has the value true if bit POS of I has the value 1 and has the value
10 false if bit POS of I has the value 0. The model for the interpretation of an integer value as a
11 sequence of bits is in 13.3.

12 **Examples.** BTEST (8, 3) has the value true. If A has the value $\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$, the value of
13 BTEST (A, 2) is $\begin{bmatrix} \text{false} & \text{false} \\ \text{false} & \text{true} \end{bmatrix}$ and the value of BTEST (2, A) is $\begin{bmatrix} \text{true} & \text{false} \\ \text{false} & \text{false} \end{bmatrix}$.
14 BTEST (B'01000', 3) has the value true.

15 **13.7.29 CEILING (A [, KIND])**

16 **Description.** Returns the least integer greater than or equal to its argument.

17 **Class.** Elemental function.

18 **Arguments.**

19 A shall be of type real.

1 KIND (optional) shall be a scalar integer initialization expression.

2 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
3 by the value of KIND; otherwise, the kind type parameter is that of default integer type.

4 **Result Value.** The result has a value equal to the least integer greater than or equal to A.

5 **Examples.** CEILING (3.7) has the value 4. CEILING (-3.7) has the value -3.

6 13.7.30 CHAR (I [, KIND])

7 **Description.** Returns the character in a given position of the processor collating sequence asso-
8 ciated with the specified kind type parameter. It is the inverse of the ICHAR function.

9 **Class.** Elemental function.

10 **Arguments.**

I shall be of type integer or bits. If I is of type bits, it is interpreted as a non-
negative integer as described in 13.3. The value of I shall be in the range
11 $0 \leq I \leq n - 1$, where n is the number of characters in the collating sequence
associated with the specified kind type parameter.

12 KIND (optional) shall be a scalar integer initialization expression.

13 **Result Characteristics.** Character of length one. If KIND is present, the kind type parameter
14 is that specified by the value of KIND; otherwise, the kind type parameter is that of default
15 character type.

16 **Result Value.** The result is the character in position I of the collating sequence associated with
17 the specified kind type parameter. ICHAR (CHAR (I, KIND (C))) shall have the value I for
18 $0 \leq I \leq n - 1$ and CHAR (ICHAR (C), KIND (C)) shall have the value C for any character C
19 capable of representation in the processor.

20 **Examples.** CHAR (88) has the value 'X' on a processor using the ASCII collating sequence
21 for the default character type. CHAR (Z'41') has the value 'A' on a processor using the ASCII
22 collating sequence.

23 13.7.31 CMPLX (X [, Y, KIND])

24 **Description.** Convert to complex type.

25 **Class.** Elemental function.

26 **Arguments.**

27 X shall be of type integer, real, or complex, or bits.

28 Y (optional) shall be of type integer or real, or bits. If X is of type complex, Y shall not be
present, nor shall Y be associated with an optional dummy argument.

29 KIND (optional) shall be a scalar integer initialization expression.

30 **Result Characteristics.** The result is of type complex. If KIND is present, the kind type
31 parameter is that specified by the value of KIND; otherwise, the kind type parameter is that of
32 default real type.

33 **Result Value.** If Y is absent and X is not complex, it is as if Y were present with the value zero.
34 If X is complex, it is as if X were real with the value REAL (X, KIND) and Y were present with

1 the value AIMAG (X, KIND). CMPLX (X, Y, KIND) has the complex value whose real part is
2 REAL (X, KIND) and whose imaginary part is REAL (Y, KIND).

3 **Example.** CMPLX (-3) has the value (-3.0, 0.0).

4 **13.7.32 CO_ALL (MASK, RESULT [, TEAM])**

5 **Description.** Determine whether all values are true on a team of images.

6 **Class.** Collective subroutine.

7 **Arguments.**

8 MASK shall be a co-array of type logical. It may be a scalar or an array. It is an
INTENT(IN) argument.

9 RESULT shall be of type logical and have the same shape as MASK. It is an IN-
TENT(OUT) argument. If it is scalar, it is assigned the value true if the value
of MASK is true on all the images of the team, and the value false otherwise.
If it is an array, each element is assigned the value true if all corresponding
elements of MASK are true on all the images of the team, and the value false
otherwise.

10 TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN)
argument that specifies the team for which CO_ALL is performed. If TEAM
is not present, the team consists of all images.

11 **Example.** If the number of images is two and MASK is the array [true, false, true] on one image
12 and [true, true, true] on the other image, the value of RESULT after the statement CALL CO-
13 ALL (MASK, RESULT) is [true, false, true].

14 **13.7.33 CO_ANY (MASK, RESULT [, TEAM])**

15 **Description.** Determine whether any value is true on a team of images.

16 **Class.** Collective subroutine.

17 **Arguments.**

18 MASK shall be a co-array of type logical. It may be a scalar or an array. It is an
INTENT(IN) argument.

19 RESULT shall be of type logical and have the same shape as MASK. It is an IN-
TENT(OUT) argument. If it is scalar it is assigned the value true if any
value of MASK is true on any image of the team, and false otherwise. If it is
an array, each element is assigned the value true if any of the corresponding
elements of MASK are true on any image of the team, and false otherwise.

20 TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN)
argument that specifies the team for which CO_ANY is performed. If TEAM
is not present, the team consists of all images.

21 **Example.** If the number of images is two and MASK is the array [true, false, false] on one image
22 and [true, true, false] on the other image, the value of RESULT after the statement CALL CO-
23 ANY (MASK, RESULT) is [true, true, false].

24 **13.7.34 CO_COUNT (MASK, RESULT [, TEAM])**

Description. Count the numbers of true elements on a team of images.

1

2

Class. Collective subroutine.

3

Arguments.

4

MASK shall be a co-array of type logical. It may be a scalar or an array. It is an INTENT(IN) argument.

5

RESULT shall be of type integer and have the same shape as MASK. It is an INTENT(OUT) argument. If it is scalar, it is assigned a value equal to the number of images of the team for which MASK has the value true. If it is an array, each element is assigned a value equal to the number of corresponding elements of MASK on the images of the team that have the value true.

6

TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN) argument that specifies the team for which CO_COUNT is performed. If TEAM is not present, the team consists of all images.

7

Example. If the number of images is two and MASK is the array [true, false, false] on one image and [true, true, false] on the other image, the value of result after the statement CALL CO_COUNT (MASK, RESULT) is [2, 1, 0].

9

10 13.7.35 CO_LBOUND (CO_ARRAY [, DIM, KIND])

11

Description. Returns all the lower co-bounds or a specified lower co-bound of a co-array.

12

Class. Inquiry function.

13

Arguments.

14

CO_ARRAY shall be a co-array and may be of any type. It may be a scalar or an array. If it is allocatable it shall be allocated.

15

DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where n is the co-rank of CO_ARRAY. The corresponding actual argument shall not be an optional dummy argument.

16

KIND (optional) shall be a scalar integer initialization expression.

17

Result Characteristics. Integer. If KIND is present, the kind type parameter is that specified by the value of KIND; otherwise, the kind type parameter is that of default integer type. The result is scalar if DIM is present; otherwise, the result is an array of rank one and size n , where n is the co-rank of CO_ARRAY.

21

Result Value.

22

Case (i): CO_LBOUND (CO_ARRAY, DIM) has a value equal to the lower co-bound for co-subscript DIM of CO_ARRAY.

23

24

Case (ii): CO_LBOUND (CO_ARRAY) has a value whose i^{th} element is equal to CO_LBOUND (CO_ARRAY, i), for $i = 1, 2, \dots, n$, where n is the co-rank of CO_ARRAY.

25

26

27

Examples. If A is allocated by the statement ALLOCATE (A [2:3, 7:*]) then CO_LBOUND (A) is [2, 7] and CO_LBOUND (A, DIM=2) is 7.

28

29 13.7.36 CO_MAXLOC (CO_ARRAY, RESULT [, TEAM])

30

Description. Image indices of the maximum values of the elements on a team of images.

1 **Class.** Collective subroutine.

2 **Arguments.**

3 CO_ARRAY shall be a co-array of type integer, real, bits, or character. It may be a scalar
or an array. It is an INTENT(IN) argument.

4 RESULT shall be of type integer and have the same shape as CO_ARRAY. It is an
INTENT(OUT) argument. If it is scalar, it is assigned a value equal to the
image index of the maximum value of CO_ARRAY on the images of the team;
if more than one image has the maximum value, the smallest such image index
is assigned. If RESULT is an array, each element of RESULT is assigned a
value equal to the image index of the maximum value of all the corresponding
elements of CO_ARRAY on the images of the team; if more than one image
has the maximum value, the smallest such image index is assigned.
If CO_ARRAY is of type character, the result is the value that would be selected
by application of intrinsic relational operators; that is, the collating sequence
for characters with the kind type parameter of the argument is applied.

5 TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN)
argument that specifies the team for which CO_MAXLOC is performed. If
TEAM is not present, the team consists of all images.

6 **Example.** If the number of images is two and CO_ARRAY is the array [1, 5, 6] on one image
7 and [4, 1, 6] on the other image, the value of RESULT after the statement
8 CALL CO_MAXLOC (CO_ARRAY, RESULT) is [2, 1, 1].

9 **13.7.37 CO_MAXVAL (CO_ARRAY, RESULT [, TEAM])**

10 **Description.** Maximum values of the elements on a team of images.

11 **Class.** Collective subroutine.

12 **Arguments.**

13 CO_ARRAY shall be a co-array of type integer, real, bits, or character. It may be a scalar
or an array. It is an INTENT(IN) argument.

14 RESULT shall be of the same type and type parameters as CO_ARRAY, and shall have
the same shape as CO_ARRAY. It is an INTENT(OUT) argument. If it is
scalar, it is assigned a value equal to the maximum value of CO_ARRAY on all
the images of the team. If it is an array, each element is assigned a value equal
to the maximum value of all the corresponding elements of CO_ARRAY on all
the images of the team.
If CO_ARRAY is of type character, the result is the value that would be selected
by application of intrinsic relational operators; that is, the collating sequence
for characters with the kind type parameter of the argument is applied.

15 TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN)
argument that specifies the team for which CO_MAXVAL is performed. If
TEAM is not present, the team consists of all images.

16 **Example.** If the number of images is two and CO_ARRAY is the array [1, 5, 3] on one image
17 and [4, 1, 6] on the other image, the value of RESULT after the statement
18 CALL CO_MAXVAL (CO_ARRAY, RESULT) is [4, 5, 6].

19 **13.7.38 CO_MINLOC (CO_ARRAY, RESULT [, TEAM])**

1 **Description.** Image indices of the minimum values of the elements on a team of images.

2 **Class.** Collective subroutine.

3 **Arguments.**

4 CO_ARRAY shall a co-array of type integer, real, bits, or character. It may be a scalar or
an array. It is an INTENT(IN) argument.

5 RESULT shall be of type integer and have the same shape as CO_ARRAY. It is an
INTENT(OUT) argument. If it is scalar, it is assigned a value equal to the
image index of the minimum value of CO_ARRAY on all the images of the
team; if more than one image has the minimum value, the smallest such image
index is assigned. If it is an array, each element is assigned a value equal to the
image index of the minimum value of all the corresponding elements of CO_
ARRAY on the images of the team; if more than one image has the minimum
value, the smallest such image index is assigned.
If CO_ARRAY is of type character, the result is the value that would be selected
by application of intrinsic relational operators; that is, the collating sequence
for characters with the kind type parameter of the argument is applied.

6 TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN)
argument that specifies the team for which CO_MINLOC is performed. If
TEAM is not present, the team consists of all images.

7 **Example.** If the number of images is two and CO_ARRAY is the array [1, 5, 6] on one image
8 and [4, 1, 6] on the other image, the value of RESULT after the statement
9 CALL CO_MINLOC (ARRAY, RESULT) is [1, 2, 1].

10 **13.7.39 CO_MINVAL (CO_ARRAY, RESULT [, TEAM])**

11 **Description.** Minimum values of the elements on a team of images.

12 **Class.** Collective subroutine.

13 **Arguments.**

14 CO_ARRAY shall be a co-array of type integer, real, bits, or character. It may be a scalar
or an array. It is an INTENT(IN) argument.

15 RESULT shall be of the same type and type parameters as CO_ARRAY, and shall have
the same shape as CO_ARRAY. It is an INTENT(OUT) argument. If it is
scalar it is assigned a value equal to the minimum value of CO_ARRAY on all
the images of the team. If it is an array, each element is assigned a value equal
to the minimum value of all the corresponding elements of CO_ARRAY on all
the images of the team.
If CO_ARRAY is of type character, the result is the value that would be selected
by application of intrinsic relational operators; that is, the collating sequence
for characters with the kind type parameter of the argument is applied.

16 TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN)
argument that specifies the team for which CO_MINVAL is performed. If
TEAM is not present, the team consists of all images.

17 **Example.** If the number of images is two and CO_ARRAY is the array [1, 5, 3] on one image
18 and [4, 1, 6] on the other image, the value of RESULT after the statement
19 CALL CO_MINVAL (CO_ARRAY, RESULT) is [1, 1, 3].

13.7.40 CO_PRODUCT (CO_ARRAY, RESULT [, TEAM])

Description. Products of elements on a team of images.

Class. Collective subroutine.

Arguments.

CO_ARRAY shall be a co-array of type integer, real, or complex. It may be a scalar or an array. It is an INTENT(IN) argument.

RESULT shall be of the same type and type parameters as CO_ARRAY, and shall have the same shape as CO_ARRAY. It is an INTENT(OUT) argument. If it is scalar, it is assigned a value equal to a processor-dependent and image-dependent approximation to the product of the values of CO_ARRAY on all the images of the team. If it is an array, each element is assigned a value equal to a processor-dependent and image-dependent approximation to the product of all the corresponding elements of CO_ARRAY on the images of the team.

TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN) argument that specifies the team for which CO_PRODUCT is performed. If TEAM is not present, the team consists of all images.

Example. If the number of images is two and CO_ARRAY is the array [1, 5, 3] on one image and [4, 1, 6] on the other image, the value of RESULT after the statement CALL CO_PRODUCT (CO_ARRAY, RESULT) is [4, 5, 18].

13.7.41 CO_SUM (CO_ARRAY, RESULT [, TEAM])

Description. Sums of elements on a team of images.

Class. Collective subroutine.

Arguments.

CO_ARRAY shall be a co-array of type integer, real, or complex. It may be a scalar or an array. It is an INTENT(IN) argument.

RESULT shall be of the same type and type parameters as CO_ARRAY, and shall have the same shape as CO_ARRAY. It is an INTENT(OUT) argument. If it is scalar, it is assigned a value equal to a processor-dependent and image-dependent approximation to the sum of the values of CO_ARRAY on all the images of the team. If it is an array, each element is assigned a value equal to a processor-dependent and image-dependent approximation to the sum of all the corresponding elements of CO_ARRAY on the images of the team.

TEAM (optional) shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(IN) argument that specifies the team for which CO_SUM is performed. If TEAM is not present, the team consists of all images.

Example. If the number of images is two and CO_ARRAY is the array [1, 5, 3] on one image and [4, 1, 6] on the other image, the value of RESULT after the statement CALL CO_SUM (CO_ARRAY, RESULT) is [5, 6, 9].

13.7.42 CO_UBOUND (CO_ARRAY [, DIM, KIND])

Description. Returns all the upper co-bounds or a specified upper co-bound of a co-array of co-rank greater than one.

1 **Class.** Inquiry function.

2 **Arguments.**

3 CO_ARRAY shall be a co-array of any type and of co-rank greater than one. It may be a scalar or an array. If it is allocatable it shall be allocated.

4 DIM (optional) d shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where n is the co-rank of CO_ARRAY. The corresponding actual argument shall not be an optional dummy argument.

5 KIND (optional) shall be a scalar integer initialization expression.

6 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified by the value of KIND; otherwise, the kind type parameter is that of default integer type. The result is scalar if DIM is present; otherwise, the result is an array of rank one and size $n - 1$, where n is the co-rank of CO_ARRAY.

10 **Result Value.**

11 *Case (i):* CO_UBOUND (CO_ARRAY, DIM) has a value equal to the upper co-bound for co-subscript DIM of CO_ARRAY.

12
13 *Case (ii):* CO_UBOUND (CO_ARRAY) has a value whose i^{th} element is equal to
14 CO_UBOUND (CO_ARRAY, i), for $i = 1, 2, \dots, n - 1$, where n is the co-rank of
15 CO_ARRAY.

16 **Examples.** If A is allocated by the statement ALLOCATE (A [2:3, 0:7, *]) then
17 CO_UBOUND (A) is [3, 7] and CO_UBOUND (A, DIM=2) is 7.

18 13.7.43 COMMAND_ARGUMENT_COUNT ()

19 **Description.** Returns the number of command arguments.

20 **Class.** Inquiry function.

21 **Argument.** None.

22 **Result Characteristics.** Scalar default integer.

23 **Result Value.** The result value is equal to the number of command arguments available. If there
24 are no command arguments available or if the processor does not support command arguments,
25 then the result has the value zero. If the processor has a concept of a command name, the
26 command name does not count as one of the command arguments.

27 **Example.** See 13.7.73.

28 13.7.44 CONJG (Z)

29 **Description.** Conjugate of a complex number.

30 **Class.** Elemental function.

31 **Argument.** Z shall be of type complex.

32 **Result Characteristics.** Same as Z.

33 **Result Value.** If Z has the value (x, y) , the result has the value $(x, -y)$.

34 **Example.** CONJG ((2.0, 3.0)) has the value (2.0, -3.0).

1 13.7.45 COS (X)

2 **Description.** Cosine function.

3 **Class.** Elemental function.

4 **Argument.** X shall be of type real or complex.

5 **Result Characteristics.** Same as X.

6 **Result Value.** The result has a value equal to a processor-dependent approximation to $\cos(X)$.
 7 If X is of type real, it is regarded as a value in radians. If X is of type complex, its real part is
 8 regarded as a value in radians.

9 **Example.** COS (1.0) has the value 0.54030231 (approximately).

10 13.7.46 COSH (X)

11 **Description.** Hyperbolic cosine function.

12 **Class.** Elemental function.

13 **Argument.** X shall be of type real or complex.

14 **Result Characteristics.** Same as X.

15 **Result Value.** The result has a value equal to a processor-dependent approximation to $\cosh(X)$.
 16 If X is of type complex its imaginary part is regarded as a value in radians.

17 **Example.** COSH (1.0) has the value 1.5430806 (approximately).

18 13.7.47 COUNT (MASK [, DIM, KIND])

19 **Description.** Count the number of true elements of MASK along dimension DIM.

20 **Class.** Transformational function.

21 **Arguments.**

22 MASK shall be of type logical. It shall be an array.

23 DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
 where n is the rank of MASK. The corresponding actual argument shall not be
 an optional dummy argument.

24 KIND (optional) shall be a scalar integer initialization expression.

25 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified by
 26 the value of KIND; otherwise the kind type parameter is that of default integer type. The result is
 27 scalar if DIM is absent; otherwise, the result has rank $n-1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1},$
 28 $\dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of MASK.

29 **Result Value.**

30 *Case (i):* The result of COUNT (MASK) has a value equal to the number of true elements
 31 of MASK or has the value zero if MASK has size zero.

32 *Case (ii):* If MASK has rank one, COUNT (MASK, DIM) has a value equal to that of
 33 COUNT (MASK). Otherwise, the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1},$
 34 $\dots, s_n)$ of COUNT (MASK, DIM) is equal to COUNT (MASK $(s_1, s_2, \dots, s_{\text{DIM}-1},$

1 :, s_{DIM+1}, \dots, s_n)).

2 **Examples.**

3 *Case (i):* The value of COUNT ((/ .TRUE., .FALSE., .TRUE. /)) is 2.

4 *Case (ii):* If B is the array $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$ and C is the array $\begin{bmatrix} 0 & 3 & 5 \\ 7 & 4 & 8 \end{bmatrix}$, COUNT (B /= C,
5 DIM = 1) is [2, 0, 1] and COUNT (B /= C, DIM = 2) is [1, 2].

6 **13.7.48 CPU_TIME (TIME)**

7 **Description.** Returns the processor time.

8 **Class.** Subroutine.

9 **Argument.** TIME shall be scalar and of type real. It is an INTENT(OUT) argument that is
10 assigned a processor-dependent approximation to the processor time in seconds. If the processor
11 cannot return a meaningful time, a processor-dependent negative value is returned.

12 **Example.**

```
13       REAL T1, T2
14       ...
15       CALL CPU_TIME(T1)
16       ... ! Code to be timed.
17       CALL CPU_TIME(T2)
18       WRITE (*,*) 'Time taken by code was ', T2-T1, ' seconds'
```

19 writes the processor time taken by a piece of code.

NOTE 13.7

A processor for which a single result is inadequate (for example, a parallel processor) might choose to provide an additional version for which time is an array.

The exact definition of time is left imprecise because of the variability in what different processors are able to provide. The primary purpose is to compare different algorithms on the same processor or discover which parts of a calculation are the most expensive.

The start time is left imprecise because the purpose is to time sections of code, as in the example.

Most computer systems have multiple concepts of time. One common concept is that of time expended by the processor for a given program. This might or might not include system overhead, and has no obvious connection to elapsed “wall clock” time.

20 **13.7.49 CSHIFT (ARRAY, SHIFT [, DIM])**

21 **Description.** Perform a circular shift on an array expression of rank one or perform circular
22 shifts on all the complete rank one sections along a given dimension of an array expression of
23 rank two or greater. Elements shifted out at one end of a section are shifted in at the other end.
24 Different sections may be shifted by different amounts and in different directions.

25 **Class.** Transformational function.

Arguments.

1

2

ARRAY may be of any type. It shall be an array.

3

SHIFT shall be of type integer and shall be scalar if ARRAY has rank one; otherwise, it shall be scalar or of rank $n - 1$ and of shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

4

DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where n is the rank of ARRAY. If DIM is omitted, it is as if it were present with the value 1.

5

Result Characteristics. The result is of the type and type parameters of ARRAY, and has the shape of ARRAY.

6

7

Result Value.

8

Case (i): If ARRAY has rank one, element i of the result is $\text{ARRAY}(1 + \text{MODULO}(i + \text{SHIFT} - 1, \text{SIZE}(\text{ARRAY})))$.

9

10

Case (ii): If ARRAY has rank greater than one, section $(s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)$ of the result has a value equal to $\text{CSHIFT}(\text{ARRAY}(s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n), sh, 1)$, where sh is SHIFT or $\text{SHIFT}(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$.

11

12

13

Examples.

14

Case (i): If V is the array [1, 2, 3, 4, 5, 6], the effect of shifting V circularly to the left by two positions is achieved by $\text{CSHIFT}(V, \text{SHIFT} = 2)$ which has the value [3, 4, 5, 6, 1, 2]; $\text{CSHIFT}(V, \text{SHIFT} = -2)$ achieves a circular shift to the right by two positions and has the value [5, 6, 1, 2, 3, 4].

15

16

17

18

Case (ii): The rows of an array of rank two may all be shifted by the same amount or by

19

different amounts. If M is the array $\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$, the value of

20

$\text{CSHIFT}(M, \text{SHIFT} = -1, \text{DIM} = 2)$ is $\begin{bmatrix} 3 & 1 & 2 \\ 6 & 4 & 5 \\ 9 & 7 & 8 \end{bmatrix}$, and the value of

21

$\text{CSHIFT}(M, \text{SHIFT} = (/ -1, 1, 0 /), \text{DIM} = 2)$ is $\begin{bmatrix} 3 & 1 & 2 \\ 5 & 6 & 4 \\ 7 & 8 & 9 \end{bmatrix}$.

22

23

13.7.50 DATE_AND_TIME ([DATE, TIME, ZONE, VALUES])

24

Description. Returns data about the real-time clock and date in a form compatible with the representations defined in ISO 8601:1988.

25

26

Class. Subroutine.

27

Arguments.

DATE (optional) shall be scalar and of type default character. It is an INTENT (OUT) argument. It is assigned a value of the form *CCYYMMDD*, where *CC* is the century, *YY* is the year within the century, *MM* is the month within the year, and *DD* is the day within the month. If there is no date available, DATE is assigned all blanks.

28

1 TIME (optional) shall be scalar and of type default character. It is an INTENT (OUT) argument. It is assigned a value of the form *hhmmss.sss*, where *hh* is the hour of the day, *mm* is the minutes of the hour, and *ss.sss* is the seconds and milliseconds of the minute. If there is no clock available, TIME is assigned all blanks.

2 ZONE (optional) shall be scalar and of type default character. It is an INTENT (OUT) argument. It is assigned a value of the form *+hhmm* or *-hhmm*, where *hh* and *mm* are the time difference with respect to Coordinated Universal Time (UTC) in hours and minutes, respectively. If this information is not available, ZONE is assigned all blanks.

3 VALUES (optional) shall be of type default integer and of rank one. It is an INTENT (OUT) argument. Its size shall be at least 8. The values returned in VALUES are as follows:

4 VALUES (1) the year, including the century (for example, 1990), or -HUGE (0) if there is no date available;

5 VALUES (2) the month of the year, or -HUGE (0) if there is no date available;

6 VALUES (3) the day of the month, or -HUGE (0) if there is no date available;

7 VALUES (4) the time difference with respect to Coordinated Universal Time (UTC) in minutes, or -HUGE (0) if this information is not available;

8 VALUES (5) the hour of the day, in the range of 0 to 23, or -HUGE (0) if there is no clock;

9 VALUES (6) the minutes of the hour, in the range 0 to 59, or -HUGE (0) if there is no clock;

10 VALUES (7) the seconds of the minute, in the range 0 to 60, or -HUGE (0) if there is no clock;

11 VALUES (8) the milliseconds of the second, in the range 0 to 999, or -HUGE (0) if there is no clock.

12 **Example.**

```
13     INTEGER DATE_TIME (8)
14     CHARACTER (LEN = 10) BIG_BEN (3)
15     CALL DATE_AND_TIME (BIG_BEN (1), BIG_BEN (2), &
16                        BIG_BEN (3), DATE_TIME)
```

17 If run in Geneva, Switzerland on April 12, 1985 at 15:27:35.5 with a system configured for the
18 local time zone, this sample would have assigned the value 19850412 to BIG_BEN (1), the value
19 152735.500 to BIG_BEN (2), the value +0100 to BIG_BEN (3), and the value (/ 1985, 4, 12, 60,
20 15, 27, 35, 500 /) to DATE.TIME.

NOTE 13.8

UTC is defined by ISO 8601:1988.

21 **13.7.51 DBLE (A)**

22 **Description.** Convert to double precision real type.

23 **Class.** Elemental function.

Argument. A shall be of type integer, real, or complex, or bits.

1

2 **Result Characteristics.** Double precision real.3 **Result Value.** The result has the value REAL (A, KIND (0.0D0)).4 **Example.** DBLE (-3) has the value -3.0D0.5 **13.7.52 DIGITS (X)**6 **Description.** Returns the number of significant digits of the model representing numbers of the
7 same type and kind type parameter as the argument.8 **Class.** Inquiry function.9 **Argument.** X shall be of type integer or real. It may be a scalar or an array.10 **Result Characteristics.** Default integer scalar.11 **Result Value.** The result has the value q if X is of type integer and p if X is of type real, where
12 q and p are as defined in 13.4 for the model representing numbers of the same type and kind type
13 parameter as X.14 **Example.** DIGITS (X) has the value 24 for real X whose model is as in Note 13.4.15 **13.7.53 DIM (X, Y)**16 **Description.** The difference X-Y if it is positive; otherwise zero.17 **Class.** Elemental function.18 **Arguments.**

19 X shall be of type integer or real.

20 Y shall be of the same type and kind type parameter as X.

21 **Result Characteristics.** Same as X.22 **Result Value.** The value of the result is X-Y if X>Y and zero otherwise.23 **Example.** DIM (-3.0, 2.0) has the value 0.0.24 **13.7.54 DOT_PRODUCT (VECTOR_A, VECTOR_B)**25 **Description.** Performs dot-product multiplication of numeric or logical vectors.26 **Class.** Transformational function.27 **Arguments.**28 VECTOR_A shall be of numeric type (integer, real, or complex) or of logical type. It shall
28 be a rank-one array.29 VECTOR_B shall be of numeric type if VECTOR_A is of numeric type or of type logical if
29 VECTOR_A is of type logical. It shall be a rank-one array. It shall be of the
29 same size as VECTOR_A.30 **Result Characteristics.** If the arguments are of numeric type, the type and kind type parameter
31 of the result are those of the expression VECTOR_A * VECTOR_B determined by the types of the

1 arguments according to 7.1.4.2. If the arguments are of type logical, the result is of type logical
 2 with the kind type parameter of the expression VECTOR_A .AND. VECTOR_B according to
 3 7.1.4.2. The result is scalar.

4 **Result Value.**

5 *Case (i):* If VECTOR_A is of type integer or real, the result has the value SUM (VECTOR_
 6 A*VECTOR_B). If the vectors have size zero, the result has the value zero.

7 *Case (ii):* If VECTOR_A is of type complex, the result has the value SUM (CONJG (VEC-
 8 TOR_A)*VECTOR_B). If the vectors have size zero, the result has the value zero.

9 *Case (iii):* If VECTOR_A is of type logical, the result has the value ANY (VECTOR_A .AND.
 10 VECTOR_B). If the vectors have size zero, the result has the value false.

11 **Example.** DOT_PRODUCT ((/ 1, 2, 3 /), (/ 2, 3, 4 /)) has the value 20.

12 **13.7.55 DPROD (X, Y)**

13 **Description.** Double precision real product.

14 **Class.** Elemental function.

15 **Arguments.**

16 X shall be of type default real.

17 Y shall be of type default real.

18 **Result Characteristics.** Double precision real.

19 **Result Value.** The result has a value equal to a processor-dependent approximation to the
 20 product of X and Y.

21 **Example.** DPROD (-3.0, 2.0) has the value -6.0D0.

22 **13.7.56 DSHIFTL (I, J, SHIFT)**

23 **Description.** Performs a combined left shift.

24 **Class.** Elemental function.

25 **Arguments.**

26 I shall be of type integer or bits.

27 J shall be of the same type and kind as I.

28 SHIFT shall be of type integer. It shall be nonnegative and less than or equal to
 29 BIT_SIZE (I).

29 **Result Characteristics.** Same as I.

30 **Result Value.** The rightmost SHIFT bits of the result value are the same as the leftmost bits of
 31 J, and the remaining bits of the result value are the same as the rightmost bits of I. This is equal
 32 to IOR (SHIFTL (I, SHIFT), SHIFTR (J, BIT_SIZE (J)-SHIFT)).

33 **Examples.** DSHIFTL (Z'01234567', Z'89ABCDEF', 8) has the value Z'23456789'.

34 DSHIFTL (I, I, SHIFT) has the same result value as ISHFTC (I, SHIFT).

35 **13.7.57 DSHIFTR (I, J, SHIFT)**

1 **Description.** Performs a combined right shift.

2 **Class.** Elemental function.

3 **Arguments.**

4 I shall be of type integer or bits.

5 J shall be of the same type and kind as I.

6 SHIFT shall be of type integer. It shall be nonnegative and less than or equal to BIT_SIZE (I).

7 **Result Characteristics.** Same as I.

8 **Result Value.** The leftmost SHIFT bits of the result value are the same as the rightmost bits of I, and the remaining bits of the result value are the same as the leftmost bits of J. This is equal to IOR (SHIFTL (I, BIT_SIZE (I)–SHIFT), SHIFTR (J, SHIFT)).

11 **Examples.** DSHIFTR(Z'01234567', Z'89ABCDEF', 8) has the value Z'6789ABCD'.
 12 DSHIFTR (B'111', B'000', 2) has the value B'110'. DSHIFTR (I, I, SHIFT) has the same result
 13 value as ISHFTC (I,–SHIFT).

14 13.7.58 EOSHIFT (ARRAY, SHIFT [, BOUNDARY, DIM])

15 **Description.** Perform an end-off shift on an array expression of rank one or perform end-off
 16 shifts on all the complete rank-one sections along a given dimension of an array expression of rank
 17 two or greater. Elements are shifted off at one end of a section and copies of a boundary value
 18 are shifted in at the other end. Different sections may have different boundary values and may
 19 be shifted by different amounts and in different directions.

20 **Class.** Transformational function.

21 **Arguments.**

22 ARRAY may be of any type. It shall be an array.

23 SHIFT shall be of type integer and shall be scalar if ARRAY has rank one; otherwise,
 it shall be scalar or of rank $n - 1$ and of shape $(d_1, d_2, \dots, d_{DIM-1}, d_{DIM+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

BOUNDARY shall be of the same type and type parameters as ARRAY and shall be scalar if
 (optional) ARRAY has rank one; otherwise, it shall be either scalar or of rank $n - 1$ and
 of shape $(d_1, d_2, \dots, d_{DIM-1}, d_{DIM+1}, \dots, d_n)$. BOUNDARY may be omitted
 for the types in the following table and, in this case, it is as if it were present
 with the scalar value shown.

Type of ARRAY	Value of BOUNDARY
Integer	0
Real	0.0
Complex	(0.0, 0.0)
Logical	false
Character (<i>len</i>)	<i>len</i> blanks
Bits	B'0'

24 DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq DIM \leq n$,
 25 where n is the rank of ARRAY. If DIM is omitted, it is as if it were present
 with the value 1.

26 **Result Characteristics.** The result has the type, type parameters, and shape of ARRAY.

1 **Result Value.** Element (s_1, s_2, \dots, s_n) of the result has the value ARRAY $(s_1, s_2, \dots, s_{\text{DIM}-1},$
 2 $s_{\text{DIM}} + sh, s_{\text{DIM}+1}, \dots, s_n)$ where sh is SHIFT or SHIFT $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$
 3 provided the inequality LBOUND (ARRAY, DIM) $\leq s_{\text{DIM}} + sh \leq$ UBOUND (ARRAY, DIM)
 4 holds and is otherwise BOUNDARY or BOUNDARY $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$.

5 **Examples.**

6 *Case (i):* If V is the array [1, 2, 3, 4, 5, 6], the effect of shifting V end-off to the left by 3
 7 positions is achieved by EOSHIFT (V, SHIFT = 3), which has the value [4, 5, 6,
 8 0, 0, 0]; EOSHIFT (V, SHIFT = -2, BOUNDARY = 99) achieves an end-off shift
 9 to the right by 2 positions with the boundary value of 99 and has the value [99, 99,
 10 1, 2, 3, 4].

11 *Case (ii):* The rows of an array of rank two may all be shifted by the same amount or by
 12 different amounts and the boundary elements can be the same or different. If M is

13 the array $\begin{bmatrix} A & B & C \\ D & E & F \\ G & H & I \end{bmatrix}$, then the value of EOSHIFT (M, SHIFT = -1, BOUND-

14 ARY = '*', DIM = 2) is $\begin{bmatrix} * & A & B \\ * & D & E \\ * & G & H \end{bmatrix}$, and the value of EOSHIFT (M, SHIFT =

15 (/ -1, 1, 0 /), BOUNDARY = (/ '*', '/', '?' /), DIM = 2) is $\begin{bmatrix} * & A & B \\ E & F & / \\ G & H & I \end{bmatrix}$.

16 13.7.59 EPSILON (X)

17 **Description.** Returns a positive model number that is almost negligible compared to unity of
 18 the model representing numbers of the same type and kind type parameter as the argument.

19 **Class.** Inquiry function.

20 **Argument.** X shall be of type real. It may be a scalar or an array.

21 **Result Characteristics.** Scalar of the same type and kind type parameter as X.

22 **Result Value.** The result has the value b^{1-p} where b and p are as defined in 13.4 for the model
 23 representing numbers of the same type and kind type parameter as X.

24 **Example.** EPSILON (X) has the value 2^{-23} for real X whose model is as in Note 13.4.

25 13.7.60 ERF (X)

26 **Description.** Error function.

27 **Class.** Elemental function.

28 **Argument.** X shall be of type real.

29 **Result Characteristics.** Same as X.

30 **Result Value.** The result has a value equal to a processor-dependent approximation to the error
 31 function of X, $\frac{2}{\sqrt{\pi}} \int_0^X \exp(-t^2) dt$.

32 **Example.** ERF (1.0) has the value 0.843 (approximately).

33 13.7.61 ERFC (X)

1 **Description.** Complementary error function.

2 **Class.** Elemental function.

3 **Argument.** X shall be of type real.

4 **Result Characteristics.** Same as X.

5 **Result Value.** The result has a value equal to a processor-dependent approximation to the
6 complementary error function of X, that is, $1 - \text{ERF}(X)$.

7 **Example.** `ERFC(1.0)` has the value 0.157 (approximately).

8 **13.7.62 ERFC_SCALED (X)**

9 **Description.** Exponentially-scaled complementary error function.

10 **Class.** Elemental function.

11 **Argument.** X shall be of type real.

12 **Result Characteristics.** Same as X.

13 **Result Value.** The result has a value equal to a processor-dependent approximation to the
14 exponentially-scaled complementary error function of X, $\exp(X^2) \frac{2}{\sqrt{\pi}} \int_X^\infty \exp(-t^2) dt$.

15 **Example.** `ERFC_SCALED(20.0)` has the value 0.02817434874 (approximately).

NOTE 13.9

The complementary error function is asymptotic to $\exp(-X^2)/(X\sqrt{\pi})$. As such it underflows for $X > \approx 9$ when using IEEE single precision arithmetic. The exponentially-scaled complementary error function is asymptotic to $1/(X\sqrt{\pi})$. As such it does not underflow until $X > \text{HUGE}(X)/\sqrt{\pi}$.

16 **13.7.63 EXECUTE_COMMAND_LINE (COMMAND [, WAIT, EXITSTAT, CMDSTAT, CMDMSG])**

17 **Description.** Execute the command line specified by the string `COMMAND`. If the processor
18 supports command line execution, it shall support synchronous and may support asynchronous
19 execution of the command line.

20 **Class.** Subroutine.

21 **Arguments.**

22 `COMMAND` shall be of type default character and shall be scalar. It is an `INTENT(IN)`
argument. Its value is the command line to be executed. The interpretation is
processor-dependent.

23 `WAIT` (optional) shall be of type default logical and shall be scalar. It is an `INTENT(IN)` argu-
ment. If `WAIT` is present with the value `false`, and the processor supports asyn-
chronous execution of the command, the command is executed asynchronously;
otherwise it is executed synchronously.

24 `EXITSTAT` (optional) shall be of type default integer and shall be a scalar. It is an `INTENT(INOUT)`
argument. If the command is executed synchronously, it is assigned the value
of the processor-dependent exit status. Otherwise, the value of `EXITSTAT` is
unchanged.

CMDSTAT shall be of type default integer and shall be a scalar. It is an INTENT(OUT) (optional) argument. It is assigned the value -1 if the processor does not support command line execution, a processor-dependent positive value if an error condition occurs, or the value -2 if no error condition occurs but WAIT is present with the value false and the processor does not support asynchronous execution. Otherwise it is assigned the value 0.

CMDMSG shall be of type default character and shall be a scalar. It is an INTENT(INOUT) argument. If an error condition occurs, it is assigned a processor-dependent explanatory message. Otherwise, it is unchanged.

When the command is executed synchronously, EXECUTE_COMMAND_LINE returns after the command line has completed execution. Otherwise, EXECUTE_COMMAND_LINE returns without waiting.

If an error condition occurs and CMDSTAT is not present, execution of the program is terminated.

13.7.64 EXP (X)

Description. Exponential.

Class. Elemental function.

Argument. X shall be of type real or complex.

Result Characteristics. Same as X.

Result Value. The result has a value equal to a processor-dependent approximation to e^X . If X is of type complex, its imaginary part is regarded as a value in radians.

Example. EXP (1.0) has the value 2.7182818 (approximately).

13.7.65 EXPONENT (X)

Description. Returns the exponent part of the argument when represented as a model number.

Class. Elemental function.

Argument. X shall be of type real.

Result Characteristics. Default integer.

Result Value. The result has a value equal to the exponent e of the representation for the value of X in the model (13.4) that has the radix of X but no limits on exponent values, provided X is nonzero and e is within the range for default integers. If X has the value zero, the result has the value zero. If X is an IEEE infinity or NaN, the result has the value HUGE(0).

Examples. EXPONENT (1.0) has the value 1 and EXPONENT (4.1) has the value 3 for reals whose model is as in Note 13.4.

13.7.66 EXTENDS_TYPE_OF (A, MOLD)

Description. Inquires whether the dynamic type of A is an extension type (4.5.7) of the dynamic type of MOLD.

Class. Inquiry function.

Arguments.

1 A shall be an object of extensible type. If it is a pointer, it shall not have an
undefined association status.

2 MOLD shall be an object of extensible type. If it is a pointer, it shall not have an
undefined association status.

3 **Result Characteristics.** Default logical scalar.

4 **Result Value.** If MOLD is unlimited polymorphic and is either a disassociated pointer or
5 unallocated allocatable, the result is true; otherwise if A is unlimited polymorphic and is either a
6 disassociated pointer or unallocated allocatable, the result is false; otherwise the result is true if
7 and only if the dynamic type of A is an extension type of the dynamic type of MOLD.

NOTE 13.10

The dynamic type of a disassociated pointer or unallocated allocatable is its declared type.

8 **13.7.67 FINDLOC (ARRAY, VALUE, DIM, [,MASK, KIND, BACK]) or
FINDLOC (ARRAY, VALUE, [, MASK, KIND, BACK])**

9 **Description.** Determine the location of the first element of ARRAY identified by MASK along
10 dimension DIM having a value equal to VALUE.

11 **Class.** Transformational function.

12 **Arguments.**

13 ARRAY shall be of intrinsic type. It shall be an array.

14 VALUE shall be in type conformance with ARRAY, as specified in Table 7.11.

15 DIM shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
where n is the rank of ARRAY. The corresponding actual argument shall not
be an optional dummy argument.

16 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

17 KIND (optional) shall be a scalar integer initialization expression.

18 BACK (optional) shall be scalar and of type logical.

19 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
20 by the value of KIND; otherwise the kind type parameter is that of default integer type. If DIM
21 is absent, the result is an array of rank one and of size equal to the rank of ARRAY; otherwise,
22 the result is of rank $n - 1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$, where $(d_1, d_2, \dots,$
23 $d_n)$ is the shape of ARRAY.

24 **Result Value.**

25 *Case (i):* The result of FINDLOC (ARRAY, VALUE) is a rank-one array whose element
26 values are the values of the subscripts of an element of ARRAY whose value matches
27 VALUE. The i^{th} subscript returned lies in the range 1 to e_i , where e_i is the extent
28 of the i^{th} dimension of ARRAY. If no elements match VALUE or ARRAY has size
29 zero, all elements of the result are zero.

30 *Case (ii):* The result of FINDLOC (ARRAY, VALUE, MASK = MASK) is a rank-one array
31 whose element values are the values of the subscripts of an element of ARRAY,
32 corresponding to a true element of MASK, whose value matches VALUE. The i^{th}
33 subscript returned lies in the range 1 to e_i , where e_i is the extent of the i^{th} dimension
34 of ARRAY. If no elements match VALUE, ARRAY has size zero, or every element

1 of MASK has the value false, all elements of the result are zero.
 2 *Case (iii):* If ARRAY has rank one, the result of
 3 FINDLOC (ARRAY, VALUE, DIM=DIM [, MASK = MASK]) is a scalar whose
 4 value is equal to that of the first element of FINDLOC (ARRAY, VALUE [, MASK
 5 = MASK]). Otherwise, the value of element ($s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n$
 6) of the result is equal to FINDLOC (ARRAY ($s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots,$
 7 s_n), VALUE, DIM=1, [, MASK = MASK ($s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n$))).
 8 If both ARRAY and VALUE are of type logical, the comparison is performed as
 9 array-element .EQV. VALUE; otherwise, the comparison is performed as array-
 10 element == VALUE. If the value of the comparison is true, array-element matches
 11 VALUE.
 12 If only one element matches VALUE, that element's subscripts are returned. Oth-
 13 erwise, if more than one element matches VALUE and BACK is absent or present
 14 with the value false, the element whose subscripts are returned is the first such
 15 element, taken in array element order. If BACK is present with the value true,
 16 the element whose subscripts are returned is the last such element, taken in array
 17 element order.

18 Examples.

19 *Case (i):* The value of FINDLOC ((/2, 6, 4, 6, /), VALUE=6) is [2], and the value of FIND-
 20 LOC ((/2, 6, 4, 6 /), VALUE=6, BACK=.TRUE.) is [4].

21 *Case (ii):* If A has the value $\begin{bmatrix} 0 & -5 & 7 & 7 \\ 3 & 4 & -1 & 2 \\ 1 & 5 & 6 & 7 \end{bmatrix}$, and M has the value $\begin{bmatrix} T & T & F & T \\ T & T & F & T \\ T & T & F & T \end{bmatrix}$,
 22 FINDLOC (A, 7, MASK = M) has the value [1, 4] and
 23 FINDLOC (A, 7, MASK = M, BACK = .TRUE.) has the value [3, 4]. Note that
 24 this is independent of the declared lower bounds for A.

25 *Case (iii):* The value of FINDLOC ([2, 6, 4], VALUE=6, DIM=1) is 2. If B has the value
 26 $\begin{bmatrix} 1 & 2 & -9 \\ 2 & 2 & 6 \end{bmatrix}$, FINDLOC (B, VALUE=2, DIM=1) has the value [2, 1, 0] and
 27 FINDLOC (B, VALUE=2, DIM=2) has the value [2, 1]. Note that this is true even
 28 if B has a declared lower bound other than 1.

29 13.7.68 FLOOR (A [, KIND])

30 **Description.** Returns the greatest integer less than or equal to its argument.

31 **Class.** Elemental function.

32 **Arguments.**

33 A shall be of type real.

34 KIND (optional) shall be a scalar integer initialization expression.

35 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
 36 by the value of KIND; otherwise, the kind type parameter is that of default integer type.

37 **Result Value.** The result has a value equal to the greatest integer less than or equal to A.

38 **Examples.** FLOOR (3.7) has the value 3. FLOOR (-3.7) has the value -4.

39 13.7.69 FORM_TEAM (TEAM, IMAGES)

1 **Description.** Form a team of images.

2 **Class.** Collective subroutine.

J3 internal note

Unresolved Technical Issue 063

I don't see why FORM_TEAM needs synchronisation – that seems completely unnecessary (it has no co-array argument, no need for communication – its actual arguments are all simple local variables).

If this needs synchronisation, why does SYNC IMAGES not need synchronisation?

3 **Arguments.**

4 TEAM shall be a scalar of type IMAGE_TEAM (13.8.3.7). It is an INTENT(OUT) argument.

5 IMAGES shall be of rank one and type integer. It is an INTENT(IN) argument that specifies the image indices of the team members, and shall have the same value on all images of the team. All the elements shall have values in the range 1, . . . , NUM_IMAGES () and there shall be no repeated values. It shall not have size zero.

6 **Example.**

7 The following code fragment splits images into two equal groups and implicitly synchronizes
8 each of the teams if there are two or more images. If there is only one image, that image
9 becomes the only team member.

```
10 USE, INTRINSIC :: ISO_FORTRAN_ENV
11 INTEGER :: I
12 TYPE(IMAGE_TEAM) :: TEAM
13 IF (THIS_IMAGE() <= NUM_IMAGES() / 2) THEN
14     CALL FORM_TEAM(TEAM, [(I, I=1, NUM_IMAGES() / 2)])
15 ELSE
16     CALL FORM_TEAM(TEAM, [(I, I=NUM_IMAGES() / 2 + 1, NUM_IMAGES())])
17 END IF
```

18 13.7.70 FRACTION (X)

19 **Description.** Returns the fractional part of the model representation of the argument value.

20 **Class.** Elemental function.

21 **Argument.** X shall be of type real.

22 **Result Characteristics.** Same as X.

23 **Result Value.** The result has the value $X \times b^{-e}$, where b and e are as defined in 13.4 for the
24 representation of X in the model that has the radix of X but no limits on exponent values. If X
25 has the value zero, or is an IEEE infinity or NaN, the result has the same value as X.

26 **Example.** FRACTION (3.0) has the value 0.75 for reals whose model is as in Note 13.4.

27 13.7.71 GAMMA (X)

Description. Gamma function.

1

2

Class. Elemental function.

3

Argument. X shall be of type real. Its value shall not be a negative integer or zero.

4

Result Characteristics. Same as X.

5

6

Result Value. The result has a value equal to a processor-dependent approximation to the gamma function of X, $\Gamma(X) = \int_0^{\infty} t^{X-1} \exp(-t) dt$.

7

Example. GAMMA (1.0) has the value 1.000 (approximately).

8

13.7.72 GET_COMMAND ([COMMAND, LENGTH, STATUS])

9

Description. Returns the entire command by which the program was invoked.

10

Class. Subroutine.

11

Arguments.

COMMAND
(optional)

shall be scalar and of type default character. It is an INTENT(OUT) argument. It is assigned the entire command by which the program was invoked. If the command cannot be determined, COMMAND is assigned all blanks.

12

LENGTH
(optional)

shall be scalar and of type default integer. It is an INTENT(OUT) argument. It is assigned the significant length of the command by which the program was invoked. The significant length may include trailing blanks if the processor allows commands with significant trailing blanks. This length does not consider any possible truncation or padding in assigning the command to the COMMAND argument; in fact the COMMAND argument need not even be present. If the command length cannot be determined, a length of 0 is assigned.

13

STATUS
(optional)

shall be scalar and of type default integer. It is an INTENT(OUT) argument. It is assigned the value -1 if the COMMAND argument is present and has a length less than the significant length of the command. It is assigned a processor-dependent positive value if the command retrieval fails. Otherwise it is assigned the value 0.

14

13.7.73 GET_COMMAND_ARGUMENT (NUMBER [, VALUE, LENGTH, STATUS])

15

Description. Returns a command argument.

16

17

Class. Subroutine.

18

Arguments.

19

NUMBER

shall be scalar and of type default integer. It is an INTENT(IN) argument.

20

It specifies the number of the command argument that the other arguments give information about. Useful values of NUMBER are those between 0 and the argument count returned by the COMMAND_ARGUMENT_COUNT intrinsic. Other values are allowed, but will result in error status return (see below).

Command argument 0 is defined to be the command name by which the program was invoked if the processor has such a concept. It is allowed to call the GET_COMMAND_ARGUMENT procedure for command argument number 0, even if the processor does not define command names or other command arguments.

The remaining command arguments are numbered consecutively from 1 to the argument count in an order determined by the processor.

VALUE
(optional) shall be scalar and of type default character. It is an INTENT(OUT) argument. It is assigned the value of the command argument specified by NUMBER. If the command argument value cannot be determined, VALUE is assigned all blanks.

LENGTH
(optional) shall be scalar and of type default integer. It is an INTENT(OUT) argument. It is assigned the significant length of the command argument specified by NUMBER. The significant length may include trailing blanks if the processor allows command arguments with significant trailing blanks. This length does not consider any possible truncation or padding in assigning the command argument value to the VALUE argument; in fact the VALUE argument need not even be present. If the command argument length cannot be determined, a length of 0 is assigned.

STATUS
(optional) shall be scalar and of type default integer. It is an INTENT(OUT) argument. It is assigned the value -1 if the VALUE argument is present and has a length less than the significant length of the command argument specified by NUMBER. It is assigned a processor-dependent positive value if the argument retrieval fails. Otherwise it is assigned the value 0.

NOTE 13.11

One possible reason for failure is that NUMBER is negative or greater than COMMAND_ARGUMENT_COUNT().

Example.

```

Program echo
  integer :: i
  character :: command*32, arg*128
  call get_command_argument(0, command)
  write (*,*) "Command name is: ", command
  do i = 1 , command_argument_count()
    call get_command_argument(i, arg)
    write (*,*) "Argument ", i, " is ", arg
  end do
end program echo

```

13.7.74 GET_ENVIRONMENT_VARIABLE (NAME [, VALUE, LENGTH, STATUS, TRIM_NAME])

Description. Gets the value of an environment variable.

Class. Subroutine.

1 **Arguments.**

2 NAME shall be scalar and of type default character. It is an INTENT(IN) argument.
The interpretation of case is processor dependent

3 VALUE shall be a scalar of type default character. It is an INTENT(OUT) argument.
(optional) It is assigned the value of the environment variable specified by NAME. VALUE
is assigned all blanks if the environment variable does not exist or does not have
a value or if the processor does not support environment variables.

4 LENGTH shall be a scalar of type default integer. It is an INTENT(OUT) argument. If
(optional) the specified environment variable exists and has a value, LENGTH is set to
the length of that value. Otherwise LENGTH is set to 0.

5 STATUS shall be scalar and of type default integer. It is an INTENT(OUT) argument.
(optional) If the environment variable exists and either has no value or its value is suc-
cessfully assigned to VALUE, STATUS is set to zero. STATUS is set to -1
if the VALUE argument is present and has a length less than the significant
length of the environment variable. It is assigned the value 1 if the specified
environment variable does not exist, or 2 if the processor does not support envi-
ronment variables. Processor-dependent values greater than 2 may be returned
for other error conditions.

6 TRIM_NAME shall be a scalar of type logical. It is an INTENT(IN) argument. If TRIM_-
(optional) NAME is present with the value false then trailing blanks in NAME are consid-
ered significant if the processor supports trailing blanks in environment variable
names. Otherwise trailing blanks in NAME are not considered part of the en-
vironment variable's name.

7 **13.7.75 HUGE (X)**

8 **Description.** Returns the largest number of the model representing numbers of the same type
9 and kind type parameter as the argument.

10 **Class.** Inquiry function.

11 **Argument.** X shall be of type integer, real, or bits. It may be a scalar or an array.

12 **Result Characteristics.** Scalar of the same type and kind type parameter as X.

13 **Result Value.** The result has the value $r^q - 1$ if X is of type integer and $(1 - b^{-p})b^{e_{\max}}$ if X is
14 of type real, where r , q , b , p , and e_{\max} are as defined in 13.4 for the model representing numbers
15 of the same type and kind type parameter as X. If X is of type bits, the result value has all of its
16 bits set to 1.

17 **Example.** HUGE (X) has the value $(1 - 2^{-24}) \times 2^{127}$ for real X whose model is as in Note 13.4.

18 **13.7.76 HYPOT (X, Y)**

19 **Description.** Euclidean distance function.

20 **Class.** Elemental function.

21 **Arguments.**

22 X shall be of type real.

23 Y shall be of type real with the same kind type parameter as X.

1 **Result Characteristics.** Same as X.

2 **Result Value.** The result has a value equal to a processor-dependent approximation to the
3 Euclidean distance, $\sqrt{X^2 + Y^2}$, without undue overflow or underflow.

4 **Example.** HYPOT (2.0, 1.0) has the value 2.236 (approximately).

5 13.7.77 IACHAR (C [, KIND])

6 **Description.** Returns the position of a character in the ASCII collating sequence. This is the
7 inverse of the ACHAR function.

8 **Class.** Elemental function.

9 **Arguments.**

10 C shall be of type character and of length one.

11 KIND (optional) shall be a scalar integer initialization expression.

12 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
13 by the value of KIND; otherwise, the kind type parameter is that of default integer type.

14 **Result Value.** If C is in the collating sequence defined by the codes specified in ISO/IEC 646:1991
15 (International Reference Version), the result is the position of C in that sequence and satisfies the
16 inequality $(0 \leq \text{IACHAR}(C) \leq 127)$. A processor-dependent value is returned if C is not in the
17 ASCII collating sequence. The results are consistent with the LGE, LGT, LLE, and LLT lexical
18 comparison functions. For example, if LLE (C, D) is true, IACHAR (C) <= IACHAR (D) is true
19 where C and D are any two characters representable by the processor.

20 **Example.** IACHAR ('X') has the value 88.

21 13.7.78 IALL (ARRAY, DIM [, MASK]) or IALL (ARRAY, [MASK])

22 **Description.** Bitwise AND of all the elements of ARRAY along dimension DIM corresponding
23 to the true elements of MASK.

24 **Class.** Transformational function.

25 **Arguments.**

26 ARRAY shall be of type integer or bits. It shall be an array.

27 DIM shall be a scalar of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where
28 n is the rank of ARRAY. The corresponding actual argument shall not be an
29 optional dummy argument.

30 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

31 **Result Characteristics.** The result is of the same type and kind type parameter as ARRAY.
32 It is scalar if DIM is absent or if ARRAY has rank one; otherwise, the result is an array of rank
33 $n - 1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of
34 ARRAY.

35 **Result Value.**

36 *Case (i):* If ARRAY has size zero the result value is equal to NOT (INT (0, KIND (ARRAY)))
 if ARRAY is of type integer, and NOT (BITS (B'0', KIND (ARRAY))) if array
 is of type bits. Otherwise, the result of IALL (ARRAY) has a value equal to the

1 bitwise AND of all the elements of ARRAY.

2 *Case (ii):* The result of IALL (ARRAY, MASK=MASK) has a value equal to
3 IALL (PACK (ARRAY, MASK)).

4 *Case (iii):* The result of IALL (ARRAY, DIM=DIM [, MASK=MASK]) has a value equal to
5 that of IALL (ARRAY, [, MASK=MASK]) if ARRAY has rank one. Otherwise,
6 the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$ of the result is equal to
7 IALL (ARRAY $(s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)$ [, MASK = MASK $(s_1, s_2,$
8 $\dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)$]).

9 **Examples.** IALL ([B'1110', B'1101', B'1011']) has the value B'1000'. IALL ([B'1110', B'1101',
10 B'1011'], MASK=[.true., .false., .true]) has the value B'1010'. IALL([14, 13, 11]) has the value 8.
11 IALL([14, 13, 11], MASK=[.true., .false., .true]) has the value 10.

12 13.7.79 IAND (I, J)

13 **Description.** Performs a bitwise AND.

14 **Class.** Elemental function.

15 **Arguments.**

16 I shall be of type integer or bits.

J shall be of type integer or bits. If both I and J are of type integer, they shall
17 have the same kind type parameter; otherwise BITS_KIND (I) shall be equal
to BITS_KIND (J).

18 **Result Characteristics.** If I and J have the same type, the result characteristics are the same
19 as I. Otherwise, the result characteristics are the same as those of the argument with integer type.

20 **Result Value.** The result has the value obtained by combining I and J bit-by-bit according to
21 the following truth table:

I	J	IAND (I, J)
1	1	1
1	0	0
0	1	0
0	0	0

22 The model for the interpretation of an integer value as a sequence of bits is in 13.3.

23 **Examples.** IAND (1, 3) has the value 1. IAND (Z'12345678', Z'0000FFFF') has the value
24 Z'00005678'.

25 13.7.80 IANY (ARRAY, DIM [, MASK]) or IANY (ARRAY, [MASK])

26 **Description.** Bitwise OR of all the elements of ARRAY along dimension DIM corresponding to
27 the true elements of MASK.

28 **Class.** Transformational function.

29 **Arguments.**

30 ARRAY shall be of type integer or bits. It shall be an array.

DIM (optional) shall be a scalar and of type integer with a value in the range $1 \leq DIM \leq n$, where n is the rank of ARRAY. The corresponding actual argument shall not be an optional dummy argument.

MASK (optional) shall be of type logical and shall be conformable with ARRAY.

Result Characteristics. The result is of the same type and kind type parameter as ARRAY. It is scalar if DIM is absent or if ARRAY has rank one; otherwise, the result is an array of rank $n - 1$ and shape $(d_1, d_2, \dots, d_{DIM-1}, d_{DIM+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

Result Value.

Case (i): The result of IANY (ARRAY) is the bitwise OR of all the elements of ARRAY. If ARRAY has size zero the result value is equal to zero if ARRAY is of type integer, and BITS (B'0', KIND(ARRAY)) otherwise.

Case (ii): The result of IANY (ARRAY, MASK=MASK) has a value equal to IANY (PACK (ARRAY, MASK)).

Case (iii): The result of IANY (ARRAY, DIM=DIM [, MASK=MASK]) has a value equal to that of IANY (ARRAY, [, MASK=MASK]) if ARRAY has rank one. Otherwise, the value of element $(s_1, s_2, \dots, s_{DIM-1}, s_{DIM+1}, \dots, s_n)$ of the result is equal to IANY (ARRAY ($s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n$) [, MASK = MASK($s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n$)])).

Examples. IANY ([B'1110', B'1101', B'1000']) has the value B'1111'. IANY ([B'1110', B'1101', B'1000'], MASK=[.true., .false., .true]) has the value B'1110'. IANY ([14, 13, 8]) has the value 15. IANY ([14, 13, 8], MASK=[.true., .false., .true]) has the value 14.

13.7.81 IBCLR (I, POS)

Description. Clears one bit to zero.

Class. Elemental function.

Arguments.

I shall be of type integer or bits.

POS shall be of type integer. It shall be nonnegative and less than BIT_SIZE (I).

Result Characteristics. Same as I.

Result Value. The result has the value of the sequence of bits of I, except that bit POS is zero. The model for the interpretation of an integer value as a sequence of bits is in 13.3.

Examples. IBCLR (14, 1) has the value 12. If V has the value [1, 2, 3, 4], the value of IBCLR (POS = V, I = 31) is [29, 27, 23, 15]. IBCLR (B'11111', 3) has the value B'10111'.

13.7.82 IBITS (I, POS, LEN)

Description. Extracts a sequence of bits.

Class. Elemental function.

Arguments.

I shall be of type integer or bits.

1 POS shall be of type integer. It shall be nonnegative and POS + LEN shall be less
than or equal to BIT_SIZE (I).

2 LEN shall be of type integer and nonnegative.

3 **Result Characteristics.** Same as I.

4 **Result Value.** The result has the value of the sequence of LEN bits in I beginning at bit POS,
5 right-adjusted and with all other bits zero. The model for the interpretation of an integer value
6 as a sequence of bits is in 13.3.

7 **Examples.** IBITS (14, 1, 3) has the value 7. IBITS (Z'ABCD', 4, 8) has the value Z'00BC'.

8 13.7.83 IBSET (I, POS)

9 **Description.** Sets one bit to one.

10 **Class.** Elemental function.

11 **Arguments.**

12 I shall be of type integer or bits.

13 POS shall be of type integer. It shall be nonnegative and less than BIT_SIZE (I).

14 **Result Characteristics.** Same as I.

15 **Result Value.** The result has the value of the sequence of bits of I, except that bit POS is one.
16 The model for the interpretation of an integer value as a sequence of bits is in 13.3.

17 **Examples.** IBSET (12, 1) has the value 14. If V has the value [1, 2, 3, 4], the value of IB-
18 SET (POS = V, I = 0) is [2, 4, 8, 16]. IBSET (B'00000', 3) has the value B'01000'.

19 13.7.84 ICHAR (C [, KIND])

20 **Description.** Returns the position of a character in the processor collating sequence associated
21 with the kind type parameter of the character. This is the inverse of the CHAR function.

22 **Class.** Elemental function.

23 **Arguments.**

24 C shall be of type character and of length one. Its value shall be that of a character
capable of representation in the processor.

25 KIND (optional) shall be a scalar integer initialization expression.

26 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
27 by the value of KIND; otherwise, the kind type parameter is that of default integer type.

28 **Result Value.** The result is the position of C in the processor collating sequence associated with
29 the kind type parameter of C and is in the range $0 \leq \text{ICHAR}(C) \leq n - 1$, where n is the number
30 of characters in the collating sequence. For any characters C and D capable of representation in
31 the processor, $C \leq D$ is true if and only if $\text{ICHAR}(C) \leq \text{ICHAR}(D)$ is true and $C == D$ is
32 true if and only if $\text{ICHAR}(C) == \text{ICHAR}(D)$ is true.

33 **Example.** ICHAR ('X') has the value 88 on a processor using the ASCII collating sequence for
34 the default character type.

1 13.7.85 IEOR (I, J)

2 **Description.** Performs a bitwise exclusive OR.

3 **Class.** Elemental function.

4 **Arguments.**

5 I shall be of type integer or bits.

J shall be of type integer or bits. If both I and J are of type integer, they shall have the same kind type parameter; otherwise BITS_KIND (I) shall be equal to BITS_KIND (J).

7 **Result Characteristics.** If I and J have the same type, the result characteristics are the same as I. Otherwise, the result characteristics are the same as those of the argument with integer type.

9 **Result Value.** The result has the value obtained by combining I and J bit-by-bit according to the following truth table:

I	J	IEOR (I, J)
1	1	0
1	0	1
0	1	1
0	0	0

11 The model for the interpretation of an integer value as a sequence of bits is in 13.3.

12 **Examples.** IEOR (1, 3) has the value 2. IEOR (O'5', B'110') has the value O'3'.

13 13.7.86 IMAGE_INDEX (CO_ARRAY, SUB)

14 **Description.** Returns the index of the image corresponding to the co-subscripts SUB for CO-
15 ARRAY.

16 **Class.** Inquiry function.

17 **Arguments.**

18 CO_ARRAY shall be a co-array of any type.

19 SUB shall be a rank-one array of size equal to the co-rank of CO_ARRAY and shall be of type integer.

20 **Result Characteristics.** Default integer scalar.

21 **Result Value.** If the value of SUB is a valid sequence of co-subscripts for CO_ARRAY, the result
22 is the index of the corresponding image. Otherwise, the result is zero.

23 **Examples.** If A is declared A [0:*], IMAGE_INDEX (A, (/0/)) has the value 1. If B is declared
24 REAL B (10, 20) [10, 0:9, 0:*], IMAGE_INDEX (B, [3, 1, 2]) has the value 213 (on any image).

NOTE 13.12

For an example of a module that implements a function similar to the intrinsic THIS_IMAGE, see subclause C.10.1.

13.7.87 INDEX (STRING, SUBSTRING [, BACK, KIND])

1 **Description.** Returns the starting position of a substring within a string.

2 **Class.** Elemental function.

3 **Arguments.**

4 STRING shall be of type character.

5 SUBSTRING shall be of type character with the same kind type parameter as STRING.

6 BACK (optional) shall be of type logical.

7 KIND (optional) shall be a scalar integer initialization expression.

8 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
9 by the value of KIND; otherwise the kind type parameter is that of default integer type.

10 **Result Value.**

11 *Case (i):* If BACK is absent or has the value false, the result is the minimum positive value
12 of I such that $\text{STRING}(I : I + \text{LEN}(\text{SUBSTRING}) - 1) = \text{SUBSTRING}$ or zero if
13 there is no such value. Zero is returned if $\text{LEN}(\text{STRING}) < \text{LEN}(\text{SUBSTRING})$
14 and one is returned if $\text{LEN}(\text{SUBSTRING}) = 0$.

15 *Case (ii):* If BACK is present with the value true, the result is the maximum value of
16 I less than or equal to $\text{LEN}(\text{STRING}) - \text{LEN}(\text{SUBSTRING}) + 1$ such that
17 $\text{STRING}(I : I + \text{LEN}(\text{SUBSTRING}) - 1) = \text{SUBSTRING}$ or zero if there is
18 no such value. Zero is returned if $\text{LEN}(\text{STRING}) < \text{LEN}(\text{SUBSTRING})$ and
19 $\text{LEN}(\text{STRING}) + 1$ is returned if $\text{LEN}(\text{SUBSTRING}) = 0$.

20 **Examples.** INDEX ('FORTRAN', 'R') has the value 3.

21 INDEX ('FORTRAN', 'R', BACK = .TRUE.) has the value 5.

22 13.7.88 INT (A [, KIND])

23 **Description.** Convert to integer type.

24 **Class.** Elemental function.

25 **Arguments.**

26 A shall be of type integer, real, complex, or bits.

27 KIND (optional) shall be a scalar integer initialization expression.

28 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
29 by the value of KIND; otherwise, the kind type parameter is that of default integer type.

30 **Result Value.**

31 *Case (i):* If A is of type integer, $\text{INT}(A) = A$.

32 *Case (ii):* If A is of type real, there are two cases: if $|A| < 1$, INT (A) has the value 0; if
33 $|A| \geq 1$, INT (A) is the integer whose magnitude is the largest integer that does
34 not exceed the magnitude of A and whose sign is the same as the sign of A.

35 *Case (iii):* If A is of type complex, $\text{INT}(A) = \text{INT}(\text{REAL}(A, \text{KIND}(A)))$.

36 *Case (iv):* If A is of type bits, the result has the integer value such that
37 $\text{BITS}(\text{INT}(A, \text{kind})) = \text{BITS}(A, \text{BITS_KIND}(\text{INT}(0, \text{kind})))$, where *kind* is the
38 kind of the result. If the value specified by the model in 13.3 for interpreting the
39 bits as an integer is a valid value for the result, it has that value.

1 **Example.** INT (-3.7) has the value -3.

2 **13.7.89 IOR (I, J)**

3 **Description.** Performs a bitwise inclusive OR.

4 **Class.** Elemental function.

5 **Arguments.**

6 I shall be of type integer or bits.

7 J shall be of type integer or bits. If both I and J are of type integer, they shall have the same kind type parameter; otherwise BITS_KIND (I) shall be equal to BITS_KIND (J).

8 **Result Characteristics.** If I and J have the same type, the result characteristics are the same as I. Otherwise, the result characteristics are the same as those of the argument with integer type.

9 **Result Value.** The result has the value obtained by combining I and J bit-by-bit according to the following truth table:

I	J	IOR (I, J)
1	1	1
1	0	1
0	1	1
0	0	0

12 The model for the interpretation of an integer value as a sequence of bits is in 13.3.

13 **Examples.** IOR (5, 3) has the value 7. IOR (Z'1234, Z'00FF') has the value Z'12FF'.

14 **13.7.90 IPARITY (ARRAY, DIM [, MASK]) or IPARITY (ARRAY, [MASK])**

15 **Description.** Bitwise exclusive OR of all the elements of ARRAY along dimension DIM corresponding to the true elements of MASK.

16 **Class.** Transformational function.

17 **Arguments.**

18 ARRAY shall be of type integer or bits. It shall be an array.

19 DIM shall be a scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where n is the rank of ARRAY. The corresponding actual argument shall not be an optional dummy argument.

20 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

21 **Result Characteristics.** The result is of the same type and kind type parameter as ARRAY. It is scalar if DIM is absent; otherwise, the result has rank $n-1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

22 **Result Value.**

23 *Case (i):* The result of IPARITY (ARRAY) has a value equal to the bitwise exclusive OR of all the elements of ARRAY. If ARRAY has size zero the result has the value zero if ARRAY is of type integer, and BITS (B'0', KIND (ARRAY)) otherwise.

1 *Case (ii):* The result of IPARITY (ARRAY, MASK=MASK) has a value equal to that of
2 IPARITY (PACK (ARRAY, MASK)).

3 *Case (iii):* The result of IPARITY (ARRAY, DIM=DIM [, MASK=MASK]) has a value equal
4 to that of IPARITY (ARRAY, [, MASK=MASK]) if ARRAY has rank one. Oth-
5 erwise, the value of element $(s_1, s_2, \dots, s_{DIM-1}, s_{DIM+1}, \dots, s_n)$ of the result is
6 equal to IPARITY (ARRAY $(s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n)$ [, MASK =
7 MASK $(s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n)$]).

8 **Examples.** IPARITY ([B'1110', B'1101', B'1000']) has the value B'1011'.
9 IPARITY ([B'1110', B'1101', B'1000'], MASK=[.true., .false., .true]) has the value B'0110'.
10 IPARITY ([14, 13, 8]) has the value 11. IPARITY ([14, 13, 8], MASK=[.true., .false., .true]) has
11 the value 6.

12 13.7.91 ISHFT (I, SHIFT)

13 **Description.** Performs a logical shift.

14 **Class.** Elemental function.

15 **Arguments.**

16 I shall be of type integer or bits.

17 SHIFT shall be of type integer. The absolute value of SHIFT shall be less than or
18 equal to BIT_SIZE (I).

19 **Result Characteristics.** Same as I.

20 **Result Value.** The result has the value obtained by shifting the bits of I by SHIFT positions.
21 If SHIFT is positive, the shift is to the left; if SHIFT is negative, the shift is to the right; and
22 if SHIFT is zero, no shift is performed. Bits shifted out from the left or from the right, as
23 appropriate, are lost. Zeros are shifted in from the opposite end. The model for the interpretation
24 of an integer value as a sequence of bits is in 13.3.

25 **Examples.** ISHFT (3, 1) has the value 6. ISHFT (B'00000011', 1) has the value B'00000110'.

26 13.7.92 ISHFTC (I, SHIFT [, SIZE])

27 **Description.** Performs a circular shift of the rightmost bits.

28 **Class.** Elemental function.

29 **Arguments.**

30 I shall be of type integer or bits.

31 SHIFT shall be of type integer. The absolute value of SHIFT shall be less than or
32 equal to SIZE.

33 SIZE (optional) shall be of type integer. The value of SIZE shall be positive and shall not
34 exceed BIT_SIZE (I). If SIZE is absent, it is as if it were present with the value
35 of BIT_SIZE (I).

36 **Result Characteristics.** Same as I.

37 **Result Value.** The result has the value obtained by shifting the SIZE rightmost bits of I circularly
38 by SHIFT positions. If SHIFT is positive, the shift is to the left; if SHIFT is negative, the shift is
39 to the right; and if SHIFT is zero, no shift is performed. No bits are lost. The unshifted bits are

1 unaltered. The model for the interpretation of an integer value as a sequence of bits is in 13.3.

2 **Examples.** ISHFTC (3, 2, 3) has the value 5. ISHFTC (Z'ABCD', 4, 8) has the value Z'ABDC'.

3 **13.7.93 IS_CONTIGUOUS (A)**

4 **Description.** Determine whether an object is contiguous (5.3.6).

5 **Class.** Inquiry function.

6 **Argument.** A may be of any type. It shall be an assumed-shape array or an array pointer. If it
7 is a pointer it shall be associated.

8 **Result Characteristics.** Default logical scalar.

9 **Result Value.** The result has the value true if A is contiguous, and false otherwise.

10 **Example.** After the pointer assignment AP => TARGET (1:10:2), IS_CONTIGUOUS (AP) will
11 have the value false.

12 **13.7.94 IS_IOSTAT_END (I)**

13 **Description.** Determine whether a value indicates an end-of-file condition.

14 **Class.** Elemental function.

15 **Argument.** I shall be of type integer.

16 **Result Characteristics.** Default logical.

17 **Result Value.** The result has the value true if and only if I is a value for the *scalar-int-variable*
18 in an IOSTAT= specifier (9.10.4) that would indicate an end-of-file condition.

19 **13.7.95 IS_IOSTAT_EOR (I)**

20 **Description.** Determine whether a value indicates an end-of-record condition.

21 **Class.** Elemental function.

22 **Argument.** I shall be of type integer.

23 **Result Characteristics.** Default logical.

24 **Result Value.** The result has the value true if and only if I is a value for the *scalar-int-variable*
25 in an IOSTAT= specifier (9.10.4) that would indicate an end-of-record condition.

26 **13.7.96 KIND (X)**

27 **Description.** Returns the value of the kind type parameter of X.

28 **Class.** Inquiry function.

29 **Argument.** X may be of any intrinsic type. It may be a scalar or an array.

30 **Result Characteristics.** Default integer scalar.

31 **Result Value.** The result has a value equal to the kind type parameter value of X.

32 **Example.** KIND (0.0) has the kind type parameter value of default real.

1 13.7.97 LBOUND (ARRAY [, DIM, KIND])

2 **Description.** Returns all the lower bounds or a specified lower bound of an array.

3 **Class.** Inquiry function.

4 **Arguments.**

5 ARRAY may be of any type. It shall be an array. It shall not be an unallocated allocatable or a pointer that is not associated.

6 DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where n is the rank of ARRAY. The corresponding actual argument shall not be an optional dummy argument.

7 KIND (optional) shall be a scalar integer initialization expression.

8 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified by the value of KIND; otherwise the kind type parameter is that of default integer type. The result is scalar if DIM is present; otherwise, the result is an array of rank one and size n , where n is the rank of ARRAY.

12 **Result Value.**

13 *Case (i):* If ARRAY is a whole array or array structure component and either ARRAY is an assumed-size array of rank DIM or dimension DIM of ARRAY has nonzero extent, LBOUND (ARRAY, DIM) has a value equal to the lower bound for subscript DIM of ARRAY. Otherwise the result value is 1.

17 *Case (ii):* LBOUND (ARRAY) has a value whose i th element is equal to LBOUND (ARRAY, i), for $i = 1, 2, \dots, n$, where n is the rank of ARRAY.

19 **Examples.** If A is declared by the statement

20 REAL A (2:3, 7:10)

21 then LBOUND (A) is [2, 7] and LBOUND (A, DIM=2) is 7.

22 13.7.98 LEADZ (I)

23 **Description.** Count the number of leading zero bits.

24 **Class.** Elemental function.

25 **Argument.** I shall be of type integer of bits.

26 **Result Characteristics.** Default integer.

27 **Result Value.** If all of the bits of I are zero, the result has the value BIT_SIZE (I). Otherwise, the result has the value $\text{BIT_SIZE (I)} - 1 - k$, where k is the position of the leftmost 1 bit in I. The model for the interpretation of an integer value as a sequence of bits is in 13.3.

30 **Examples.** LEADZ (B'001101000') has the value 2. LEADZ (1) has the value 31 if BIT_SIZE (1) has the value 32.

32 13.7.99 LEN (STRING [, KIND])

33 **Description.** Returns the length of a character entity.

34 **Class.** Inquiry function.

1 **Arguments.**

2 STRING shall be of type character. It may be a scalar or an array. If it is an unallocated
3 allocatable or a pointer that is not associated, its length type parameter shall
4 not be deferred.

5 KIND (optional) shall be a scalar integer initialization expression.

6 **Result Characteristics.** Integer scalar. If KIND is present, the kind type parameter is that
7 specified by the value of KIND; otherwise the kind type parameter is that of default integer type.

8 **Result Value.** The result has a value equal to the number of characters in STRING if it is scalar
9 or in an element of STRING if it is an array.

10 **Example.** If C is declared by the statement

11 CHARACTER (11) C (100)

12 LEN (C) has the value 11.

13 **13.7.100 LEN_TRIM (STRING [, KIND])**

14 **Description.** Returns the length of the character argument without counting trailing blank
15 characters.

16 **Class.** Elemental function.

17 **Arguments.**

18 STRING shall be of type character.

19 KIND (optional) shall be a scalar integer initialization expression.

20 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
21 by the value of KIND; otherwise the kind type parameter is that of default integer type.

22 **Result Value.** The result has a value equal to the number of characters remaining after any
23 trailing blanks in STRING are removed. If the argument contains no nonblank characters, the
24 result is zero.

25 **Examples.** LEN_TRIM (' A B ') has the value 4 and LEN_TRIM (' ') has the value 0.

26 **13.7.101 LGE (STRING_A, STRING_B)**

27 **Description.** Test whether a string is lexically greater than or equal to another string, based on
28 the ASCII collating sequence.

29 **Class.** Elemental function.

30 **Arguments.**

31 STRING_A shall be of type default character of type ASCII character.

32 STRING_B shall be of type character with the same kind type parameter as STRING_A.

33 **Result Characteristics.** Default logical.

34 **Result Value.** If the strings are of unequal length, the comparison is made as if the shorter
string were extended on the right with blanks to the length of the longer string. If either string
contains a character not in the ASCII character set, the result is processor dependent. The result

1 is true if the strings are equal or if `STRING_A` follows `STRING_B` in the ASCII collating sequence;
2 otherwise, the result is false.

NOTE 13.13

The result is true if both `STRING_A` and `STRING_B` are of zero length.

3 **Example.** `LGE ('ONE', 'TWO')` has the value false.

13.7.102 LGT (STRING_A, STRING_B)

5 **Description.** Test whether a string is lexically greater than another string, based on the ASCII
6 collating sequence.

7 **Class.** Elemental function.

Arguments.

9 `STRING_A` shall be of type default character of type ASCII character.

10 `STRING_B` shall be of type character with the same kind type parameter as `STRING_A`.

11 **Result Characteristics.** Default logical.

12 **Result Value.** If the strings are of unequal length, the comparison is made as if the shorter
13 string were extended on the right with blanks to the length of the longer string. If either string
14 contains a character not in the ASCII character set, the result is processor dependent. The result
15 is true if `STRING_A` follows `STRING_B` in the ASCII collating sequence; otherwise, the result is
16 false.

NOTE 13.14

The result is false if both `STRING_A` and `STRING_B` are of zero length.

17 **Example.** `LGT ('ONE', 'TWO')` has the value false.

13.7.103 LLE (STRING_A, STRING_B)

19 **Description.** Test whether a string is lexically less than or equal to another string, based on the
20 ASCII collating sequence.

21 **Class.** Elemental function.

Arguments.

23 `STRING_A` shall be of type default character of type ASCII character.

24 `STRING_B` shall be of type character with the same kind type parameter as `STRING_A`.

25 **Result Characteristics.** Default logical.

26 **Result Value.** If the strings are of unequal length, the comparison is made as if the shorter string
27 were extended on the right with blanks to the length of the longer string. If either string contains
28 a character not in the ASCII character set, the result is processor dependent. The result is true
29 if the strings are equal or if `STRING_A` precedes `STRING_B` in the ASCII collating sequence;
30 otherwise, the result is false.

NOTE 13.15

The result is true if both <code>STRING_A</code> and <code>STRING_B</code> are of zero length.
--

1 **Example.** `LLE ('ONE', 'TWO')` has the value true.

2 **13.7.104 LLT (STRING_A, STRING_B)**

3 **Description.** Test whether a string is lexically less than another string, based on the .

4 **Class.** Elemental function.

5 **Arguments.**

6 `STRING_A` shall be of type default character of type ASCII character.

7 `STRING_B` shall be of type character with the same kind type parameter as `STRING_A`.

8 **Result Characteristics.** Default logical.

9 **Result Value.** If the strings are of unequal length, the comparison is made as if the shorter
10 string were extended on the right with blanks to the length of the longer string. If either string
11 contains a character not in the ASCII character set, the result is processor dependent. The result
12 is true if `STRING_A` precedes `STRING_B` in the ASCII collating sequence; otherwise, the result
13 is false.

NOTE 13.16

The result is false if both <code>STRING_A</code> and <code>STRING_B</code> are of zero length.

14 **Example.** `LLT ('ONE', 'TWO')` has the value true.

15 **13.7.105 LOG (X)**

16 **Description.** Natural logarithm.

17 **Class.** Elemental function.

18 **Argument.** `X` shall be of type real or complex. If `X` is real, its value shall be greater than zero.
19 If `X` is complex, its value shall not be zero.

20 **Result Characteristics.** Same as `X`.

21 **Result Value.** The result has a value equal to a processor-dependent approximation to $\log_e X$.
22 A result of type complex is the principal value with imaginary part ω in the range $-\pi \leq \omega \leq \pi$.
23 If the real part of `X` is less than zero and the imaginary part of `X` is zero, then the imaginary part
24 of the result is approximately π if the imaginary part of `X` is positive real zero or the processor
25 cannot distinguish between positive and negative real zero, and approximately $-\pi$ if the imaginary
26 part of `X` is negative real zero.

27 **Example.** `LOG (10.0)` has the value 2.3025851 (approximately).

28 **13.7.106 LOG_GAMMA (X)**

29 **Description.** Logarithm of the absolute value of the gamma function.

30 **Class.** Elemental function.

Argument. `X` shall be of type real. Its value shall not be a negative integer or zero.

1 **Result Characteristics.** Same as X.

2 **Result Value.** The result has a value equal to a processor-dependent approximation to the
3 natural logarithm of the absolute value of the gamma function of X.

4 **Example.** LOG_GAMMA (3.0) has the value 0.693 (approximately).

5 **13.7.107 LOG10 (X)**

6 **Description.** Common logarithm.

7 **Class.** Elemental function.

8 **Argument.** X shall be of type real. The value of X shall be greater than zero.

9 **Result Characteristics.** Same as X.

10 **Result Value.** The result has a value equal to a processor-dependent approximation to $\log_{10}X$.

11 **Example.** LOG10 (10.0) has the value 1.0 (approximately).

12 **13.7.108 LOGICAL (L [, KIND])**

13 **Description.** Converts between kinds of logical or from bits to logical.

14 **Class.** Elemental function.

15 **Arguments.**

16 L shall be of type logical or bits.

17 KIND (optional) shall be a scalar integer initialization expression.

18 **Result Characteristics.** Logical. If KIND is present, the kind type parameter is that specified
19 by the value of KIND; otherwise, the kind type parameter is that of default logical.

20 **Result Value.**

21 *Case (i):* If L is of type logical, the value is that of L.

22 *Case (ii):* If L is of type bits and KIND (L) is greater than or equal to BITS_KIND (*result*),
23 the physical representation of the result is the same as that of the rightmost bits
24 of L.

25 *Case (iii):* If L is of type bits and KIND (L) is less than BITS_KIND (*result*), the rightmost
26 KIND(L) bits of the physical representation of the result are the same as those of
27 L, and the remaining bits of the physical representation of the result are zero.

NOTE 13.17

The result of a bits to logical conversion may be used in a context requiring a logical value only if the physical representation of the result is valid as a logical value.

28 **Examples.** LOGICAL (L .OR. .NOT. L) has the value true and is of type default logical,
29 regardless of the kind type parameter of the logical variable L. LOGICAL (BITS (.TRUE.)) has
30 the value true.

31 **13.7.109 MASKL (I, [, KIND])**

32 **Description.** Generate a left justified mask.

1 **Class.** Elemental function.

2 **Arguments.**

3 I shall be of type integer. It shall be nonnegative and less than or equal to the
4 kind type parameter of the result.

5 KIND (optional) shall be a scalar integer initialization expression.

6 **Result Characteristics.** Bits. If KIND is present, the kind type parameter is that specified by
7 the value of KIND; otherwise, the kind type parameter is that of default bits type.

8 **Result Value.** The result value has its leftmost I bits set to 1 and the remaining bits set to 0.

9 **Example.** MASKL (12) has the value Z'FFF00000' if the default bits kind type parameter value
10 is 32.

10 13.7.110 MASKR (I, [, KIND])

11 **Description.** Generate a right justified mask.

12 **Class.** Elemental function.

13 **Arguments.**

14 I shall be of type integer. It shall be nonnegative and less than or equal to the
15 kind type parameter of the result.

16 KIND (optional) shall be a scalar integer initialization expression.

17 **Result Characteristics.** Bits. If KIND is present, the kind type parameter is that specified by
18 the value of KIND; otherwise, the kind type parameter is that of default bits type.

19 **Result Value.** The result value has its rightmost I bits set to 1 and the remaining bits set to 0.

20 **Example.** MASKR (12) has the value Z'00000FFF' if the default bits kind type parameter value
21 is 32.

21 13.7.111 MATMUL (MATRIX_A, MATRIX_B)

22 **Description.** Performs matrix multiplication of numeric or logical matrices.

23 **Class.** Transformational function.

24 **Arguments.**

25 MATRIX_A shall be of numeric type (integer, real, or complex) or of logical type. It shall
26 be an array of rank one or two.

27 MATRIX_B shall be of numeric type if MATRIX_A is of numeric type and of logical type
28 if MATRIX_A is of logical type. It shall be an array of rank one or two. If
29 MATRIX_A has rank one, MATRIX_B shall have rank two. If MATRIX_B
30 has rank one, MATRIX_A shall have rank two. The size of the first (or only)
31 dimension of MATRIX_B shall equal the size of the last (or only) dimension of
32 MATRIX_A.

33 **Result Characteristics.** If the arguments are of numeric type, the type and kind type parameter
34 of the result are determined by the types of the arguments as specified in 7.1.4.2 for the * operator.
35 If the arguments are of type logical, the result is of type logical with the kind type parameter of

1 the arguments as specified in 7.1.4.2 for the .AND. operator. The shape of the result depends on
2 the shapes of the arguments as follows:

3 *Case (i):* If MATRIX_A has shape (n, m) and MATRIX_B has shape (m, k) , the result has
4 shape (n, k) .

5 *Case (ii):* If MATRIX_A has shape (m) and MATRIX_B has shape (m, k) , the result has
6 shape (k) .

7 *Case (iii):* If MATRIX_A has shape (n, m) and MATRIX_B has shape (m) , the result has
8 shape (n) .

9 **Result Value.**

10 *Case (i):* Element (i, j) of the result has the value SUM (MATRIX_A $(i, :)$ * MATRIX_B $(:, j)$) if the arguments are of numeric type and has the value ANY (MATRIX_A $(i, :)$.AND. MATRIX_B $(:, j)$) if the arguments are of logical type.

13 *Case (ii):* Element (j) of the result has the value SUM (MATRIX_A $(:, j)$ * MATRIX_B $(:, j)$) if the arguments are of numeric type and has the value ANY (MATRIX_A $(:, j)$.AND. MATRIX_B $(:, j)$) if the arguments are of logical type.

16 *Case (iii):* Element (i) of the result has the value SUM (MATRIX_A $(i, :)$ * MATRIX_B $(:, :)$) if the arguments are of numeric type and has the value ANY (MATRIX_A $(i, :)$.AND. MATRIX_B $(:, :)$) if the arguments are of logical type.

19 **Examples.** Let A and B be the matrices $\begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \end{bmatrix}$ and $\begin{bmatrix} 1 & 2 \\ 2 & 3 \\ 3 & 4 \end{bmatrix}$; let X and Y be the vectors

20 $[1, 2]$ and $[1, 2, 3]$.

21 *Case (i):* The result of MATMUL (A, B) is the matrix-matrix product AB with the value
22 $\begin{bmatrix} 14 & 20 \\ 20 & 29 \end{bmatrix}$.

23 *Case (ii):* The result of MATMUL (X, A) is the vector-matrix product XA with the value
24 $[5, 8, 11]$.

25 *Case (iii):* The result of MATMUL (A, Y) is the matrix-vector product AY with the value
26 $[14, 20]$.

27 **13.7.112 MAX (A1, A2 [, A3, ...])**

28 **Description.** Maximum value.

29 **Class.** Elemental function.

30 **Arguments.** The arguments shall all have the same type which shall be integer, real, bits, or
31 character and they shall all have the same kind type parameter.

32 **Result Characteristics.** The type and kind type parameter of the result are the same as those
33 of the arguments. For arguments of character type, the length of the result is the length of the
34 longest argument.

35 **Result Value.** The value of the result is that of the largest argument. For arguments of character
36 type, the result is the value that would be selected by application of intrinsic relational operators;
37 that is, the collating sequence for characters with the kind type parameter of the arguments is
38 applied. If the selected argument is shorter than the longest argument, the result is extended with
39 blanks on the right to the length of the longest argument.

40 **Examples.** MAX (-9.0, 7.0, 2.0) has the value 7.0, MAX ('Z', 'BB') has the value 'Z ', and
41 MAX (('A', 'Z'), ('BB', 'Y ')) has the value ('BB', 'Z '). MAX (B'10000', B'01111') has the

1 value B'10000'.

2 **13.7.113 MAXEXPONENT (X)**

3 **Description.** Returns the maximum exponent of the model representing numbers of the same
4 type and kind type parameter as the argument.

5 **Class.** Inquiry function.

6 **Argument.** X shall be of type real. It may be a scalar or an array.

7 **Result Characteristics.** Default integer scalar.

8 **Result Value.** The result has the value e_{\max} , as defined in 13.4 for the model representing
9 numbers of the same type and kind type parameter as X.

10 **Example.** MAXEXPONENT (X) has the value 127 for real X whose model is as in Note 13.4.

11 **13.7.114 MAXLOC (ARRAY, DIM [, MASK, KIND, BACK]) or 12 MAXLOC (ARRAY [, MASK, KIND, BACK])**

13 **Description.** Determine the location of the first element of ARRAY along dimension DIM having
14 the maximum value of the elements identified by MASK.

15 **Class.** Transformational function.

16 **Arguments.**

17 ARRAY shall be of type integer, real, bits, or character. It shall be an array.

18 DIM shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
19 where n is the rank of ARRAY. The corresponding actual argument shall not
20 be an optional dummy argument.

21 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

22 KIND (optional) shall be a scalar integer initialization expression.

23 BACK (optional) shall be scalar and of type logical.

24 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
25 by the value of KIND; otherwise the kind type parameter is that of default integer type. If DIM
26 is absent, the result is an array of rank one and of size equal to the rank of ARRAY; otherwise,
27 the result is of rank $n - 1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$, where $(d_1, d_2, \dots,$
28 $d_n)$ is the shape of ARRAY.

29 **Result Value.**

30 *Case (i):* The result of MAXLOC (ARRAY) is a rank-one array whose element values are the
31 values of the subscripts of an element of ARRAY whose value equals the maximum
32 value of all of the elements of ARRAY. The i th subscript returned lies in the range
33 1 to e_i , where e_i is the extent of the i th dimension of ARRAY. If ARRAY has size
34 zero, all elements of the result are zero.

35 *Case (ii):* The result of MAXLOC (ARRAY, MASK = MASK) is a rank-one array whose
36 element values are the values of the subscripts of an element of ARRAY, corre-
sponding to a true element of MASK, whose value equals the maximum value of
all such elements of ARRAY. The i th subscript returned lies in the range 1 to e_i ,
where e_i is the extent of the i th dimension of ARRAY. If ARRAY has size zero or

1 every element of MASK has the value false, all elements of the result are zero.
 2 *Case (iii):* If ARRAY has rank one, MAXLOC (ARRAY, DIM = DIM [, MASK = MASK]) is
 3 a scalar whose value is equal to that of the first element of MAXLOC (ARRAY [,
 4 MASK = MASK]). Otherwise, the value of element $(s_1, s_2, \dots, s_{DIM-1}, s_{DIM+1},$
 5 $\dots, s_n)$ of the result is equal to
 6
$$\text{MAXLOC (ARRAY (} s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n), \text{DIM}=1$$

 7
$$[, \text{MASK = MASK (} s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n)])$$
.

8 If only one element has the maximum value, that element's subscripts are returned. Otherwise,
 9 if more than one element has the maximum value and BACK is absent or present with the value
 10 false, the element whose subscripts are returned is the first such element, taken in array element
 11 order. If BACK is present with the value true, the element whose subscripts are returned is the
 12 last such element, taken in array element order.

13 If ARRAY has type character, the result is the value that would be selected by application of
 14 intrinsic relational operators; that is, the collating sequence for characters with the kind type
 15 parameter of the arguments is applied.

16 **Examples.**

17 *Case (i):* The value of MAXLOC ((/ 2, 6, 4, 6 /)) is [2] and the value of MAXLOC ([2, 6, 4, 6],
 18 BACK=.TRUE.) is [4].

19 *Case (ii):* If A has the value $\begin{bmatrix} 0 & -5 & 8 & -3 \\ 3 & 4 & -1 & 2 \\ 1 & 5 & 6 & -4 \end{bmatrix}$, MAXLOC (A, MASK = A < 6) has the
 20 value [3, 2]. Note that this is independent of the declared lower bounds for A.

21 *Case (iii):* The value of MAXLOC ((/ 5, -9, 3 /), DIM = 1) is 1. If B has the value
 22 $\begin{bmatrix} 1 & 3 & -9 \\ 2 & 2 & 6 \end{bmatrix}$, MAXLOC (B, DIM = 1) is [2, 1, 2] and MAXLOC (B, DIM = 2)
 23 is [2, 3]. Note that this is independent of the declared lower bounds for B.

24 **13.7.115 MAXVAL (ARRAY, DIM [, MASK]) or MAXVAL (ARRAY [, MASK])**

25 **Description.** Maximum value of the elements of ARRAY along dimension DIM corresponding
 26 to the true elements of MASK.

27 **Class.** Transformational function.

28 **Arguments.**

29 ARRAY shall be of type integer, real, bits, or character. It shall be an array.

30 DIM shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
 where n is the rank of ARRAY. The corresponding actual argument shall not
 be an optional dummy argument.

31 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

32 **Result Characteristics.** The result is of the same type and type parameters as ARRAY. It is
 33 scalar if DIM is absent; otherwise, the result has rank $n-1$ and shape $(d_1, d_2, \dots, d_{DIM-1}, d_{DIM+1},$
 34 $\dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

35 **Result Value.**

36 *Case (i):* The result of MAXVAL (ARRAY) has a value equal to the maximum value of all
 37 the elements of ARRAY if the size of ARRAY is not zero. If ARRAY has size zero

and type integer or real, the result has the value of the negative number of the largest magnitude supported by the processor for numbers of the type and kind type parameter of ARRAY. If ARRAY has size zero and type character, the result has the value of a string of characters of length LEN (ARRAY), with each character equal to CHAR (0, KIND = KIND (ARRAY)).

Case (ii): The result of MAXVAL (ARRAY, MASK = MASK) has a value equal to that of MAXVAL (PACK (ARRAY, MASK)).

Case (iii): The result of MAXVAL (ARRAY, DIM = DIM [,MASK = MASK]) has a value equal to that of MAXVAL (ARRAY [,MASK = MASK]) if ARRAY has rank one. Otherwise, the value of element $(s_1, s_2, \dots, s_{DIM-1}, s_{DIM+1}, \dots, s_n)$ of the result is equal to

$$\text{MAXVAL} (\text{ARRAY} (s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n) \\ [, \text{MASK} = \text{MASK} (s_1, s_2, \dots, s_{DIM-1}, :, s_{DIM+1}, \dots, s_n)]).$$

If ARRAY has type character, the result is the value that would be selected by application of intrinsic relational operators; that is, the collating sequence for characters with the kind type parameter of the arguments is applied.

Examples.

Case (i): The value of MAXVAL ((/ 1, 2, 3 /)) is 3.

Case (ii): MAXVAL (C, MASK = C < 0.0) finds the maximum of the negative elements of C.

Case (iii): If B is the array $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$, MAXVAL (B, DIM = 1) is [2, 4, 6] and MAXVAL (B, DIM = 2) is [5, 6].

13.7.116 MERGE (TSOURCE, FSOURCE, MASK)

Description. Choose alternative value according to the value of a mask.

Class. Elemental function.

Arguments.

TSOURCE may be of any type.

FSOURCE shall be of the same type and type parameters as TSOURCE.

MASK shall be of type logical.

Result Characteristics. Same as TSOURCE.

Result Value. The result is TSOURCE if MASK is true and FSOURCE otherwise.

Examples. If TSOURCE is the array $\begin{bmatrix} 1 & 6 & 5 \\ 2 & 4 & 6 \end{bmatrix}$, FSOURCE is the array $\begin{bmatrix} 0 & 3 & 2 \\ 7 & 4 & 8 \end{bmatrix}$ and MASK is the array $\begin{bmatrix} T & . & T \\ . & . & T \end{bmatrix}$, where “T” represents true and “.” represents false, then MERGE (TSOURCE, FSOURCE, MASK) is $\begin{bmatrix} 1 & 3 & 5 \\ 7 & 4 & 6 \end{bmatrix}$. The value of MERGE (1.0, 0.0, K > 0) is 1.0 for K = 5 and 0.0 for K = -2.

13.7.117 MERGE_BITS (I, J, MASK)

Description. Merge bits under mask.

1

2 **Class.** Elemental function.3 **Arguments.**

4 I shall be of type bits or integer.

5 J shall be of type integer or bits.
BITS_KIND (I) shall be equal to BITS_KIND (J).6 MASK shall be of type integer or bits.
BITS_KIND (I) shall be equal to BITS_KIND (MASK).

7 All integer arguments shall have the same kind type parameter.

8 **Result Characteristics.** If I, J and MASK all have the same type, the result characteristics are
9 the same as I. Otherwise, the result characteristics are the same as those of the argument(s) with
10 integer type.11 **Result Value.** The result has the value of IOR (IAND (I, MASK), IAND (J, NOT (MASK))).12 **Example.** MERGE_BITS(Z'0123', Z'89AB', Z'FF00') has the value Z'01AB'.13 **13.7.118 MIN (A1, A2 [, A3, ...])**14 **Description.** Minimum value.15 **Class.** Elemental function.16 **Arguments.** The arguments shall all be of the same type which shall be integer, real, bits, or
17 character and they shall all have the same kind type parameter.18 **Result Characteristics.** The type and kind type parameter of the result are the same as those
19 of the arguments. For arguments of character type, the length of the result is the length of the
20 longest argument.21 **Result Value.** The value of the result is that of the smallest argument. For arguments of
22 character type, the result is the value that would be selected by application of intrinsic relational
23 operators; that is, the collating sequence for characters with the kind type parameter of the
24 arguments is applied. If the selected argument is shorter than the longest argument, the result is
25 extended with blanks on the right to the length of the longest argument.26 **Examples.** MIN (-9.0, 7.0, 2.0) has the value -9.0, MIN ('A', 'YY') has the value 'A ', and
27 MIN ((/'Z', 'A'/), (/ 'YY', 'B ' /)) has the value (/ 'YY', 'A ' /). MIN (B'10000', B'01111') has the
28 value B'01111'.29 **13.7.119 MINEXPONENT (X)**30 **Description.** Returns the minimum (most negative) exponent of the model representing numbers
31 of the same type and kind type parameter as the argument.32 **Class.** Inquiry function.33 **Argument.** X shall be of type real. It may be a scalar or an array.34 **Result Characteristics.** Default integer scalar.35 **Result Value.** The result has the value e_{\min} , as defined in 13.4 for the model representing

1 numbers of the same type and kind type parameter as X.

2 **Example.** MINEXPONENT (X) has the value -126 for real X whose model is as in Note 13.4.

3 **13.7.120 MINLOC (ARRAY, DIM [, MASK, KIND, BACK]) or** **4 MINLOC (ARRAY [, MASK, KIND, BACK])**

4 **Description.** Determine the location of the first element of ARRAY along dimension DIM having
 5 the minimum value of the elements identified by MASK.

6 **Class.** Transformational function.

7 **Arguments.**

8 ARRAY shall be of type integer, real, bits, or character. It shall be an array.

9 DIM shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
 where n is the rank of ARRAY. The corresponding actual argument shall not
 be an optional dummy argument.

10 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

11 KIND (optional) shall be a scalar integer initialization expression.

12 BACK (optional) shall be scalar and of type logical.

13 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
 14 by the value of KIND; otherwise the kind type parameter is that of default integer type. If DIM
 15 is absent, the result is an array of rank one and of size equal to the rank of ARRAY; otherwise,
 16 the result is of rank $n - 1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n)
 17 is the shape of ARRAY.

18 **Result Value.**

19 *Case (i):* The result of MINLOC (ARRAY) is a rank-one array whose element values are the
 20 values of the subscripts of an element of ARRAY whose value equals the minimum
 21 value of all the elements of ARRAY. The i th subscript returned lies in the range 1
 22 to e_i , where e_i is the extent of the i th dimension of ARRAY. If ARRAY has size
 23 zero, all elements of the result are zero.

24 *Case (ii):* The result of MINLOC (ARRAY, MASK = MASK) is a rank-one array whose
 25 element values are the values of the subscripts of an element of ARRAY, corre-
 26 sponding to a true element of MASK, whose value equals the minimum value of
 27 all such elements of ARRAY. The i th subscript returned lies in the range 1 to e_i ,
 28 where e_i is the extent of the i th dimension of ARRAY. If ARRAY has size zero or
 29 every element of MASK has the value false, all elements of the result are zero.

30 *Case (iii):* If ARRAY has rank one, MINLOC (ARRAY, DIM = DIM [, MASK = MASK]) is
 31 a scalar whose value is equal to that of the first element of MINLOC (ARRAY [,
 32 MASK = MASK]). Otherwise, the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1},$
 33 $\dots, s_n)$ of the result is equal to

34 $\text{MINLOC (ARRAY (} s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n), \text{DIM}=1$
 35 $[\text{, MASK = MASK (} s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)])$.

36 If only one element has the minimum value, that elements subscripts are returned. Otherwise,
 37 if more than one element has the minimum value and BACK is absent or present with the value
 38 false, the element whose subscripts are returned is the first such element, taken in array element
 39 order. If BACK is present with the value true, the element whose subscripts are returned is the

1 last such element, taken in array element order.

2 If ARRAY has type character, the result is the value that would be selected by application of
3 intrinsic relational operators; that is, the collating sequence for characters with the kind type
4 parameter of the arguments is applied.

5 **Examples.**

6 *Case (i):* The value of MINLOC ((/ 4, 3, 6, 3 /)) is [2] and the value of MINLOC ([4, 3, 6, 3],
7 BACK = .TRUE.) is [4].

8 *Case (ii):* If A has the value $\begin{bmatrix} 0 & -5 & 8 & -3 \\ 3 & 4 & -1 & 2 \\ 1 & 5 & 6 & -4 \end{bmatrix}$, MINLOC (A, MASK = A > -4) has the
9 value [1, 4]. Note that this is independent of the declared lower bounds for A.

10 *Case (iii):* The value of MINLOC ((/ 5, -9, 3 /), DIM = 1) is 2. If B has the value $\begin{bmatrix} 1 & 3 & -9 \\ 2 & 2 & 6 \end{bmatrix}$,
11 MINLOC (B, DIM = 1) is [1, 2, 1] and MINLOC (B, DIM = 2) is [3, 1]. Note that
12 this is independent of the declared lower bounds for B.

13 **13.7.121 MINVAL (ARRAY, DIM [, MASK]) or MINVAL (ARRAY [, MASK])**

14 **Description.** Minimum value of all the elements of ARRAY along dimension DIM corresponding
15 to true elements of MASK.

16 **Class.** Transformational function.

17 **Arguments.**

18 ARRAY shall be of type integer, real, bits, or character. It shall be an array.

19 DIM shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
20 where n is the rank of ARRAY. The corresponding actual argument shall not
21 be an optional dummy argument.

22 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

23 **Result Characteristics.** The result is of the same type and type parameters as ARRAY. It is
24 scalar if DIM is absent; otherwise, the result has rank $n-1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1},$
25 $\dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

26 **Result Value.**

27 *Case (i):* The result of MINVAL (ARRAY) has a value equal to the minimum value of all
28 the elements of ARRAY if the size of ARRAY is not zero. If ARRAY has size
29 zero and type integer or real, the result has the value of the positive number of
30 the largest magnitude supported by the processor for numbers of the type and
31 kind type parameter of ARRAY. If ARRAY has size zero and type character, the
32 result has the value of a string of characters of length LEN (ARRAY), with each
33 character equal to CHAR ($n-1$, KIND = KIND (ARRAY)), where n is the number
34 of characters in the collating sequence for characters with the kind type parameter
35 of ARRAY.

36 *Case (ii):* The result of MINVAL (ARRAY, MASK = MASK) has a value equal to that of
37 MINVAL (PACK (ARRAY, MASK)).

38 *Case (iii):* The result of MINVAL (ARRAY, DIM = DIM [, MASK = MASK]) has a value
equal to that of MINVAL (ARRAY [, MASK = MASK]) if ARRAY has rank one.
Otherwise, the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$ of the result
is equal to

1

2

$$\text{MINVAL} (\text{ARRAY} (s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n) \\ [, \text{MASK} = \text{MASK} (s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)]).$$

3

4

If ARRAY has type character, the result is the value that would be selected by application of intrinsic relational operators; that is, the collating sequence for characters with the kind type parameter of the arguments is applied.

5

6

Examples.

7

Case (i): The value of MINVAL ((/ 1, 2, 3 /)) is 1.

8

Case (ii): MINVAL (C, MASK = C > 0.0) forms the minimum of the positive elements of C.

9

Case (iii): If B is the array $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$, MINVAL (B, DIM = 1) is [1, 3, 5] and MINVAL (B, DIM = 2) is [1, 2].

10

11

13.7.122 MOD (A, P)

12

Description. Remainder function.

13

Class. Elemental function.

14

Arguments.

15

A shall be of type integer or real.

16

P shall be of the same type and kind type parameter as A. P shall not be zero.

17

Result Characteristics. Same as A.

18

Result Value. The value of the result is A-INT (A/P) * P.

19

Examples. MOD (3.0, 2.0) has the value 1.0 (approximately). MOD (8, 5) has the value 3. MOD (-8, 5) has the value -3. MOD (8, -5) has the value 3. MOD (-8, -5) has the value -3.

20

21

13.7.123 MODULO (A, P)

22

Description. Modulo function.

23

Class. Elemental function.

24

Arguments.

25

A shall be of type integer or real.

26

P shall be of the same type and kind type parameter as A. P shall not be zero.

27

Result Characteristics. Same as A.

28

Result Value.

29

Case (i): A is of type integer. MODULO (A, P) has the value R such that $A = Q \times P + R$, where Q is an integer, the inequalities $0 \leq R < P$ hold if $P > 0$, and $P < R \leq 0$ hold if $P < 0$.

30

31

Case (ii): A is of type real. The value of the result is $A - \text{FLOOR} (A / P) * P$.

32

Examples. MODULO (8, 5) has the value 3. MODULO (-8, 5) has the value 2. MODULO (8, -5) has the value -2. MODULO (-8, -5) has the value -3.

33

1 13.7.124 MOVE_ALLOC (FROM, TO)

2 **Description.** Moves an allocation from one allocatable object to another.

3 **Class.** Pure subroutine.

4 **Arguments.**

5 FROM may be of any type and rank. It shall be allocatable. It is an INTENT(INOUT) argument.

6 TO shall be type compatible (4.3.1.3) with FROM and have the same rank. It shall be allocatable. It shall be polymorphic if FROM is polymorphic. It is an INTENT(OUT) argument. Each nondeferred parameter of the declared type of TO shall have the same value as the corresponding parameter of the declared type of FROM.

7 The allocation status of TO becomes unallocated if FROM is unallocated on entry to MOVE_-
8 ALLOC. Otherwise, TO becomes allocated with dynamic type, type parameters, array bounds,
9 and value identical to those that FROM had on entry to MOVE_ALLOC.

10 If TO has the TARGET attribute, any pointer associated with FROM on entry to MOVE_ALLOC
11 becomes correspondingly associated with TO. If TO does not have the TARGET attribute, the
12 pointer association status of any pointer associated with FROM on entry becomes undefined.

13 The allocation status of FROM becomes unallocated.

14 **Example.**

```
15 REAL,ALLOCATABLE :: GRID(:),TEMPGRID(:)
16 ...
17 ALLOCATE(GRID(-N:N))          ! initial allocation of GRID
18 ...
19 ! "reallocation" of GRID to allow intermediate points
20 ALLOCATE(TEMPGRID(-2*N:2*N)) ! allocate bigger grid
21 TEMPGRID(:,2)=GRID ! distribute values to new locations
22 CALL MOVE_ALLOC(TO=GRID,FROM=TEMPGRID)
23             ! old grid is deallocated because TO is
24             ! INTENT(OUT), and GRID then "takes over"
25             ! new grid allocation
```

NOTE 13.18

It is expected that the implementation of allocatable objects will typically involve descriptors to locate the allocated storage; MOVE_ALLOC could then be implemented by transferring the contents of the descriptor for FROM to the descriptor for TO and clearing the descriptor for FROM.

26 13.7.125 MVBITS (FROM, FROMPOS, LEN, TO, TOPOS)

27 **Description.** Copies a sequence of bits from one data object to another.

28 **Class.** Elemental subroutine.

Arguments.

1		
2	FROM	shall be of type integer or bits. It is an INTENT (IN) argument.
	FROMPOS	shall be of type integer and nonnegative. It is an INTENT (IN) argument. FROMPOS + LEN shall be less than or equal to BIT_SIZE (FROM). The model for the interpretation of an integer value as a sequence of bits is in 13.3.
3		
4	LEN	shall be of type integer and nonnegative. It is an INTENT (IN) argument.
	TO	shall be a variable of the same type and kind type parameter value as FROM and may be associated with FROM (12.8.3). It is an INTENT (INOUT) argument. TO is defined by copying the sequence of bits of length LEN, starting at position FROMPOS of FROM to position TOPOS of TO. No other bits of TO are altered. On return, the LEN bits of TO starting at TOPOS are equal to the value that the LEN bits of FROM starting at FROMPOS had on entry. The model for the interpretation of an integer value as a sequence of bits is in 13.3.
5		
	TOPOS	shall be of type integer and nonnegative. It is an INTENT (IN) argument. TOPOS + LEN shall be less than or equal to BIT_SIZE (TO).
6		
7	Examples.	If TO has the initial value 6, the value of TO after the statement
8		CALL MVBITS (7, 2, 2, TO, 0) is 5. If TO has the initial value B'000000111111', the value of
9		TO after the statement CALL MVBITS (B'000000000011', 0, 2, TO, 8) is B'001100111111'.

10 13.7.126 NEAREST (X, S)

11 **Description.** Returns the nearest different machine-representable number in a given direction.

12 **Class.** Elemental function.

13 **Arguments.**

14 X shall be of type real.

15 S shall be of type real and not equal to zero.

16 **Result Characteristics.** Same as X.

17 **Result Value.** The result has a value equal to the machine-representable number distinct from
18 X and nearest to it in the direction of the infinity with the same sign as S.

19 **Example.** NEAREST (3.0, 2.0) has the value $3 + 2^{-22}$ on a machine whose representation is that
20 of the model in Note 13.4.

NOTE 13.19

Unlike other floating-point manipulation functions, NEAREST operates on machine-representable numbers rather than model numbers. On many systems there are machine-representable numbers that lie between adjacent model numbers.

21 13.7.127 NEW_LINE (A)

22 **Description.** Returns a newline character.

23 **Class.** Inquiry function.

24 **Argument.** A shall be of type character. It may be a scalar or an array.

1 **Result Characteristics.** Character scalar of length one with the same kind type parameter as
2 A.

3 **Result Value.**

4 *Case (i):* If A is of the default character type and the character in position 10 of the ASCII
5 collating sequence is representable in the default character set, then the result is
6 ACHAR(10).

7 *Case (ii):* If A is of the ASCII character type or the ISO 10646 character type, then the result
8 is CHAR(10,KIND(A)).

9 *Case (iii):* Otherwise, the result is a processor-dependent character that represents a newline
10 in output to files connected for formatted stream output if there is such a character.

11 *Case (iv):* Otherwise, the result is the blank character.

12 **13.7.128 NINT (A [, KIND])**

13 **Description.** Nearest integer.

14 **Class.** Elemental function.

15 **Arguments.**

16 A shall be of type real.

17 KIND (optional) shall be a scalar integer initialization expression.

18 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
19 by the value of KIND; otherwise, the kind type parameter is that of default integer type.

20 **Result Value.** The result is the integer nearest A, or if there are two integers equally near A,
21 the result is whichever such integer has the greater magnitude.

22 **Example.** NINT (2.783) has the value 3.

23 **13.7.129 NOT (I)**

24 **Description.** Performs a bitwise complement.

25 **Class.** Elemental function.

26 **Argument.** I shall be of type integer or bits.

27 **Result Characteristics.** Same as I.

28 **Result Value.** The result has the value obtained by complementing I bit-by-bit according to the
29 following truth table:

I	NOT (I)
1	0
0	1

30 The model for the interpretation of an integer value as a sequence of bits is in 13.3.

31 **Examples.** If I is represented by the string of bits 01010101, NOT (I) has the binary value
32 10101010. NOT (Z'FFFF0000') has the value Z'0000FFFF'.

1 13.7.130 NORM2 (X)

2 **Description.** L_2 norm of an array.

3 **Class.** Transformational function.

4 **Argument.** X shall be of type real. It shall be an array.

5 **Result Characteristics.** Scalar of the same type and kind type parameter value as X.

6 **Result Value.** The result has a value equal to a processor-dependent approximation to the
7 generalized L_2 norm of X, which is the square root of the sum of the squares of the elements of X.
8 It is recommended that the processor compute the result without undue overflow or underflow.

9 **Example.** The value of NORM2 ((/ 3.0, 4.0 /)) is 5.0 (approximately).

10 13.7.131 NULL ([MOLD])

11 **Description.** Returns a disassociated pointer or designates an unallocated allocatable component
12 of a structure constructor.

13 **Class.** Transformational function.

14 **Argument.** MOLD shall be a pointer or allocatable. It may be of any type or may be a procedure
15 pointer. If MOLD is a pointer its pointer association status may be undefined, disassociated, or
16 associated. If MOLD is allocatable its allocation status may be allocated or unallocated. It need
17 not be defined with a value.

18 **Result Characteristics.** If MOLD is present, the characteristics are the same as MOLD. If
19 MOLD has deferred type parameters, those type parameters of the result are deferred.

20 If MOLD is absent, the characteristics of the result are determined by the entity with which the
21 reference is associated. See Table 13.1. MOLD shall not be absent in any other context. If any
22 type parameters of the contextual entity are deferred, those type parameters of the result are
23 deferred. If any type parameters of the contextual entity are assumed, MOLD shall be present.

24 If the context of the reference to NULL is an actual argument to a generic procedure, MOLD
25 shall be present if the type, type parameters, or rank are required to resolve the generic reference.
26 MOLD shall also be present if the reference appears as an actual argument corresponding to a
27 dummy argument with assumed character length.

Table 13.1: Characteristics of the result of NULL ()

Appearance of NULL ()	Type, type parameters, and rank of result:
right side of a pointer assignment	pointer on the left side
initialization for an object in a declaration	the object
default initialization for a component	the component
in a structure constructor	the corresponding component
as an actual argument	the corresponding dummy argument
in a DATA statement	the corresponding pointer object

28 **Result.** The result is a disassociated pointer or an unallocated allocatable entity.

29 **Examples.**

30 *Case (i):* REAL, POINTER, DIMENSION(:) :: VEC => NULL () defines the initial asso-
31 ciation status of VEC to be disassociated.

```

1      Case (ii):   The MOLD argument is required in the following:
2
3      INTERFACE GEN
4          SUBROUTINE S1 (J, PI)
5              INTEGER J
6              INTEGER, POINTER :: PI
7          END SUBROUTINE S1
8          SUBROUTINE S2 (K, PR)
9              INTEGER K
10             REAL, POINTER :: PR
11         END SUBROUTINE S2
12     END INTERFACE
13     REAL, POINTER :: REAL_PTR
14     CALL GEN (7, NULL (REAL_PTR) )      ! Invokes S2

```

14 13.7.132 NUM_IMAGES ()

15 **Description.** Returns the number of images.

16 **Class.** Inquiry function.

17 **Argument.** None.

18 **Result Characteristics.** Default integer scalar.

19 **Result Value.** The number of images.

20 **Example.** The following code uses image 1 to read data and broadcast it to other images.

```

21 REAL :: P[*]
22 ...
23 IF (THIS_IMAGE()==1) THEN
24     READ (6,*) P
25     DO I = 2, NUM_IMAGES()
26         P[I] = P
27     END DO
28 END IF
29 SYNC ALL

```

30 13.7.133 PACK (ARRAY, MASK [, VECTOR])

31 **Description.** Pack an array into an array of rank one under the control of a mask.

32 **Class.** Transformational function.

33 **Arguments.**

34 **ARRAY** may be of any type. It shall be an array.

35 **MASK** shall be of type logical and shall be conformable with ARRAY.

VECTOR shall be of the same type and type parameters as ARRAY and shall have rank one. VECTOR shall have at least as many elements as there are true elements in MASK. If MASK is scalar with the value true, VECTOR shall have at least as many elements as there are in ARRAY.

Result Characteristics. The result is an array of rank one with the same type and type parameters as ARRAY. If VECTOR is present, the result size is that of VECTOR; otherwise, the result size is the number t of true elements in MASK unless MASK is scalar with the value true, in which case the result size is the size of ARRAY.

Result Value. Element i of the result is the element of ARRAY that corresponds to the i th true element of MASK, taking elements in array element order, for $i = 1, 2, \dots, t$. If VECTOR is present and has size $n > t$, element i of the result has the value VECTOR (i), for $i = t + 1, \dots, n$.

Examples. The nonzero elements of an array M with the value $\begin{bmatrix} 0 & 0 & 0 \\ 9 & 0 & 0 \\ 0 & 0 & 7 \end{bmatrix}$ may be “gathered” by the function PACK. The result of PACK (M, MASK = M /= 0) is [9, 7] and the result of PACK (M, M /= 0, VECTOR = (/ 2, 4, 6, 8, 10, 12 /)) is [9, 7, 6, 8, 10, 12].

13.7.134 PARITY (MASK [, DIM])

Description. Determine whether an odd number of values are true in MASK along dimension DIM.

Class. Transformational function.

Arguments.

MASK shall be of type logical. It shall be an array.

DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where n is the rank of MASK. The corresponding actual argument shall not be an optional dummy argument.

Result Characteristics. The result is of type logical with the same kind type parameter as MASK. It is scalar if DIM is absent; otherwise, the result has rank $n - 1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1}, \dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of MASK.

Result Value.

Case (i): The result of PARITY (MASK) has the value true if an odd number of the elements of MASK are true, and false otherwise.

Case (ii): If MASK has rank one, PARITY (MASK, DIM) is equal to PARITY (MASK). Otherwise, the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$ of PARITY (MASK, DIM) is equal to PARITY (MASK ($s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n$)).

Examples.

Case (i): The value of PARITY ([T, T, T, F]) is true if T has the value true and F has the value false.

Case (ii): If B is the array $\begin{bmatrix} T & T & F \\ T & T & T \end{bmatrix}$, where T has the value true and F has the value false, then PARITY (B, DIM=1) has the value [F, F, T] and PARITY (B, DIM=2) has the value [F, T].

1 13.7.135 POPCNT (I)

2 **Description.** Return the number of one bits.

3 **Class.** Elemental function.

4 **Argument.** I shall be of type integer or bits.

5 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
6 by the value of KIND; otherwise, the kind type parameter is that of the default integer type.

7 **Result Value.** If I is of type integer, the result value is equal to the number of one bits in the
8 sequence of bits of I. The model for the interpretation of an integer value as a sequence of bits is
9 in 13.3. If I is of type bits, the result value is the number of one bits in I.

10 **Examples.** POPCNT ([1, 2, 3, 4, 5, 6]) has the value [1, 1, 2, 1, 2, 2]. POPCNT (Z'FFFF0000')
11 has the value 16. If B is of type bits, POPCNT (HUGE (B)) has the same value as KIND (B).

12 13.7.136 POPPAR (I)

13 **Description.** Return the bitwise parity.

14 **Class.** Elemental function.

15 **Argument.** I shall be of type integer or bits.

16 **Result Characteristics.** Default integer.

17 **Result Value.** POPPAR (I) has the value 1 if POPCNT (I) is odd, and 0 if POPCNT (I) is
18 even.

19 **Examples.** POPPAR ([1, 2, 3, 4, 5, 6]) has the value [1, 1, 0, 1, 0, 0]. POPPAR (Z'FFFF0000')
20 has the value 0.

21 13.7.137 PRECISION (X)

22 **Description.** Returns the decimal precision of the model representing real numbers with the
23 same kind type parameter as the argument.

24 **Class.** Inquiry function.

25 **Argument.** X shall be of type real or complex. It may be a scalar or an array.

26 **Result Characteristics.** Default integer scalar.

27 **Result Value.** The result has the value $\text{INT}((p - 1) * \text{LOG}_{10}(b)) + k$, where b and p are as
28 defined in 13.4 for the model representing real numbers with the same value for the kind type
29 parameter as X, and where k is 1 if b is an integral power of 10 and 0 otherwise.

30 **Example.** PRECISION (X) has the value $\text{INT}(23 * \text{LOG}_{10}(2.)) = \text{INT}(6.92\dots) = 6$ for real
31 X whose model is as in Note 13.4.

32 13.7.138 PRESENT (A)

33 **Description.** Determine whether an optional argument is present.

34 **Class.** Inquiry function.

35 **Argument.** A shall be the name of an optional dummy argument that is accessible in the

1 subprogram in which the PRESENT function reference appears. It may be of any type and it
 2 may be a pointer. It may be a scalar or an array. It may be a dummy procedure. The dummy
 3 argument A has no INTENT attribute.

4 **Result Characteristics.** Default logical scalar.

5 **Result Value.** The result has the value true if A is present (12.5.2.13) and otherwise has the
 6 value false.

7 13.7.139 PRODUCT (ARRAY, DIM [, MASK]) or PRODUCT (ARRAY [, MASK])

8 **Description.** Product of all the elements of ARRAY along dimension DIM corresponding to the
 9 true elements of MASK.

10 **Class.** Transformational function.

11 **Arguments.**

12 ARRAY shall be of type integer, real, or complex. It shall be an array.

DIM shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
 where n is the rank of ARRAY. The corresponding actual argument shall not
 13 be an optional dummy argument.

14 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

15 **Result Characteristics.** The result is of the same type and kind type parameter as ARRAY. It is
 16 scalar if DIM is absent; otherwise, the result has rank $n-1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1},$
 17 $\dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

18 **Result Value.**

19 *Case (i):* The result of PRODUCT (ARRAY) has a value equal to a processor-dependent
 20 approximation to the product of all the elements of ARRAY or has the value one
 21 if ARRAY has size zero.

22 *Case (ii):* The result of PRODUCT (ARRAY, MASK = MASK) has a value equal to a
 23 processor-dependent approximation to the product of the elements of ARRAY cor-
 24 responding to the true elements of MASK or has the value one if there are no true
 25 elements.

26 *Case (iii):* If ARRAY has rank one, PRODUCT (ARRAY, DIM = DIM [, MASK = MASK])
 27 has a value equal to that of PRODUCT (ARRAY [, MASK = MASK]). Otherwise,
 28 the value of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$ of PRODUCT (ARRAY,
 29 DIM = DIM [, MASK = MASK]) is equal to

30
$$\text{PRODUCT (ARRAY (} s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n) [, \text{MASK} =$$

 31
$$\text{MASK (} s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)]).$$

32 **Examples.**

33 *Case (i):* The value of PRODUCT ((/ 1, 2, 3 /)) is 6.

34 *Case (ii):* PRODUCT (C, MASK = C > 0.0) forms the product of the positive elements of C.

35 *Case (iii):* If B is the array $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$, PRODUCT (B, DIM = 1) is [2, 12, 30] and PROD-
 36 UCT (B, DIM = 2) is [15, 48].

37 13.7.140 RADIX (X)

1 **Description.** Returns the base of the model representing numbers of the same type and kind
2 type parameter as the argument.

3 **Class.** Inquiry function.

4 **Argument.** X shall be of type integer or real. It may be a scalar or an array.

5 **Result Characteristics.** Default integer scalar.

6 **Result Value.** The result has the value r if X is of type integer and the value b if X is of type
7 real, where r and b are as defined in 13.4 for the model representing numbers of the same type
8 and kind type parameter as X.

9 **Example.** RADIX (X) has the value 2 for real X whose model is as in Note 13.4.

10 13.7.141 RANDOM_NUMBER (HARVEST)

11 **Description.** Returns one pseudorandom number or an array of pseudorandom numbers.

12 **Class.** Subroutine.

13 **Argument.** HARVEST shall be of type real or bits. It is an INTENT (OUT) argument. It may
14 be a scalar or an array variable. If it is real, it is assigned pseudorandom numbers from the uniform
15 distribution in the interval $0 \leq x < 1$. If it is of type bits, it is assigned pseudorandom values
16 with each of the KIND (HARVEST) bits of each value having a probability of approximately 0.5
17 of being 1.

18 **Examples.**

```
19         REAL X, Y (10, 10)
20         BITS B
21         ! Initialize X with a pseudorandom number
22         CALL RANDOM_NUMBER (HARVEST = X)
23         CALL RANDOM_NUMBER (Y)
24         ! X and Y contain uniformly distributed random numbers
25         CALL RANDOM_NUMBER(B)
26         ! B contains a uniformly random collection of 0 and 1 bits.
```

27 13.7.142 RANDOM_SEED ([SIZE, PUT, GET])

28 **Description.** Restarts or queries the pseudorandom number generator used by RANDOM_NUM-
29 BER.

30 **Class.** Subroutine.

31 **Arguments.** There shall either be exactly one or no arguments present.

 SIZE (optional) shall be scalar and of type default integer. It is an INTENT (OUT) argument.
 It is assigned the number N of integers that the processor uses to hold the
32 value of the seed.

 PUT (optional) shall be a default integer array of rank one and size $\geq N$. It is an INTENT (IN)
 argument. It is used in a processor-dependent manner to compute the seed
33 value accessed by the pseudorandom number generator.

1 GET (optional) shall be a default integer array of rank one and size $\geq N$. It is an INTENT (OUT) argument. It is assigned the current value of the seed.

2 If no argument is present, the processor assigns a processor-dependent value to the seed.

3 The pseudorandom number generator used by RANDOM_NUMBER maintains a seed that is
4 updated during the execution of RANDOM_NUMBER and that may be specified or returned by
5 RANDOM_SEED. Computation of the seed from the argument PUT is performed in a processor-
6 dependent manner. The value returned by GET need not be the same as the value specified by
7 PUT in an immediately preceding reference to RANDOM_SEED. For example, following execution
8 of the statements

```
9      CALL RANDOM_SEED (PUT=SEED1)
10     CALL RANDOM_SEED (GET=SEED2)
```

11 SEED2 need not equal SEED1. When the values differ, the use of either value as the PUT argu-
12 ment in a subsequent call to RANDOM_SEED shall result in the same sequence of pseudorandom
13 numbers being generated. For example, after execution of the statements

```
14     CALL RANDOM_SEED (PUT=SEED1)
15     CALL RANDOM_SEED (GET=SEED2)
16     CALL RANDOM_NUMBER (X1)
17     CALL RANDOM_SEED (PUT=SEED2)
18     CALL RANDOM_NUMBER (X2)
```

19 X2 equals X1.

20 **Examples.**

```
21     CALL RANDOM_SEED                ! Processor initialization
22     CALL RANDOM_SEED (SIZE = K)      ! Puts size of seed in K
23     CALL RANDOM_SEED (PUT = SEED (1 : K)) ! Define seed
24     CALL RANDOM_SEED (GET = OLD (1 : K)) ! Read current seed
```

25 13.7.143 RANGE (X)

26 **Description.** Returns the decimal exponent range of the model representing integer or real
27 numbers with the same kind type parameter as the argument.

28 **Class.** Inquiry function.

29 **Argument.** X shall be of type integer, real, or complex. It may be a scalar or an array.

30 **Result Characteristics.** Default integer scalar.

31 **Result Value.**

32 *Case (i):* For an integer argument, the result has the value INT (LOG10 (HUGE(X))).

33 *Case (ii):* For a real argument, the result has the value INT (MIN (LOG10 (HUGE(X)),
34 -LOG10 (TINY(X)))).

35 *Case (iii):* For a complex argument, the result has the value RANGE(REAL(X)).

1 **Examples.** RANGE (X) has the value 38 for real X whose model is as in Note 13.4, because in
2 this case $HUGE(X) = (1 - 2^{-24}) \times 2^{127}$ and $TINY(X) = 2^{-127}$.

3 **13.7.144 REAL (A [, KIND])**

4 **Description.** Convert to real type.

5 **Class.** Elemental function.

6 **Arguments.**

7 A shall be of type integer, real, complex, or bits.

8 KIND (optional) shall be a scalar integer initialization expression.

9 **Result Characteristics.** Real.

10 *Case (i):* If A is of type integer, real, or bits, and KIND is present, the kind type parameter
11 is that specified by the value of KIND. If A is of type integer, real, or bits, and
12 KIND is not present, the kind type parameter is that of default real type.

13 *Case (ii):* If A is of type complex and KIND is present, the kind type parameter is that
14 specified by the value of KIND. If A is of type complex and KIND is not present,
15 the kind type parameter is the kind type parameter of A.

16 **Result Value.**

17 *Case (i):* If A is of type integer or real, the result is equal to a processor-dependent approx-
18 imation to A.

19 *Case (ii):* If A is of type complex, the result is equal to a processor-dependent approximation
20 to the real part of A.

21 *Case (iii):* If A is of type bits and KIND (A) is greater than or equal to BITS_KIND (*result*),
22 the result has the value whose physical representation is the same as the rightmost
23 bits of A.

24 *Case (iv):* If A is of type bits and KIND (A) is less than BITS_KIND (*result*), the rightmost
25 KIND (A) bits of the physical representation of the result value are the same as
26 those of A, and the remaining bits of the physical representation of the result value
27 are zero.

28 **Examples.** REAL (-3) has the value -3.0. REAL (Z) has the same kind type parameter and the
29 same value as the real part of the complex variable Z. REAL (Z'7F800000') has the value positive
30 infinity if the default real kind is IEEE single precision.

31 **13.7.145 REPEAT (STRING, NCOPIES)**

32 **Description.** Concatenate several copies of a string.

33 **Class.** Transformational function.

34 **Arguments.**

35 STRING shall be scalar and of type character.

36 NCOPIES shall be scalar and of type integer. Its value shall not be negative.

37 **Result Characteristics.** Character scalar of length NCOPIES times that of STRING, with the
38 same kind type parameter as STRING.

1 **Result Value.** The value of the result is the concatenation of NCOPIES copies of STRING.

2 **Examples.** REPEAT ('H', 2) has the value HH. REPEAT ('XYZ', 0) has the value of a zero-
3 length string.

4 **13.7.146 RESHAPE (SOURCE, SHAPE [, PAD, ORDER])**

5 **Description.** Constructs an array of a specified shape from the elements of a given array.

6 **Class.** Transformational function.

7 **Arguments.**

SOURCE may be of any type. It shall be an array. If PAD is absent or of size zero, the
8 size of SOURCE shall be greater than or equal to PRODUCT (SHAPE). The
size of the result is the product of the values of the elements of SHAPE.

SHAPE shall be of type integer, rank one, and constant size. Its size shall be positive
9 and less than 8. It shall not have an element whose value is negative.

10 PAD (optional) shall be of the same type and type parameters as SOURCE. It shall be an array.

ORDER shall be of type integer, shall have the same shape as SHAPE, and its value
11 (optional) shall be a permutation of $(1, 2, \dots, n)$, where n is the size of SHAPE. If absent,
it is as if it were present with value $(1, 2, \dots, n)$.

12 **Result Characteristics.** The result is an array of shape SHAPE (that is, SHAPE (RE-
13 SHAPE (SOURCE, SHAPE, PAD, ORDER)) is equal to SHAPE) with the same type and type
14 parameters as SOURCE.

15 **Result Value.** The elements of the result, taken in permuted subscript order ORDER (1),
16 \dots , ORDER (n), are those of SOURCE in normal array element order followed if necessary by
17 those of PAD in array element order, followed if necessary by additional copies of PAD in array
18 element order.

19 **Examples.** RESHAPE ((/ 1, 2, 3, 4, 5, 6 /), (/ 2, 3 /)) has the value $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$.
20 RESHAPE ((/ 1, 2, 3, 4, 5, 6 /), (/ 2, 4 /), (/ 0, 0 /), (/ 2, 1 /)) has the value $\begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 0 & 0 \end{bmatrix}$.

21 **13.7.147 RRSPACING (X)**

22 **Description.** Returns the reciprocal of the relative spacing of model numbers near the argument
23 value.

24 **Class.** Elemental function.

25 **Argument.** X shall be of type real.

26 **Result Characteristics.** Same as X.

27 **Result Value.** The result has the value $|X \times b^{-e}| \times b^p$, where b , e , and p are as defined in 13.4
28 for the value nearest to X in the model for real values whose kind type parameter is that of X; if
29 there are two such values, the value of greater absolute value is taken. If X is an IEEE infinity,
30 the result is zero. If X is an IEEE NaN, the result is that NaN.

J3 internal note

Unresolved Technical Issue 092

This is contradictory for IEEE infinity: the specific sentence contradicts the formula, which would imply that $RRSPACING(\pm\infty)=NaN$.

Moreover, it appears to me that there is no limit as $X \rightarrow \infty$ for $RRSPACING$, and so the value implied by the formula (NaN) is indeed the correct value for $RRSPACING(\pm\infty)$.

1 **Example.** $RRSPACING(-3.0)$ has the value 0.75×2^{24} for reals whose model is as in Note 13.4.

2 **13.7.148 SAME_TYPE_AS (A, B)**

3 **Description.** Inquires whether the dynamic type of A is the same as the dynamic type of B.

4 **Class.** Inquiry function.

5 **Arguments.**

6 A shall be an object of extensible type. If it is a pointer, it shall not have an undefined association status.

7 B shall be an object of extensible type. If it is a pointer, it shall not have an undefined association status.

8 **Result Characteristics.** Default logical scalar.

9 **Result Value.** The result is true if and only if the dynamic type of A is the same as the dynamic type of B.
10

NOTE 13.20

The dynamic type of a disassociated pointer or unallocated allocatable is its declared type.

11 **13.7.149 SCALE (X, I)**

12 **Description.** Returns $X \times b^I$ where b is the base of the model representation of X.

13 **Class.** Elemental function.

14 **Arguments.**

15 X shall be of type real.

16 I shall be of type integer.

17 **Result Characteristics.** Same as X.

18 **Result Value.** The result has the value $X \times b^I$, where b is defined in 13.4 for model numbers representing values of X, provided this result is within range; if not, the result is processor dependent.
19
20

21 **Example.** $SCALE(3.0, 2)$ has the value 12.0 for reals whose model is as in Note 13.4.

22 **13.7.150 SCAN (STRING, SET [, BACK, KIND])**

23 **Description.** Scan a string for any one of the characters in a set of characters.

24 **Class.** Elemental function.

25 **Arguments.**

1 STRING shall be of type character.
 2 SET shall be of type character with the same kind type parameter as STRING.
 3 BACK (optional) shall be of type logical.
 4 KIND (optional) shall be a scalar integer initialization expression.

5 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
 6 by the value of KIND; otherwise the kind type parameter is that of default integer type.

7 **Result Value.**

8 *Case (i):* If BACK is absent or is present with the value false and if STRING contains at
 9 least one character that is in SET, the value of the result is the position of the
 10 leftmost character of STRING that is in SET.

11 *Case (ii):* If BACK is present with the value true and if STRING contains at least one charac-
 12 ter that is in SET, the value of the result is the position of the rightmost character
 13 of STRING that is in SET.

14 *Case (iii):* The value of the result is zero if no character of STRING is in SET or if the length
 15 of STRING or SET is zero.

16 **Examples.**

17 *Case (i):* SCAN ('FORTRAN', 'TR') has the value 3.

18 *Case (ii):* SCAN ('FORTRAN', 'TR', BACK = .TRUE.) has the value 5.

19 *Case (iii):* SCAN ('FORTRAN', 'BCD') has the value 0.

20 13.7.151 SELECTED_BITS_KIND (N)

21 **Description.** Returns a value of the kind type parameter of a bits type with N bits.

22 **Class.** Transformational function.

23 **Argument.** N shall be scalar and of type integer.

24 **Result Characteristics.** Default integer scalar.

25 **Result Value.** The result value is N if the processor supports a bits type with a kind type
 26 parameter equal to N; otherwise the result value is -1 .

27 **Example.** If the value of NUMERIC_STORAGE_SIZE is 32, SELECTED_BITS_KIND (43) has
 28 the value 43.

29 13.7.152 SELECTED_CHAR_KIND (NAME)

30 **Description.** Returns the value of the kind type parameter of the character set named by the
 31 argument.

32 **Class.** Transformational function.

33 **Argument.** NAME shall be scalar and of type default character.

34 **Result Characteristics.** Default integer scalar.

35 **Result Value.** If NAME has the value DEFAULT, then the result has a value equal to that of
 36 the kind type parameter of the default character type. If NAME has the value ASCII, then the
 37 result has a value equal to that of the kind type parameter of the ASCII character type if the

1 processor supports such a type; otherwise the result has the value -1. If NAME has the value
 2 ISO_10646, then the result has a value equal to that of the kind type parameter of the ISO/IEC
 3 10646-1:2000 UCS-4 character type if the processor supports such a type; otherwise the result has
 4 the value -1. If NAME is a processor-defined name of some other character type supported by the
 5 processor, then the result has a value equal to that of the kind type parameter of that character
 6 type. If NAME is not the name of a supported character type, then the result has the value -1.
 7 The NAME is interpreted without respect to case or trailing blanks.

8 **Examples.** SELECTED_CHAR_KIND ('ASCII') has the value 1 on a processor that uses 1 as the
 9 kind type parameter for the ASCII character set. The following subroutine produces a Japanese
 10 date stamp.

```

11     SUBROUTINE create_date_string(string)
12         INTRINSIC date_and_time,selected_char_kind
13         INTEGER,PARAMETER :: ucs4 = selected_char_kind("ISO_10646")
14         CHARACTER(1,UCS4),PARAMETER :: nen=CHAR(INT(Z'5e74'),UCS4), & !year
15             gatsu=CHAR(INT(Z'6708'),UCS4), & !month
16             nichi=CHAR(INT(Z'65e5'),UCS4) !day
17         CHARACTER(len= *, kind= ucs4) string
18         INTEGER values(8)
19         CALL date_and_time(values=values)
20         WRITE(string,1) values(1),nen,values(2),gatsu,values(3),nichi
21     1 FORMAT(IO,A,IO,A,IO,A)
22     END SUBROUTINE"

```

23 13.7.153 SELECTED_INT_KIND (R)

24 **Description.** Returns a value of the kind type parameter of an integer type that represents all
 25 integer values n with $-10^R < n < 10^R$.

26 **Class.** Transformational function.

27 **Argument.** R shall be scalar and of type integer.

28 **Result Characteristics.** Default integer scalar.

29 **Result Value.** The result has a value equal to the value of the kind type parameter of an integer
 30 type that represents all values n in the range $-10^R < n < 10^R$, or if no such kind type parameter
 31 is available on the processor, the result is -1. If more than one kind type parameter meets the
 32 criterion, the value returned is the one with the smallest decimal exponent range, unless there are
 33 several such values, in which case the smallest of these kind values is returned.

34 **Example.** Assume a processor supports two integer kinds, 32 with representation method $r = 2$
 35 and $q = 31$, and 64 with representation method $r = 2$ and $q = 63$. On this processor SELECTED_
 36 INT_KIND(9) has the value 32 and SELECTED_INT_KIND(10) has the value 64.

37 13.7.154 SELECTED_REAL_KIND ([P, R, RADIX])

38 **Description.** Returns a value of the kind type parameter of a real type with decimal precision
 39 of at least P digits, a decimal exponent range of at least R, and a radix of RADIX.

40 **Class.** Transformational function.

1 **Arguments.** At least one argument shall be present.

2 P (optional) shall be scalar and of type integer.

3 R (optional) shall be scalar and of type integer.

4 RADIX (optional) shall be scalar and of type integer.

5 **Result Characteristics.** Default integer scalar.

6 **Result Value.** If P or R is absent, the result value is the same as if it were present with the
7 value zero. If RADIX is absent, there is no requirement on the radix of the selected kind.

8 The result has a value equal to a value of the kind type parameter of a real type with decimal precision,
9 as returned by the function PRECISION, of at least P digits, a decimal exponent range, as returned by
10 the function RANGE, of at least R, and a radix, as returned by the function RADIX, of RADIX, if such
11 a kind type parameter is available on the processor.

12 Otherwise, the result is -1 if the processor supports a real type with radix RADIX and exponent range
13 of at least R but not with precision of at least P, -2 if the processor supports a real type with radix
14 RADIX and precision of at least P but not with exponent range of at least R, -3 if the processor
15 supports a real type with radix RADIX but with neither precision of at least P nor exponent range of
16 at least R, -4 if the processor supports a real type with radix RADIX and either precision of at least
17 P or exponent range of at least R but not both together, and -5 if the processor supports no real type
18 with radix RADIX.

19 If more than one kind type parameter value meets the criteria, the value returned is the one with the
20 smallest decimal precision, unless there are several such values, in which case the smallest of these kind
21 values is returned.

22 **Example.** SELECTED_REAL_KIND (6, 70) has the value KIND (0.0) on a machine that sup-
23 ports a default real approximation method with $b = 16$, $p = 6$, $e_{\min} = -64$, and $e_{\max} = 63$ and
24 does not have a less precise approximation method.

25 13.7.155 SET_EXPONENT (X, I)

26 **Description.** Returns the model number whose fractional part is the fractional part of the model
27 representation of X and whose exponent part is I.

28 **Class.** Elemental function.

29 **Arguments.**

30 X shall be of type real.

31 I shall be of type integer.

32 **Result Characteristics.** Same as X.

33 **Result Value.** The result has the value $X \times b^{I-e}$, where b and e are as defined in 13.4 for the
34 representation for the value of X in the model that has the radix of X but no limits on exponent
35 values. If X has the value zero, the result has the same value as X.

J3 internal note

Unresolved Technical Issue 093

The above is broken for Inf and NaN. I suggest SET_EXPONENT(Inf or NaN,any) = NaN.

36 **Example.** SET_EXPONENT (3.0, 1) has the value 1.5 for reals whose model is as in Note 13.4.

1 13.7.156 SHAPE (SOURCE [, KIND])

2 **Description.** Returns the shape of an array or a scalar.

3 **Class.** Inquiry function.

4 **Arguments.**

SOURCE may be of any type. It may be a scalar or an array. It shall not be an unallocated allocatable or a pointer that is not associated. It shall not be an assumed-size array.

6 KIND (optional) shall be a scalar integer initialization expression.

7 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified by the value of KIND; otherwise the kind type parameter is that of default integer type. The result is an array of rank one whose size is equal to the rank of SOURCE.

10 **Result Value.** The value of the result is the shape of SOURCE.

11 **Examples.** The value of SHAPE (A (2:5, -1:1)) is [4, 3]. The value of SHAPE (3) is the rank-one array of size zero.

13 13.7.157 SHIFTA (I, SHIFT)

14 **Description.** Performs a right shift with fill.

15 **Class.** Elemental function.

16 **Arguments.**

17 I shall be of type integer or bits.

18 SHIFT shall be of type integer. It shall be nonnegative and less than or equal to BIT_SIZE (I).

19 **Result Characteristics.** Same as I.

20 **Result Value.** The result has the value obtained by shifting the bits of I to the right SHIFT bits and replicating the leftmost bit of I in the left SHIFT bits.

22 If SHIFT is zero the result is I. Bits shifted out from the right are lost. The model for the interpretation of an integer value as a sequence of bits is in 13.3.

24 **Example.** SHIFTA (B'10000', 2) has the value B'11100'.

25 13.7.158 SHIFTL (I, SHIFT)

26 **Description.** Performs a left shift.

27 **Class.** Elemental function.

28 **Arguments.**

29 I shall be of type integer or bits.

30 SHIFT shall be of type integer. It shall be nonnegative and less than or equal to BIT_SIZE (I).

31 **Result Characteristics.** Same as I.

1 **Result Value.** The result has a value equal to ISHFT (I, SHIFT).

2 **Examples.** SHIFTL (3, 1) has the value 6. SHIFTL (B'00001', 2) has the value B'00100'.

3 **13.7.159 SHIFTR (I, SHIFT)**

4 **Description.** Performs a right shift.

5 **Class.** Elemental function.

6 **Arguments.**

7 I shall be of type integer or bits.

8 SHIFT shall be of type integer. It shall be nonnegative and less than or equal to
BIT_SIZE (I).

9 **Result Characteristics.** Same as I.

10 **Result Value.** The value of the result is ISHFT (I, -SHIFT).

11 **Examples.** SHIFTR (3, 1) has the value 1. SHIFTR (B'10000', 2) has the value B'00100'.

12 **13.7.160 SIGN (A, B)**

13 **Description.** Magnitude of A with the sign of B.

14 **Class.** Elemental function.

15 **Arguments.**

16 A shall be of type integer or real.

17 B shall be of the same type and kind type parameter as A.

18 **Result Characteristics.** Same as A.

19 **Result Value.**

20 *Case (i):* If $B > 0$, the value of the result is $|A|$.

21 *Case (ii):* If $B < 0$, the value of the result is $-|A|$.

22 *Case (iii):* If B is of type integer and $B=0$, the value of the result is $|A|$.

23 *Case (iv):* If B is of type real and is zero, then

24 (1) If the processor cannot distinguish between positive and negative real zero,
25 the value of the result is $|A|$.

26 (2) If B is positive real zero, the value of the result is $|A|$.

27 (3) If B is negative real zero, the value of the result is $-|A|$.

28 **Example.** SIGN (-3.0, 2.0) has the value 3.0.

29 **13.7.161 SIN (X)**

30 **Description.** Sine function.

31 **Class.** Elemental function.

32 **Argument.** X shall be of type real or complex.

1 **Result Characteristics.** Same as X.

2 **Result Value.** The result has a value equal to a processor-dependent approximation to $\sin(X)$.
 3 If X is of type real, it is regarded as a value in radians. If X is of type complex, its real part is
 4 regarded as a value in radians.

5 **Example.** SIN (1.0) has the value 0.84147098 (approximately).

6 **13.7.162 SINH (X)**

7 **Description.** Hyperbolic sine function.

8 **Class.** Elemental function.

9 **Argument.** X shall be of type real or complex.

10 **Result Characteristics.** Same as X.

11 **Result Value.** The result has a value equal to a processor-dependent approximation to $\sinh(X)$.
 12 If X is of type complex its imaginary part is regarded as a value in radians.

13 **Example.** SINH (1.0) has the value 1.1752012 (approximately).

14 **13.7.163 SIZE (ARRAY [, DIM, KIND])**

15 **Description.** Returns the extent of an array along a specified dimension or the total number of
 16 elements in the array.

17 **Class.** Inquiry function.

18 **Arguments.**

19 ARRAY may be of any type. It shall be an array. It shall not be an unallocated
 allocatable or a pointer that is not associated. If ARRAY is an assumed-size
 array, DIM shall be present with a value less than the rank of ARRAY.

20 DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
 where n is the rank of ARRAY.

21 KIND (optional) shall be a scalar integer initialization expression.

22 **Result Characteristics.** Integer scalar. If KIND is present, the kind type parameter is that
 23 specified by the value of KIND; otherwise the kind type parameter is that of default integer type.

24 **Result Value.** The result has a value equal to the extent of dimension DIM of ARRAY or, if
 25 DIM is absent, the total number of elements of ARRAY.

26 **Examples.** The value of SIZE (A (2:5, -1:1), DIM=2) is 3. The value of SIZE (A (2:5, -1:1)) is
 27 12.

28 **13.7.164 SPACING (X)**

29 **Description.** Returns the absolute spacing of model numbers near the argument value.

30 **Class.** Elemental function.

31 **Argument.** X shall be of type real.

32 **Result Characteristics.** Same as X.

1 **Result Value.** If X does not have the value zero and is not an IEEE infinity or NaN, the result
 2 has the value $b^{\max(e-p, e_{\text{MIN}}-1)}$, where b , e , and p are as defined in 13.4 for the value nearest to X
 3 in the model for real values whose kind type parameter is that of X; if there are two such values
 4 the value of greater absolute value is taken. If X has the value zero, the result is the same as that
 5 of TINY (X). If X is an IEEE infinity, the result is positive infinity. If X is an IEEE NaN, the
 6 result is that NaN.

7 **Example.** SPACING (3.0) has the value 2^{-22} for reals whose model is as in Note 13.4.

8 13.7.165 SPREAD (SOURCE, DIM, NCOPIES)

9 **Description.** Replicates an array by adding a dimension. Broadcasts several copies of SOURCE
 10 along a specified dimension (as in forming a book from copies of a single page) and thus forms an
 11 array of rank one greater.

12 **Class.** Transformational function.

13 **Arguments.**

14 SOURCE may be of any type. It may be a scalar or an array. The rank of SOURCE
 shall be less than 15.

15 DIM shall be scalar and of type integer with value in the range $1 \leq \text{DIM} \leq n + 1$,
 where n is the rank of SOURCE.

16 NCOPIES shall be scalar and of type integer.

17 **Result Characteristics.** The result is an array of the same type and type parameters as
 18 SOURCE and of rank $n + 1$, where n is the rank of SOURCE.

19 *Case (i):* If SOURCE is scalar, the shape of the result is (MAX (NCOPIES, 0)).

20 *Case (ii):* If SOURCE is an array with shape (d_1, d_2, \dots, d_n) , the shape of the result is
 21 $(d_1, d_2, \dots, d_{\text{DIM}-1}, \text{MAX}(\text{NCOPIES}, 0), d_{\text{DIM}}, \dots, d_n)$.

22 **Result Value.**

23 *Case (i):* If SOURCE is scalar, each element of the result has a value equal to SOURCE.

24 *Case (ii):* If SOURCE is an array, the element of the result with subscripts $(r_1, r_2, \dots, r_{n+1})$
 25 has the value SOURCE $(r_1, r_2, \dots, r_{\text{DIM}-1}, r_{\text{DIM}+1}, \dots, r_{n+1})$.

26 **Examples.** If A is the array [2, 3, 4], SPREAD (A, DIM=1, NCOPIES=NC) is the array

27 $\begin{bmatrix} 2 & 3 & 4 \\ 2 & 3 & 4 \\ 2 & 3 & 4 \end{bmatrix}$ if NC has the value 3 and is a zero-sized array if NC has the value 0.

28 13.7.166 SQRT (X)

29 **Description.** Square root.

30 **Class.** Elemental function.

31 **Argument.** X shall be of type real or complex. Unless X is complex, its value shall be greater
 32 than or equal to zero.

33 **Result Characteristics.** Same as X.

34 **Result Value.** The result has a value equal to a processor-dependent approximation to the
 35 square root of X. A result of type complex is the principal value with the real part greater than

1 or equal to zero. When the real part of the result is zero, the imaginary part has the same sign
2 as the imaginary part of X.

3 **Example.** SQRT (4.0) has the value 2.0 (approximately).

4 **13.7.167 STORAGE_SIZE (A [, KIND])**

5 **Description.** Returns the storage size in bits that an array element of the same dynamic type
6 and type parameters of A would have.

7 **Class.** Inquiry function.

8 **Arguments.**

A may be of any type. It may be a scalar or an array. If it is polymorphic it shall
not be an undefined pointer. If it has any deferred type parameters it shall not
be an unallocated allocatable or a disassociated or undefined pointer.

9
10 KIND (optional) shall be a scalar integer initialization expression.

11 **Result Characteristics.** Integer scalar. If KIND is present, the kind type parameter is that
12 specified by the value of KIND; otherwise, the kind type parameter is that of default integer type.

13 **Result Value.** The result value is the size expressed in bits for an element of an array that
14 has the dynamic type and type parameters of A. If the type and type parameters are such that
15 storage association (16.5.3) applies, the result is consistent with the named constants defined in
16 the intrinsic module ISO_FORTRAN_ENV.

NOTE 13.21

An array element might take more bits to store than an isolated scalar, since any hardware-imposed alignment requirements for array elements might not apply to a simple scalar variable.

NOTE 13.22

This is intended to be the size in memory that an object takes when it is stored; this might differ from the size it takes during expression handling (which might be the native register size) or when stored in a file. If an object is never stored in memory but only in a register, this function nonetheless returns the size it would take if it were stored in memory.

17 **Example.** STORAGE_SIZE(1.0) has the same value as the named constant NUMERIC_STOR-
18 AGE_SIZE in the intrinsic module ISO_FORTRAN_ENV.

19 **13.7.168 SUM (ARRAY, DIM [, MASK]) or SUM (ARRAY [, MASK])**

20 **Description.** Sum all the elements of ARRAY along dimension DIM corresponding to the true
21 elements of MASK.

22 **Class.** Transformational function.

23 **Arguments.**

24 ARRAY shall be of type integer, real, or complex. It shall be an array.

DIM shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$,
where n is the rank of ARRAY. The corresponding actual argument shall not
25 be an optional dummy argument.

1 MASK (optional) shall be of type logical and shall be conformable with ARRAY.

2 **Result Characteristics.** The result is of the same type and kind type parameter as ARRAY. It is
3 scalar if DIM is absent; otherwise, the result has rank $n-1$ and shape $(d_1, d_2, \dots, d_{\text{DIM}-1}, d_{\text{DIM}+1},$
4 $\dots, d_n)$ where (d_1, d_2, \dots, d_n) is the shape of ARRAY.

5 **Result Value.**

6 *Case (i):* The result of SUM (ARRAY) has a value equal to a processor-dependent approxi-
7 mation to the sum of all the elements of ARRAY or has the value zero if ARRAY
8 has size zero.

9 *Case (ii):* The result of SUM (ARRAY, MASK = MASK) has a value equal to a processor-
10 dependent approximation to the sum of the elements of ARRAY corresponding to
11 the true elements of MASK or has the value zero if there are no true elements.

12 *Case (iii):* If ARRAY has rank one, SUM (ARRAY, DIM = DIM [, MASK = MASK]) has a
13 value equal to that of SUM (ARRAY [, MASK = MASK]). Otherwise, the value
14 of element $(s_1, s_2, \dots, s_{\text{DIM}-1}, s_{\text{DIM}+1}, \dots, s_n)$ SUM (ARRAY, DIM = DIM [
15 , MASK = MASK]) is equal to

16 $\text{SUM (ARRAY } (s_1, s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n) [, \text{ MASK= MASK } (s_1,$
17 $s_2, \dots, s_{\text{DIM}-1}, :, s_{\text{DIM}+1}, \dots, s_n)])$.

18 **Examples.**

19 *Case (i):* The value of SUM ((/ 1, 2, 3 /)) is 6.

20 *Case (ii):* SUM (C, MASK= C > 0.0) forms the sum of the positive elements of C.

21 *Case (iii):* If B is the array $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix}$, SUM (B, DIM = 1) is [3, 7, 11] and SUM (B, DIM =
22 2) is [9, 12].

23 13.7.169 SYSTEM_CLOCK ([COUNT, COUNT_RATE, COUNT_MAX])

24 **Description.** Returns numeric data from a real-time clock.

25 **Class.** Subroutine.

26 **Arguments.**

27 COUNT shall be scalar and of type integer. It is an INTENT (OUT) argument. It is
(optional) assigned a processor-dependent value based on the current value of the proces-
sor clock, or -HUGE (COUNT) if there is no clock. The processor-dependent
value is incremented by one for each clock count until the value COUNT_MAX
is reached and is reset to zero at the next count. It lies in the range 0 to
COUNT_MAX if there is a clock.

28 COUNT_RATE shall be scalar and of type integer or real. It is an INTENT (OUT) argument.
(optional) It is assigned a processor-dependent approximation to the number of processor
clock counts per second, or zero if there is no clock.

29 COUNT_MAX shall be scalar and of type integer. It is an INTENT(OUT) argument. It is
(optional) assigned the maximum value that COUNT can have, or zero if there is no clock.

30 **Example.** If the processor clock is a 24-hour clock that registers time at approximately 18.20648193
31 ticks per second, at 11:30 A.M. the reference

32 CALL SYSTEM_CLOCK (COUNT = C, COUNT_RATE = R, COUNT_MAX = M)

1 defines $C = (11 \times 3600 + 30 \times 60) \times 18.20648193 = 753748$, $R = 18.20648193$, and $M = 24 \times$
 2 $3600 \times 18.20648193 - 1 = 1573039$.

3 **13.7.170 TAN (X)**

4 **Description.** Tangent function.

5 **Class.** Elemental function.

6 **Argument.** X shall be of type real or complex.

7 **Result Characteristics.** Same as X.

8 **Result Value.** The result has a value equal to a processor-dependent approximation to $\tan(X)$.
 9 If X is of type real, it is regarded as a value in radians. If X is of type complex, its real part is
 10 regarded as a value in radians.

11 **Example.** TAN (1.0) has the value 1.5574077 (approximately).

12 **13.7.171 TANH (X)**

13 **Description.** Hyperbolic tangent function.

14 **Class.** Elemental function.

15 **Argument.** X shall be of type real or complex.

16 **Result Characteristics.** Same as X.

17 **Result Value.** The result has a value equal to a processor-dependent approximation to $\tanh(X)$.
 18 If X is of type complex its imaginary part is regarded as a value in radians.

19 **Example.** TANH (1.0) has the value 0.76159416 (approximately).

20 **13.7.172 TEAM_IMAGES (TEAM)**

21 **Description.** Returns a list of the images in a team.

22 **Class.** Transformational function.

23 **Argument.** TEAM shall be a scalar of type IMAGE_TEAM (13.8.3.7).

24 **Result Characteristics.** The result is a default integer array of rank one and of size equal to
 25 the number of images in the team identified by TEAM.

26 **Result Value.** The result is a rank-one array whose element values are the image indices of the
 27 images in the team identified by TEAM.

J3 internal note

Unresolved Technical Issue 085

Is the order of images (to SYNC_IMAGES and FORM_TEAM) significant?

If image 7 does FORM_TEAM with [7,8] and image 8 does FORM_TEAM with [8,7] is this ok?

What order do they come back from TEAM_IMAGES?

The editor recommends “no”, “yes”, and “monotonic increasing”.

28 **Examples.** If the value of TEAM was defined by the statement

29 CALL FORM_TEAM(TEAM, [1,2,3,4])

1 then TEAM_IMAGES (TEAM) has the value [1, 2, 3, 4]. The value of TEAM_IMAGES (NULL-
2 IMAGE_TEAM) is an array of size zero.

3 13.7.173 THIS_IMAGE () or THIS_IMAGE (CO_ARRAY [, DIM])

4 **Description.** Returns the index of the invoking image, a single co-subscript, or a list of co-
5 subscripts.

6 **Class.** Inquiry function.

7 **Arguments.**

8 CO_ARRAY shall be a co-array and may be of any type.
(optional)

9 DIM (optional) shall be scalar and of type default integer. Its value shall be in the range
10 $1 \leq \text{DIM} \leq n$, where n is the co-rank of CO_ARRAY. It shall not be an absent
11 optional dummy argument.

12 **Result Characteristics.** Type default integer. It is scalar if CO_ARRAY is absent or DIM is
13 present; otherwise, the result has rank one and its size is equal to the co-rank of CO_ARRAY.

14 **Result Value.**

15 *Case (i):* The result of THIS_IMAGE () is a scalar with a value equal to the index of the
16 invoking image.

17 *Case (ii):* The result of THIS_IMAGE (CO_ARRAY) is the sequence of co-subscript values
18 for CO_ARRAY that would specify the invoking image.

19 *Case (iii):* The result of THIS_IMAGE (CO_ARRAY, DIM) is the value of co-subscript DIM in
20 the sequence of co-subscript values for CO_ARRAY that would specify the invoking
21 image.

22 **Examples.** If A is declared by the statement REAL A (10, 20) [10, 0:9, 0:*]

23 then on image 5, THIS_IMAGE () has the value 5 and THIS_IMAGE (A) has the value [5, 0, 0].
24 For the same co-array on image 213, THIS_IMAGE (A) has the value [3, 1, 2].

25 The following code uses image 1 to read data. The other images then copy the data.

```
26 IF (THIS_IMAGE()==1) READ (*,*) P
27 SYNC ALL
28 P = P[1]
```

NOTE 13.23

For an example of a module that implements a function similar to the intrinsic IMAGE_INDEX,
see subclause C.10.1.

27 13.7.174 TINY (X)

28 **Description.** Returns the smallest positive number of the model representing numbers of the
29 same type and kind type parameter as the argument.

30 **Class.** Inquiry function.

Argument. X shall be of type real. It may be a scalar or an array.

1

2 **Result Characteristics.** Scalar with the same type and kind type parameter as X.3 **Result Value.** The result has the value $b^{e_{\min}-1}$ where b and e_{\min} are as defined in 13.4 for the
4 model representing numbers of the same type and kind type parameter as X.5 **Example.** TINY (X) has the value 2^{-127} for real X whose model is as in Note 13.4.6 **13.7.175 TRAILZ (I)**7 **Description.** Count the number of trailing zero bits.8 **Class.** Elemental function.9 **Argument.** I shall be of type integer or bits.10 **Result Characteristics.** Default integer.11 **Result Value.** If all of the bits of I are zero, the result value is BIT_SIZE (I). Otherwise, the
12 result value is the position of the rightmost 1 bit in I. The model for the interpretation of an
13 integer value as a sequence of bits is in 13.3.14 **Examples.** TRAILZ (8) has the value 3. TRAILZ (B'101101000') has the value 3.15 **13.7.176 TRANSFER (SOURCE, MOLD [, SIZE])**16 **Description.** Returns a result with a physical representation identical to that of SOURCE but
17 interpreted with the type and type parameters of MOLD.18 **Class.** Transformational function.19 **Arguments.**

20 SOURCE may be of any type. It may be a scalar or an array.

21 MOLD may be of any type. It may be a scalar or an array. If it is a variable, it need
 not be defined.22 SIZE (optional) shall be scalar and of type integer. The corresponding actual argument shall
 not be an optional dummy argument.23 **Result Characteristics.** The result is of the same type and type parameters as MOLD.24 *Case (i):* If MOLD is a scalar and SIZE is absent, the result is a scalar.25 *Case (ii):* If MOLD is an array and SIZE is absent, the result is an array and of rank one. Its
26 size is as small as possible such that its physical representation is not shorter than
27 that of SOURCE.28 *Case (iii):* If SIZE is present, the result is an array of rank one and size SIZE.29 **Result Value.** If the physical representation of the result has the same length as that of SOURCE,
30 the physical representation of the result is that of SOURCE. If the physical representation of the
31 result is longer than that of SOURCE, the physical representation of the leading part is that of
32 SOURCE and the remainder is processor dependent. If the physical representation of the result
33 is shorter than that of SOURCE, the physical representation of the result is the leading part
34 of SOURCE. If D and E are scalar variables such that the physical representation of D is as
35 long as or longer than that of E, the value of TRANSFER (TRANSFER (E, D), E) shall be the
36 value of E. IF D is an array and E is an array of rank one, the value of TRANSFER (TRANS-
 FER (E, D), E, SIZE (E)) shall be the value of E.

1 **Examples.**

2 *Case (i):* TRANSFER (1082130432, 0.0) has the value 4.0 on a processor that represents the
 3 values 4.0 and 1082130432 as the string of binary digits 0100 0000 1000 0000 0000
 4 0000 0000 0000.

5 *Case (ii):* TRANSFER ((/ 1.1, 2.2, 3.3 /), (/ (0.0, 0.0) /)) is a complex rank-one array of
 6 length two whose first element has the value (1.1, 2.2) and whose second element
 7 has a real part with the value 3.3. The imaginary part of the second element is
 8 processor dependent.

9 *Case (iii):* TRANSFER ((/ 1.1, 2.2, 3.3 /), (/ (0.0, 0.0) /), 1) is a complex rank-one array of
 10 length one whose only element has the value (1.1, 2.2).

11 **13.7.177 TRANSPOSE (MATRIX)**

12 **Description.** Transpose an array of rank two.

13 **Class.** Transformational function.

14 **Argument.** MATRIX may be of any type and shall have rank two.

15 **Result Characteristics.** The result is an array of the same type and type parameters as
 16 MATRIX and with rank two and shape (n, m) where (m, n) is the shape of MATRIX.

17 **Result Value.** Element (i, j) of the result has the value MATRIX (j, i) , $i = 1, 2, \dots, n$;
 18 $j = 1, 2, \dots, m$.

19 **Example.** If A is the array $\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$, then TRANSPOSE (A) has the value $\begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix}$.

20 **13.7.178 TRIM (STRING)**

21 **Description.** Returns the argument with trailing blank characters removed.

22 **Class.** Transformational function.

23 **Argument.** STRING shall be of type character and shall be scalar.

24 **Result Characteristics.** Character with the same kind type parameter value as STRING and
 25 with a length that is the length of STRING less the number of trailing blanks in STRING.

26 **Result Value.** The value of the result is the same as STRING except any trailing blanks are
 27 removed. If STRING contains no nonblank characters, the result has zero length.

28 **Example.** TRIM (' A B ') has the value ' A B'.

29 **13.7.179 UBOUND (ARRAY [, DIM, KIND])**

30 **Description.** Returns all the upper bounds of an array or a specified upper bound.

31 **Class.** Inquiry function.

32 **Arguments.**

ARRAY may be of any type. It shall be an array. It shall not be an unallocated
 allocatable or a pointer that is not associated. If ARRAY is an assumed-size
 33 array, DIM shall be present with a value less than the rank of ARRAY.

DIM (optional) shall be scalar and of type integer with a value in the range $1 \leq \text{DIM} \leq n$, where n is the rank of ARRAY. The corresponding actual argument shall not be an optional dummy argument.

KIND (optional) shall be a scalar integer initialization expression.

Result Characteristics. Integer. If KIND is present, the kind type parameter is that specified by the value of KIND; otherwise the kind type parameter is that of default integer type. The result is scalar if DIM is present; otherwise, the result is an array of rank one and size n , where n is the rank of ARRAY.

Result Value.

Case (i): For an array section or for an array expression, other than a whole array or array structure component, UBOUND (ARRAY, DIM) has a value equal to the number of elements in the given dimension; otherwise, it has a value equal to the upper bound for subscript DIM of ARRAY if dimension DIM of ARRAY does not have size zero and has the value zero if dimension DIM has size zero.

Case (ii): UBOUND (ARRAY) has a value whose i th element is equal to UBOUND (ARRAY, i), for $i = 1, 2, \dots, n$, where n is the rank of ARRAY.

Examples. If A is declared by the statement

```
REAL A (2:3, 7:10)
```

then UBOUND (A) is [3, 10] and UBOUND (A, DIM = 2) is 10.

13.7.180 UNPACK (VECTOR, MASK, FIELD)

Description. Unpack an array of rank one into an array under the control of a mask.

Class. Transformational function.

Arguments.

VECTOR may be of any type. It shall have rank one. Its size shall be at least t where t is the number of true elements in MASK.

MASK shall be an array of type logical.

FIELD shall be of the same type and type parameters as VECTOR and shall be conformable with MASK.

Result Characteristics. The result is an array of the same type and type parameters as VECTOR and the same shape as MASK.

Result Value. The element of the result that corresponds to the i th true element of MASK, in array element order, has the value VECTOR (i) for $i = 1, 2, \dots, t$, where t is the number of true values in MASK. Each other element has a value equal to FIELD if FIELD is scalar or to the corresponding element of FIELD if it is an array.

Examples. Particular values may be “scattered” to particular positions in an array by using UNPACK. If M is the array $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$, V is the array [1, 2, 3], and Q is the logical

mask $\begin{bmatrix} . & T & . \\ T & . & . \\ . & . & T \end{bmatrix}$, where “T” represents true and “.” represents false, then the result of

1 UNPACK (V, MASK = Q, FIELD = M) has the value $\begin{bmatrix} 1 & 2 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 3 \end{bmatrix}$ and the result of UN-
 2 PACK (V, MASK = Q, FIELD = 0) has the value $\begin{bmatrix} 0 & 2 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 3 \end{bmatrix}$.

3 13.7.181 **VERIFY (STRING, SET [, BACK, KIND])**

4 **Description.** Verify that a set of characters contains all the characters in a string by identifying
 5 the position of the first character in a string of characters that does not appear in a given set of
 6 characters.

7 **Class.** Elemental function.

8 **Arguments.**

9 STRING shall be of type character.

10 SET shall be of type character with the same kind type parameter as STRING.

11 BACK (optional) shall be of type logical.

12 KIND (optional) shall be a scalar integer initialization expression.

13 **Result Characteristics.** Integer. If KIND is present, the kind type parameter is that specified
 14 by the value of KIND; otherwise the kind type parameter is that of default integer type.

15 **Result Value.**

16 *Case (i):* If BACK is absent or has the value false and if STRING contains at least one
 17 character that is not in SET, the value of the result is the position of the leftmost
 18 character of STRING that is not in SET.

19 *Case (ii):* If BACK is present with the value true and if STRING contains at least one char-
 20 acter that is not in SET, the value of the result is the position of the rightmost
 21 character of STRING that is not in SET.

22 *Case (iii):* The value of the result is zero if each character in STRING is in SET or if STRING
 23 has zero length.

24 **Examples.**

25 *Case (i):* VERIFY ('ABBA', 'A') has the value 2.

26 *Case (ii):* VERIFY ('ABBA', 'A', BACK = .TRUE.) has the value 3.

27 *Case (iii):* VERIFY ('ABBA', 'AB') has the value 0.

28 13.8 **Standard intrinsic modules**

29 13.8.1 **General**

30 This part of ISO/IEC 1539 defines five standard intrinsic modules: a Fortran environment module, a
 31 set of three modules to support exception handling and IEEE arithmetic, and a module to support
 32 interoperability with the C programming language.

33 A processor may extend the standard intrinsic modules to provide public entities in them in addition to
 34 those specified in this standard.

NOTE 13.24

To avoid potential name conflicts with program entities, it is recommended that a program use the ONLY option in any USE statement that references a standard intrinsic module.

1 13.8.2 The IEEE intrinsic modules

2 The IEEE_EXCEPTIONS, IEEE_ARITHMETIC, and IEEE_FEATURES intrinsic modules are describ-
3 ed in Clause 14.

4 13.8.3 The ISO_FORTRAN_ENV intrinsic module**5 13.8.3.1 General**

6 The intrinsic module ISO_FORTRAN_ENV provides public entities relating to the Fortran environment.

7 The processor shall provide the named constants, derived type, and procedures described in subclause
8 13.8.3.

9 13.8.3.2 CHARACTER_STORAGE_SIZE

10 The value of the default integer scalar constant CHARACTER_STORAGE_SIZE is the size expressed
11 in bits of the character storage unit (16.5.3.2).

12 13.8.3.3 COMPILER_OPTIONS ()

13 **Description.** Returns a processor-dependent string describing the options that controlled the
14 program translation phase.

15 **Class.** Inquiry function.

16 **Argument.** None.

17 **Result Characteristics.** Default character scalar with processor-dependent length.

18 **Result Value.** A processor-dependent value which describes the options that controlled the
19 translation phase of program execution.

20 **Example.** COMPILER_OPTIONS () might have the value '/OPTIMIZE /FLOAT=IEEE'.

21 13.8.3.4 COMPILER_VERSION ()

22 **Description.** Returns a processor-dependent string identifying the program translation phase.

23 **Class.** Inquiry function.

24 **Argument.** None.

25 **Result Characteristics.** Default character scalar with processor-dependent length.

26 **Result Value.** A processor-dependent value that identifies the name and version of the program
27 translation phase of the processor.

28 **Example.** COMPILER_VERSION () might have the value 'Fast KL-10 Compiler Version 7'.

NOTE 13.25

For both `COMPILER_OPTIONS` and `COMPILER_VERSION` the processor should include relevant information that could be useful in solving problems found long after the translation phase. For example, compiler release and patch level, default compiler arguments, environment variable values, and run time library requirements might be included. A processor might include this information in an object file automatically, without the user needing to save the result of this function in a variable.

1 **13.8.3.5 ERROR_UNIT**

2 The value of the default integer scalar constant `ERROR_UNIT` identifies the processor-dependent pre-
 3 connected external unit used for the purpose of error reporting (9.4). This unit may be the same as
 4 `OUTPUT_UNIT`. The value shall not be `-1`.

5 **13.8.3.6 FILE_STORAGE_SIZE**

6 The value of the default integer scalar constant `FILE_STORAGE_SIZE` is the size expressed in bits of
 7 the file storage unit (9.2.4).

8 **13.8.3.7 IMAGE_TEAM**

9 A scalar object of type `IMAGE_TEAM` identifies a team of images. It shall have only private components.

J3 internal note

Unresolved Technical Issue 090
 Is `IMAGE_TEAM` a sequence type?
 Is it a `BIND(C)` type?
 Is it extensible?
 (The answer to the last question is a consequence of the answers to the first two questions.)
 Does it have pointer components?
 Does it have allocatable components?
 Just saying “private” is simply not good enough – the processor needs to know.

J3 internal note

Unresolved Technical Issue 073
`IMAGE_TEAM` is not adequately explained.

> The “special”
 > derived types `IMAGE_TEAM`, `C_PTR`, and `C_FUNPTR` are very likely to contain
 > information that is specific to the image where it was defined.

That argument applies to *ALL* values, not just these.

> Use of these data on a different image could lead to various errors.

“Could” lead to “various” errors. Not “inevitably” leads to a “fatal” error.
 Yes, there are certainly valid cross-image uses. Some of these are even obvious (diagnostic information for example).

10 **13.8.3.8 INPUT_UNIT**

11 The value of the default integer scalar constant `INPUT_UNIT` identifies the same processor-dependent
 12 preconnected external unit as the one identified by an asterisk in a `READ` statement (9.4). The value

1 shall not be -1 .

2 13.8.3.9 INT8, INT16, INT32, and INT64

3 The values of these default integer scalar named constants shall be those of the kind type parameters that
4 specify an INTEGER type whose storage size expressed in bits is 8, 16, 32, and 64 respectively. If, for
5 any of these constants, the processor supports more than one kind of that size, it is processor-dependent
6 which kind value is provided. If the processor supports no kind of a particular size, that constant shall
7 be equal to -2 if the processor supports kinds of a larger size and -1 otherwise.

8 13.8.3.10 IOSTAT_END

9 The value of the default integer scalar constant IOSTAT_END is assigned to the variable specified in
10 an IOSTAT= specifier (9.10.4) if an end-of-file condition occurs during execution of an input/output
11 statement and no error condition occurs. This value shall be negative.

12 13.8.3.11 IOSTAT_EOR

13 The value of the default integer scalar constant IOSTAT_EOR is assigned to the variable specified in
14 an IOSTAT= specifier (9.10.4) if an end-of-record condition occurs during execution of an input/output
15 statement and no end-of-file or error condition occurs. This value shall be negative and different from
16 the value of IOSTAT_END.

17 13.8.3.12 IOSTAT_INQUIRE_INTERNAL_UNIT

18 The value of the default integer scalar constant IOSTAT_INQUIRE_INTERNAL_UNIT is assigned to
19 the variable specified in an IOSTAT= specifier in an INQUIRE statement (9.9) if a *file-unit-number*
20 identifies an internal unit in that statement.

NOTE 13.26

This can only occur when a user defined derived type input/output procedure is called by the processor as the result of executing a parent data transfer statement for an internal unit.

21 13.8.3.13 NULL_IMAGE_TEAM

22 The value of the TYPE(IMAGE_TEAM) scalar constant NULL_IMAGE_TEAM identifies a team with
23 no images.

J3 internal note

Unresolved Technical Issue 087

This constant appears to be worthless and useless. The only even semi-plausible use of it I can think of is for things like

```
TYPE(IMAGE_TEAM) :: X = NULL_IMAGE_TEAM
```

but I don't think this is in fact useful, since there is no comparison operation for this type. Even the immeasurably small functionality offered by this would be *far* better provided by default-initializing IMAGE_TEAM to be empty image sets.

Having this constant raises the question of what happens when you use it for anything. Currently this makes the standard crash (the words just don't make sense) rather than unambiguously making the program nonconforming or in error.

Oh, and by the way, it's a crap name.

24 13.8.3.14 NUMERIC_STORAGE_SIZE

1 The value of the default integer scalar constant `NUMERIC_STORAGE_SIZE` is the size expressed in
2 bits of the numeric storage unit (16.5.3.2).

3 **13.8.3.15 OUTPUT_UNIT**

4 The value of the default integer scalar constant `OUTPUT_UNIT` identifies the same processor-dependent
5 preconnected external unit as the one identified by an asterisk in a `WRITE` statement (9.4). The value
6 shall not be `-1`.

7 **13.8.3.16 REAL32, REAL64, and REAL128**

8 The values of these default integer scalar named constants shall be those of the kind type parameters
9 that specify a `REAL` type whose storage size expressed in bits is 32, 64, and 128 respectively. If, for
10 any of these constants, the processor supports more than one kind of that size, it is processor-dependent
11 which kind value is provided. If the processor supports no kind of a particular size, that constant shall
12 be equal to `-2` if the processor supports kinds of a larger size and `-1` otherwise.

13 **13.8.4 The ISO_C_BINDING intrinsic module**

14 The `ISO_C_BINDING` intrinsic module is described in Clause 15.

1 14 Exceptions and IEEE arithmetic

2 The intrinsic modules IEEE_EXCEPTIONS, IEEE_ARITHMETIC, and IEEE_FEATURES provide sup-
3 port for exceptions and IEEE arithmetic. Whether the modules are provided is processor dependent.
4 If the module IEEE_FEATURES is provided, which of the named constants defined in this standard
5 are included is processor dependent. The module IEEE_ARITHMETIC behaves as if it contained a
6 USE statement for IEEE_EXCEPTIONS; everything that is public in IEEE_EXCEPTIONS is public in
7 IEEE_ARITHMETIC.

NOTE 14.1

The types and procedures defined in these modules are not themselves intrinsic.

8 If IEEE_EXCEPTIONS or IEEE_ARITHMETIC is accessible in a scoping unit, IEEE_OVERFLOW and
9 IEEE_DIVIDE_BY_ZERO are supported in the scoping unit for all kinds of real and complex data. Which
10 other exceptions are supported can be determined by the function IEEE_SUPPORT_FLAG (14.10.26);
11 whether control of halting is supported can be determined by the function IEEE_SUPPORT_HALTING.
12 The extent of support of the other exceptions may be influenced by the accessibility of the named con-
13 stants IEEE_INEXACT_FLAG, IEEE_INVALID_FLAG, and IEEE_UNDERFLOW_FLAG of the module
14 IEEE_FEATURES. If a scoping unit has access to IEEE_UNDERFLOW_FLAG of IEEE_FEATURES,
15 within the scoping unit the processor shall support underflow and return true from IEEE_SUPPORT_-
16 FLAG(IEEE_UNDERFLOW, X) for at least one kind of real. Similarly, if IEEE_INEXACT_FLAG or
17 IEEE_INVALID_FLAG is accessible, within the scoping unit the processor shall support the exception
18 and return true from the corresponding inquiry for at least one kind of real. Also, if IEEE_HALTING
19 is accessible, within the scoping unit the processor shall support control of halting and return true from
20 IEEE_SUPPORT_HALTING(FLAGS) for the flag.

NOTE 14.2

IEEE_INVALID is not required to be supported whenever IEEE_EXCEPTIONS is accessed. This is to allow a non-IEEE processor to provide support for overflow and divide_by_zero. On an IEEE machine, invalid is an equally serious condition.

NOTE 14.3

The IEEE_FEATURES module is provided to allow a reasonable amount of cooperation between the programmer and the processor in controlling the extent of IEEE arithmetic support. On some processors some IEEE features are natural for the processor to support, others may be inefficient at run time, and others are essentially impossible to support. If IEEE_FEATURES is not used, the processor will support only the natural operations. Within IEEE_FEATURES the processor will define the named constants (14.1) corresponding to the time-consuming features (as well as the natural ones for completeness) but will not define named constants corresponding to the impossible features. If the programmer accesses IEEE_FEATURES, the processor shall provide support for all of the IEEE_FEATURES that are reasonably possible. If the programmer uses an ONLY option on a USE statement to access a particular feature name, the processor shall provide support for the corresponding feature, or issue an error message saying the name is not defined in the module.

When used this way, the named constants in the IEEE_FEATURES are similar to what are frequently called command line switches for the compiler. They can specify compilation options in a reasonably portable manner.

1 If a scoping unit does not access IEEE_FEATURES, IEEE_EXCEPTIONS, or IEEE_ARITHMETIC,
 2 the level of support is processor dependent, and need not include support for any exceptions. If a flag is
 3 signaling on entry to such a scoping unit, the processor ensures that it is signaling on exit. If a flag is
 4 quiet on entry to such a scoping unit, whether it is signaling on exit is processor dependent.

5 Further IEEE support is available through the module IEEE_ARITHMETIC. The extent of support may
 6 be influenced by the accessibility of the named constants of the module IEEE_FEATURES. If a scoping
 7 unit has access to IEEE_DATATYPE of IEEE_FEATURES, within the scoping unit the processor shall
 8 support IEEE arithmetic and return true from IEEE_SUPPORT_DATATYPE(X) (14.10.23) for at least
 9 one kind of real. Similarly, if IEEE_DENORMAL, IEEE_DIVIDE, IEEE_INF, IEEE_NAN, IEEE_-
 10 ROUNDING, or IEEE_SQRT is accessible, within the scoping unit the processor shall support the feature
 11 and return true from the corresponding inquiry function for at least one kind of real. In the case of
 12 IEEE_ROUNDING, it shall return true for all the rounding modes IEEE_NEAREST, IEEE_TO_ZERO,
 13 IEEE_UP, and IEEE_DOWN.

14 Execution might be slowed on some processors by the support of some features. If IEEE_EXCEPTIONS
 15 or IEEE_ARITHMETIC is accessed but IEEE_FEATURES is not accessed, the supported subset of fea-
 16 tures is processor dependent. The processor's fullest support is provided when all of IEEE_FEATURES
 17 is accessed as in

```
18 USE, INTRINSIC :: IEEE_ARITHMETIC; USE, INTRINSIC :: IEEE_FEATURES
```

19 but execution might then be slowed by the presence of a feature that is not needed. In all cases, the
 20 extent of support can be determined by the inquiry functions.

21 14.1 Derived types and constants defined in the modules

22 The modules IEEE_EXCEPTIONS, IEEE_ARITHMETIC, and IEEE_FEATURES define five derived
 23 types, whose components are all private.

24 The module IEEE_EXCEPTIONS defines the following types.

- 25 • IEEE_FLAG_TYPE is for identifying a particular exception flag. Its only possible values are those
 26 of named constants defined in the module: IEEE_INVALID, IEEE_OVERFLOW, IEEE_DIVIDE_-
 27 BY_ZERO, IEEE_UNDERFLOW, and IEEE_INEXACT. The module also defines the array named
 28 constants IEEE_USUAL = (/ IEEE_OVERFLOW, IEEE_DIVIDE_BY_ZERO, IEEE_INVALID /)
 29 and IEEE_ALL = (/ IEEE_USUAL, IEEE_UNDERFLOW, IEEE_INEXACT /).
- 30 • IEEE_STATUS_TYPE is for saving the current floating point status.

31 The module IEEE_ARITHMETIC defines the following.

- 32 • The type IEEE_CLASS_TYPE, for identifying a class of floating-point values. Its only possible
 33 values are those of named constants defined in the module: IEEE_SIGNALING_NAN, IEEE_QUI-
 34 ET_NAN, IEEE_NEGATIVE_INF, IEEE_NEGATIVE_NORMAL, IEEE_NEGATIVE_DENORM-
 35 AL, IEEE_NEGATIVE_ZERO, IEEE_POSITIVE_ZERO, IEEE_POSITIVE_DENORMAL, IEEE_-
 36 POSITIVE_NORMAL, IEEE_POSITIVE_INF, and IEEE_OTHER_VALUE.
- 37 • The type IEEE_ROUND_TYPE, for identifying a particular rounding mode. Its only possible
 38 values are those of named constants defined in the module: IEEE_NEAREST, IEEE_TO_ZERO,
 39 IEEE_UP, and IEEE_DOWN for the IEEE modes; and IEEE_OTHER for any other mode.

- 1 • The elemental operator `==` for two values of one of these types to return true if the values are the
2 same and false otherwise.
- 3 • The elemental operator `/=` for two values of one of these types to return true if the values differ
4 and false otherwise.

5 The module `IEEE_FEATURES` defines the type `IEEE_FEATURES_TYPE`, for expressing the need
6 for particular IEEE features. Its only possible values are those of named constants defined in the
7 module: `IEEE_DATATYPE`, `IEEE_DENORMAL`, `IEEE_DIVIDE`, `IEEE_HALTING`, `IEEE_INEXACT_-`
8 `FLAG`, `IEEE_INF`, `IEEE_INVALID_FLAG`, `IEEE_NAN`, `IEEE_ROUNDING`, `IEEE_SQRT`, and `IEEE_-`
9 `UNDERFLOW_FLAG`.

10 14.2 The exceptions

11 The exceptions are the following.

- 12 • `IEEE_OVERFLOW` occurs when the result for an intrinsic real operation or assignment has an
13 absolute value greater than a processor-dependent limit, or the real or imaginary part of the result
14 for an intrinsic complex operation or assignment has an absolute value greater than a processor-
15 dependent limit.
- 16 • `IEEE_DIVIDE_BY_ZERO` occurs when a real or complex division has a nonzero numerator and a
17 zero denominator.
- 18 • `IEEE_INVALID` occurs when a real or complex operation or assignment is invalid; possible examples
19 are `SQRT(X)` when `X` is real and has a nonzero negative value, and conversion to an integer (by
20 assignment, an intrinsic procedure, or a procedure defined in an intrinsic module) when the result
21 is too large to be representable.
- 22 • `IEEE_UNDERFLOW` occurs when the result for an intrinsic real operation or assignment has an
23 absolute value less than a processor-dependent limit and loss of accuracy is detected, or the real
24 or imaginary part of the result for an intrinsic complex operation or assignment has an absolute
25 value less than a processor-dependent limit and loss of accuracy is detected.
- 26 • `IEEE_INEXACT` occurs when the result of a real or complex operation or assignment is not exact.

27 Each exception has a flag whose value is either quiet or signaling. The value can be determined by
28 the function `IEEE_GET_FLAG`. Its initial value is quiet and it signals when the associated exception
29 occurs. Its status can also be changed by the subroutine `IEEE_SET_FLAG` or the subroutine `IEEE_-`
30 `SET_STATUS`. Once signaling within a procedure, it remains signaling unless set quiet by an invocation
31 of the subroutine `IEEE_SET_FLAG` or the subroutine `IEEE_SET_STATUS`.

32 If a flag is signaling on entry to a procedure other than `IEEE_GET_FLAG` or `IEEE_GET_STATUS`, the
33 processor will set it to quiet on entry and restore it to signaling on return.

NOTE 14.4

If a flag signals during execution of a procedure, the processor shall not set it to quiet on return.

34 Evaluation of a specification expression may cause an exception to signal.

35 In a scoping unit that has access to `IEEE_EXCEPTIONS` or `IEEE_ARITHMETIC`, if an intrinsic pro-
36 cedure or a procedure defined in an intrinsic module executes normally, the values of the flags `IEEE_-`
37 `OVERFLOW`, `IEEE_DIVIDE_BY_ZERO`, and `IEEE_INVALID` shall be as on entry to the procedure,

1 even if one or more signals during the calculation. If a real or complex result is too large for the pro-
2 cedure to handle, IEEE_OVERFLOW may signal. If a real or complex result is a NaN because of an
3 invalid operation (for example, LOG(-1.0)), IEEE_INVALID may signal. Similar rules apply to format
4 processing and to intrinsic operations: no signaling flag shall be set quiet and no quiet flag shall be set
5 signaling because of an intermediate calculation that does not affect the result.

6 In a sequence of statements that has no invocations of IEEE_GET_FLAG, IEEE_SET_FLAG, IEEE_-
7 GET_STATUS, IEEE_SET_HALTING_MODE, or IEEE_SET_STATUS, if the execution of an operation
8 would cause an exception to signal but after execution of the sequence no value of a variable depends
9 on the operation, whether the exception is signaling is processor dependent. For example, when Y has
10 the value zero, whether the code

```
11     X = 1.0/Y  
12     X = 3.0
```

13 signals IEEE_DIVIDE_BY_ZERO is processor dependent. Another example is the following:

```
14     REAL, PARAMETER :: X=0.0, Y=6.0  
15     IF (1.0/X == Y) PRINT *, 'Hello world'
```

16 where the processor is permitted to discard the IF statement because the logical expression can never
17 be true and no value of a variable depends on it.

18 An exception shall not signal if this could arise only during execution of an operation beyond those
19 required or permitted by the standard. For example, the statement

```
20     IF (F(X)>0.0) Y = 1.0/Z
```

21 shall not signal IEEE_DIVIDE_BY_ZERO when both F(X) and Z are zero and the statement

```
22     WHERE(A>0.0) A = 1.0/A
```

23 shall not signal IEEE_DIVIDE_BY_ZERO. On the other hand, when X has the value 1.0 and Y has the
24 value 0.0, the expression

```
25     X>0.00001 .OR. X/Y>0.00001
```

26 is permitted to cause the signaling of IEEE_DIVIDE_BY_ZERO.

27 The processor need not support IEEE_INVALID, IEEE_UNDERFLOW, and IEEE_INEXACT. If an
28 exception is not supported, its flag is always quiet. The function IEEE_SUPPORT_FLAG can be used
29 to inquire whether a particular flag is supported.

30 14.3 The rounding modes

31 The IEEE International Standard specifies four rounding modes.

- 1 • IEEE_NEAREST rounds the exact result to the nearest representable value.
- 2 • IEEE_TO_ZERO rounds the exact result towards zero to the next representable value.
- 3 • IEEE_UP rounds the exact result towards +infinity to the next representable value.
- 4 • IEEE_DOWN rounds the exact result towards -infinity to the next representable value.

5 The function IEEE_GET_ROUNDING_MODE can be used to inquire which rounding mode is in opera-
6 tion. Its value is one of the above four or IEEE_OTHER if the rounding mode does not conform to the
7 IEEE International Standard.

8 If the processor supports the alteration of the rounding mode during execution, the subroutine IEEE_-
9 SET_ROUNDING_MODE can be used to alter it. The function IEEE_SUPPORT_ROUNDING can be
10 used to inquire whether this facility is available for a particular mode. The function IEEE_SUPPORT_-
11 IO can be used to inquire whether rounding for base conversion in formatted input/output (9.4.5.14,
12 9.5.2.13, 10.7.2.3.7) is as specified in the IEEE International Standard.

13 In a procedure other than IEEE_SET_ROUNDING_MODE or IEEE_SET_STATUS, the processor shall
14 not change the rounding mode on entry, and on return shall ensure that the rounding mode is the same
15 as it was on entry.

NOTE 14.5

Within a program, all literal constants that have the same form have the same value (4.1.2). Therefore, the value of a literal constant is not affected by the rounding mode.
--

16 14.4 Underflow mode

17 Some processors allow control during program execution of whether underflow produces a denormalized
18 number in conformance with the IEEE International Standard (gradual underflow) or produces zero
19 instead (abrupt underflow). On some processors, floating-point performance is typically better in abrupt
20 underflow mode than in gradual underflow mode.

21 Control over the underflow mode is exercised by invocation of IEEE_SET_UNDERFLOW_MODE. The
22 function IEEE_GET_UNDERFLOW_MODE can be used to inquire which underflow mode is in opera-
23 tion. The function IEEE_SUPPORT_UNDERFLOW_MODE can be used to inquire whether this facility
24 is available. The initial underflow mode is processor dependent. In a procedure other than IEEE_SET_-
25 UNDERFLOW_MODE or IEEE_SET_STATUS, the processor shall not change the underflow mode on
26 entry, and on return shall ensure that the underflow mode is the same as it was on entry.

27 The underflow mode affects only floating-point calculations whose type is that of an X for which IEEE_-
28 SUPPORT_UNDERFLOW_CONTROL returns true.

29 14.5 Halting

30 Some processors allow control during program execution of whether to abort or continue execution after
31 an exception. Such control is exercised by invocation of the subroutine IEEE_SET_HALTING_MODE.
32 Halting is not precise and may occur any time after the exception has occurred. The function IEEE_SUP-
33 PORT_HALTING can be used to inquire whether this facility is available. The initial halting mode is
34 processor dependent. In a procedure other than IEEE_SET_HALTING_MODE or IEEE_SET_STATUS,
35 the processor shall not change the halting mode on entry, and on return shall ensure that the halting
36 mode is the same as it was on entry.

1 14.6 The floating point status

2 The values of all the supported flags for exceptions, rounding mode, underflow mode, and halting are
3 called the floating point status. The floating point status can be saved in a scalar variable of type
4 `TYPE(IEEE_STATUS_TYPE)` with the subroutine `IEEE_GET_STATUS` and restored with the subrou-
5 tine `IEEE_SET_STATUS`. There are no facilities for finding the values of particular flags held within
6 such a variable. Portions of the floating point status can be saved with the subroutines `IEEE_GET_-`
7 `FLAG`, `IEEE_GET_HALTING_MODE`, and `IEEE_GET_ROUNDING_MODE`, and set with the subrou-
8 tines `IEEE_SET_FLAG`, `IEEE_SET_HALTING_MODE`, and `IEEE_SET_ROUNDING_MODE`.

NOTE 14.6

Some processors hold all these flags in a floating point status register that can be saved and restored as a whole much faster than all individual flags can be saved and restored. These procedures are provided to exploit this feature.

NOTE 14.7

The processor is required to ensure that a call to a Fortran procedure does not change the floating point status other than by setting exception flags to signaling.

9 14.7 Exceptional values

10 The IEEE International Standard specifies the following exceptional floating point values.

- 11 • Denormalized values have very small absolute values and lowered precision.
- 12 • Infinite values (+infinity and -infinity) are created by overflow or division by zero.
- 13 • Not-a-Number (NaN) values are undefined values or values created by an invalid operation.

14 In this standard, the term **normal** is used for values that are not in one of these exceptional classes.

15 The functions `IEEE_IS_FINITE`, `IEEE_IS_NAN`, `IEEE_IS_NEGATIVE`, and `IEEE_IS_NORMAL` are pro-
16 vided to test whether a value is finite, NaN, negative, or normal. The function `IEEE_VALUE` is pro-
17 vided to generate an IEEE number of any class, including an infinity or a NaN. The functions `IEEE_-`
18 `SUPPORT_DENORMAL`, `IEEE_SUPPORT_INF`, and `IEEE_SUPPORT_NAN` are provided to determine
19 whether these facilities are available for a particular kind of real.

20 14.8 IEEE arithmetic

21 The function `IEEE_SUPPORT_DATATYPE` can be used to inquire whether IEEE arithmetic is sup-
22 ported for a particular kind of real. Complete conformance with the IEEE International Standard is not
23 required, but the normalized numbers shall be exactly those of an IEEE floating-point format; the oper-
24 ations of addition, subtraction, and multiplication shall conform with at least one of the IEEE rounding
25 modes; the IEEE operation `rem` shall be provided by the function `IEEE_REM`; and the IEEE functions
26 `copysign`, `scalb`, `logb`, `nextafter`, and `unordered` shall be provided by the functions `IEEE_COPY_SIGN`,
27 `IEEE_SCALB`, `IEEE_LOGB`, `IEEE_NEXT_AFTER`, and `IEEE_UNORDERED`, respectively. The in-
28 quiry function `IEEE_SUPPORT_DIVIDE` is provided to inquire whether the processor supports divide
29 with the accuracy specified by the IEEE International Standard. For each of the operations of addi-
30 tion, subtraction, and multiplication, the result shall be as specified in the IEEE International Standard
31 whenever the IEEE result is normalized and the operands are normalized (if floating point) or are valid
and within range (if another type).

1

2 The inquiry function IEEE_SUPPORT_NAN is provided to inquire whether the processor supports
3 IEEE NaNs. Where these are supported, their behavior for unary and binary operations, including
4 those defined by intrinsic functions and by functions in intrinsic modules, shall be consistent with the
5 specifications in the IEEE International Standard.

6 The inquiry function IEEE_SUPPORT_INF is provided to inquire whether the processor supports IEEE
7 infinities. Where these are supported, their behavior for unary and binary operations, including those
8 defined by intrinsic functions and by functions in intrinsic modules, shall be consistent with the speci-
9 fications in the IEEE International Standard.

10 The inquiry function IEEE_SUPPORT_DENORMAL is provided to inquire whether the processor sup-
11 ports IEEE denormals. Where these are supported, their behavior for unary and binary operations,
12 including those defined by intrinsic functions and by functions in intrinsic modules, shall be consistent
13 with the specifications in the IEEE International Standard.

14 The IEEE International Standard specifies a square root function that returns -0.0 for the square root
15 of -0.0 and has certain accuracy requirements. The function IEEE_SUPPORT_SQRT can be used to
16 inquire whether SQRT conforms to the IEEE International Standard for a particular kind of real.

17 The inquiry function IEEE_SUPPORT_STANDARD is provided to inquire whether the processor sup-
18 ports all the IEEE facilities defined in this standard for a particular kind of real.

19 **14.9 Tables of the procedures**

20 For all of the procedures defined in the modules, the arguments shown are the names that shall be used
21 for argument keywords if the keyword form is used for the actual arguments.

22 The procedure classification terms “inquiry function” and “transformational function” are used here
23 with the same meanings as in 13.1.

24 **14.9.1 Inquiry functions**

25 The module IEEE_EXCEPTIONS contains the following inquiry functions.

26	IEEE_SUPPORT_FLAG (FLAG [, X])	Are IEEE exceptions supported?
27	IEEE_SUPPORT_HALTING (FLAG)	Is IEEE halting control supported?

28 The module IEEE_ARITHMETIC contains the following inquiry functions.

29	IEEE_SUPPORT_DATATYPE ([X])	Is IEEE arithmetic supported?
30	IEEE_SUPPORT_DENORMAL ([X])	Are IEEE denormalized numbers supported?
31	IEEE_SUPPORT_DIVIDE ([X])	Is IEEE divide supported?
32	IEEE_SUPPORT_INF ([X])	Is IEEE infinity supported?
33	IEEE_SUPPORT_IO ([X])	Is IEEE formatting supported?
34	IEEE_SUPPORT_NAN ([X])	Are IEEE NaNs supported?
	IEEE_SUPPORT_ROUNDING	Is IEEE rounding supported?
35	(ROUND_VALUE [, X])	
36	IEEE_SUPPORT_SQRT ([X])	Is IEEE square root supported?
37	IEEE_SUPPORT_STANDARD ([X])	Are all IEEE facilities supported?
	IEEE_SUPPORT_UNDERFLOW_-	Is IEEE underflow control supported?
	CONTROL ([X])	

38

39 **14.9.2 Elemental functions**

1 The module IEEE_ARITHMETIC contains the following elemental functions for reals X and Y for which
 2 IEEE_SUPPORT_DATATYPE(X) and IEEE_SUPPORT_DATATYPE(Y) are true.

3	IEEE_CLASS (X)	IEEE class.
4	IEEE_COPY_SIGN (X,Y)	IEEE copysign function.
5	IEEE_IS_FINITE (X)	Determine if value is finite.
6	IEEE_IS_NAN (X)	Determine if value is IEEE Not-a-Number.
7	IEEE_IS_NORMAL (X)	Determine if a value is normal, that is, neither an infinity, a NaN, nor denormalized.
8		
9	IEEE_IS_NEGATIVE (X)	Determine if value is negative.
10	IEEE_LOGB (X)	Unbiased exponent in the IEEE floating point format.
11		
12	IEEE_NEXT_AFTER (X,Y)	Returns the next representable neighbor of X in the direction toward Y.
13		
14	IEEE_REM (X,Y)	The IEEE REM function, that is $X - Y*N$, where N is the integer nearest to the exact value X/Y.
15		
16		
17	IEEE_RINT (X)	Round to an integer value according to the current rounding mode.
18		
19	IEEE_SCALB (X,I)	Returns $X \times 2^I$.
20	IEEE_UNORDERED (X,Y)	IEEE unordered function. True if X or Y is a NaN and false otherwise.
21		
22	IEEE_VALUE (X,CLASS)	Generate an IEEE value.

23 14.9.3 Kind function

24 The module IEEE_ARITHMETIC contains the following transformational function.

25	IEEE_SELECTED_REAL_KIND ([P, R, RADIX])	Kind type parameter value for an IEEE real with given precision, range, and radix.
26		

27 14.9.4 Elemental subroutines

28 The module IEEE_EXCEPTIONS contains the following elemental subroutines.

29	IEEE_GET_FLAG (FLAG,FLAG_VALUE)	Get an exception flag.
30	IEEE_GET_HALTING_MODE (FLAG, HALTING)	Get halting mode for an exception.

31 14.9.5 Nonelemental subroutines

32 The module IEEE_EXCEPTIONS contains the following nonelemental subroutines.

33	IEEE_GET_STATUS (STATUS_VALUE)	Get the current state of the floating point environment.
34		
35	IEEE_SET_FLAG (FLAG,FLAG_VALUE)	Set an exception flag.
36	IEEE_SET_HALTING_MODE (FLAG, HALTING)	Controls continuation or halting on exceptions.
37	IEEE_SET_STATUS (STATUS_VALUE)	Restore the state of the floating point environment.
38		

39 The nonelemental subroutines IEEE_SET_FLAG and IEEE_SET_HALTING_MODE are pure. No other
 40 nonelemental subroutine contained in IEEE_EXCEPTIONS is pure.

1 The module IEEE_ARITHMETIC contains the following nonelemental subroutines.

2	IEEE_GET_ROUNDING_MODE (ROUND_VALUE)	Get the current IEEE rounding mode.
3	IEEE_GET_UNDERFLOW_MODE (GRADUAL)	Get the current underflow mode.
4	IEEE_SET_ROUNDING_MODE (ROUND_VALUE)	Set the current IEEE rounding mode.
5	IEEE_SET_UNDERFLOW_MODE (GRADUAL)	Set the current underflow mode.

6 No nonelemental subroutine contained in IEEE_ARITHMETIC is pure.

7 14.10 Specifications of the procedures

8 In the detailed descriptions below, procedure names are generic and are not specific. All the functions
9 are pure. The dummy arguments of the intrinsic module procedures in 14.9.1, 14.9.2, and 14.9.3 have
10 INTENT(IN). The dummy arguments of the intrinsic module procedures in 14.9.4 and 14.9.5 have
11 INTENT(IN) if the intent is not stated explicitly. In the examples, it is assumed that the processor
12 supports IEEE arithmetic for default real.

NOTE 14.8

It is intended that a processor should not check a condition given in a paragraph labeled “**Restriction**” at compile time, but rather should rely on the programmer writing code such as

```

IF (IEEE_SUPPORT_DATATYPE(X)) THEN
  C = IEEE_CLASS(X)
ELSE
  .
  .
ENDIF

```

to avoid a call being made on a processor for which the condition is violated.

13 For the elemental functions of IEEE_ARITHMETIC, as tabulated in 14.9.2, if X or Y has a value that
14 is an infinity or a NaN, the result shall be consistent with the general rules in 6.1 and 6.2 of the IEEE
15 International Standard. For example, the result for an infinity shall be constructed as the limiting case
16 of the result with a value of arbitrarily large magnitude, if such a limit exists.

17 14.10.1 IEEE_CLASS (X)

18 **Description.** IEEE class function.

19 **Class.** Elemental function.

20 **Argument.** X shall be of type real.

21 **Restriction.** IEEE_CLASS(X) shall not be invoked if IEEE_SUPPORT_DATATYPE(X) has the
22 value false.

23 **Result Characteristics.** TYPE(IEEE_CLASS_TYPE).

24 **Result Value.** The result value shall be IEEE_SIGNALING_NAN or IEEE_QUIET_NAN if

1 IEEE_SUPPORT_NAN(X) has the value true and the value of X is a signaling or quiet NaN,
 2 respectively. The result value shall be IEEE_NEGATIVE_INF or IEEE_POSITIVE_INF if IEEE-
 3 _SUPPORT_INF(X) has the value true and the value of X is negative or positive infinity, respec-
 4 tively. The result value shall be IEEE_NEGATIVE_DENORMAL or IEEE_POSITIVE_DENOR-
 5 MAL if IEEE_SUPPORT_DENORMAL(X) has the value true and the value of X is a negative or
 6 positive denormalized value, respectively. The result value shall be IEEE_NEGATIVE_NORMAL,
 7 IEEE_NEGATIVE_ZERO, IEEE_POSITIVE_ZERO, or IEEE_POSITIVE_NORMAL if value of X
 8 is negative normal, negative zero, positive zero, or positive normal, respectively. Otherwise, the
 9 result value shall be IEEE_OTHER_VALUE.

10 **Example.** IEEE_CLASS(-1.0) has the value IEEE_NEGATIVE_NORMAL.

NOTE 14.9

The result value IEEE_OTHER_VALUE is needed for implementing the module on systems which are basically IEEE, but do not implement all of it. It might be needed, for example, if an unformatted file were written in a program executing with gradual underflow enabled and read with it disabled.

11 14.10.2 IEEE_COPY_SIGN (X, Y)

12 **Description.** IEEE copysign function.

13 **Class.** Elemental function.

14 **Arguments.** The arguments shall be of type real.

15 **Restriction.** IEEE_COPY_SIGN(X,Y) shall not be invoked if IEEE_SUPPORT_DATATYPE(X)
 16 or IEEE_SUPPORT_DATATYPE(Y) has the value false.

17 **Result Characteristics.** Same as X.

18 **Result Value.** The result has the value of X with the sign of Y. This is true even for IEEE
 19 special values, such as a NaN or an infinity (on processors supporting such values).

20 **Example.** The value of IEEE_COPY_SIGN(X,1.0) is ABS(X) even when X is NaN.

21 14.10.3 IEEE_GET_FLAG (FLAG, FLAG_VALUE)

22 **Description.** Get an exception flag.

23 **Class.** Elemental subroutine.

24 **Arguments.**

25 FLAG shall be of type TYPE(IEEE_FLAG_TYPE). It specifies the IEEE flag to be
 obtained.

26 FLAG_VALUE shall be of type default logical. It is an INTENT(OUT) argument. If the
 value of FLAG is IEEE_INVALID, IEEE_OVERFLOW, IEEE_DIVIDE_BY_-
 ZERO, IEEE_UNDERFLOW, or IEEE_INEXACT, the result value is true if
 the corresponding exception flag is signaling and is false otherwise.

27 **Example.** Following CALL IEEE_GET_FLAG(IEEE_OVERFLOW,FLAG_VALUE), FLAG_-
 28 VALUE is true if the IEEE_OVERFLOW flag is signaling and is false if it is quiet.

29 14.10.4 IEEE_GET_HALTING_MODE (FLAG, HALTING)

1 **Description.** Get halting mode for an exception.

2 **Class.** Elemental subroutine.

3 **Arguments.**

4 FLAG shall be of type TYPE(IEEE_FLAG_TYPE). It specifies the IEEE flag. It shall have one of the values IEEE_INVALID, IEEE_OVERFLOW, IEEE_DIVIDE-BY_ZERO, IEEE_UNDERFLOW, or IEEE_INEXACT.

5 HALTING shall be of type default logical. It is of INTENT(OUT). The value is true if the exception specified by FLAG will cause halting. Otherwise, the value is false.

6 **Example.** To store the halting mode for IEEE_OVERFLOW, do a calculation without halting, and restore the halting mode later:

```
8            USE, INTRINSIC :: IEEE_ARITHMETIC
9            LOGICAL HALTING
10            ...
11            CALL IEEE_GET_HALTING_MODE(IEEE_OVERFLOW,HALTING) ! Store halting mode
12            CALL IEEE_SET_HALTING_MODE(IEEE_OVERFLOW,.FALSE.) ! No halting
13            ...! calculation without halting
14            CALL IEEE_SET_HALTING_MODE(IEEE_OVERFLOW,HALTING) ! Restore halting mode
```

15 14.10.5 IEEE_GET_ROUNDING_MODE (ROUND_VALUE)

16 **Description.** Get the current IEEE rounding mode.

17 **Class.** Subroutine.

18 **Argument.** ROUND_VALUE shall be scalar of type TYPE(IEEE_ROUND_TYPE). It is an INTENT(OUT) argument and returns the floating point rounding mode, with value IEEE_NEAREST, IEEE_TO_ZERO, IEEE_UP, or IEEE_DOWN if one of the IEEE modes is in operation and IEEE_OTHER otherwise.

22 **Example.** To store the rounding mode, do a calculation with round to nearest, and restore the rounding mode later:

```
24            USE, INTRINSIC :: IEEE_ARITHMETIC
25            TYPE(IEEE_ROUND_TYPE) ROUND_VALUE
26            ...
27            CALL IEEE_GET_ROUNDING_MODE(ROUND_VALUE) ! Store the rounding mode
28            CALL IEEE_SET_ROUNDING_MODE(IEEE_NEAREST)
29            ... ! calculation with round to nearest
30            CALL IEEE_SET_ROUNDING_MODE(ROUND_VALUE) ! Restore the rounding mode
```

31 14.10.6 IEEE_GET_STATUS (STATUS_VALUE)

32 **Description.** Get the current value of the floating point status (14.6).

33 **Class.** Subroutine.

1 **Argument.** STATUS_VALUE shall be scalar of type TYPE(IEEE_STATUS_TYPE). It is an
2 INTENT(OUT) argument and returns the floating point status.

3 **Example.** To store all the exception flags, do a calculation involving exception handling, and
4 restore them later:

```
5           USE, INTRINSIC :: IEEE_ARITHMETIC
6           TYPE(IEEE_STATUS_TYPE) STATUS_VALUE
7           ...
8           CALL IEEE_GET_STATUS(STATUS_VALUE) ! Get the flags
9           CALL IEEE_SET_FLAG(IEEE_ALL,.FALSE.) ! Set the flags quiet.
10          ... ! calculation involving exception handling
11          CALL IEEE_SET_STATUS(STATUS_VALUE) ! Restore the flags
```

12 14.10.7 IEEE_GET_UNDERFLOW_MODE (GRADUAL)

13 **Description.** Get the current underflow mode (14.4).

14 **Class.** Subroutine.

15 **Argument.** GRADUAL shall be scalar and of type default logical. It is an INTENT(OUT)
16 argument. The value is true if the current underflow mode is gradual underflow, and false if the
17 current underflow mode is abrupt underflow.

18 **Restriction.** IEEE_GET_UNDERFLOW_MODE shall not be invoked unless IEEE_SUPPORT-
19 _UNDERFLOW_CONTROL(X) is true for some X.

20 **Example.** After CALL IEEE_SET_UNDERFLOW_MODE(.FALSE.), a subsequent CALL
21 IEEE_GET_UNDERFLOW_MODE(GRADUAL) will set GRADUAL to false.

22 14.10.8 IEEE_IS_FINITE (X)

23 **Description.** Determine if a value is finite.

24 **Class.** Elemental function.

25 **Argument.** X shall be of type real.

26 **Restriction.** IEEE_IS_FINITE(X) shall not be invoked if IEEE_SUPPORT_DATATYPE(X) has
27 the value false.

28 **Result Characteristics.** Default logical.

29 **Result Value.** The result has the value true if the value of X is finite, that is, IEEE_CLASS(X)
30 has one of the values IEEE_NEGATIVE_NORMAL, IEEE_NEGATIVE_DENORMAL, IEEE-
31 NEGATIVE_ZERO, IEEE_POSITIVE_ZERO, IEEE_POSITIVE_DENORMAL, or IEEE_POSI-
32 TIVE_NORMAL; otherwise, the result has the value false.

33 **Example.** IEEE_IS_FINITE(1.0) has the value true.

34 14.10.9 IEEE_IS_NAN (X)

35 **Description.** Determine if a value is IEEE Not-a-Number.

36 **Class.** Elemental function.

1 **Argument.** X shall be of type real.

2 **Restriction.** IEEE_IS_NAN(X) shall not be invoked if IEEE_SUPPORT_NAN(X) has the value
3 false.

4 **Result Characteristics.** Default logical.

5 **Result Value.** The result has the value true if the value of X is an IEEE NaN; otherwise, it has
6 the value false.

7 **Example.** IEEE_IS_NAN(SQRT(-1.0)) has the value true if IEEE_SUPPORT_SQRT(1.0) has
8 the value true.

9 **14.10.10 IEEE_IS_NEGATIVE (X)**

10 **Description.** Determine if a value is negative.

11 **Class.** Elemental function.

12 **Argument.** X shall be of type real.

13 **Restriction.** IEEE_IS_NEGATIVE(X) shall not be invoked if IEEE_SUPPORT_DATATYPE(X)
14 has the value false.

15 **Result Characteristics.** Default logical.

16 **Result Value.** The result has the value true if IEEE_CLASS(X) has one of the values IEEE_-
17 NEGATIVE_NORMAL, IEEE_NEGATIVE_DENORMAL, IEEE_NEGATIVE_ZERO or IEEE_-
18 NEGATIVE_INF; otherwise, the result has the value false.

19 **Example.** IEEE_IS_NEGATIVE(0.0)) has the value false.

20 **14.10.11 IEEE_IS_NORMAL (X)**

21 **Description.** Determine if a value is normal, that is, neither an infinity, a NaN, nor denormalized.

22 **Class.** Elemental function.

23 **Argument.** X shall be of type real.

24 **Restriction.** IEEE_IS_NORMAL(X) shall not be invoked if IEEE_SUPPORT_DATATYPE(X)
25 has the value false.

26 **Result Characteristics.** Default logical.

27 **Result Value.** The result has the value true if IEEE_CLASS(X) has one of the values IEEE_NEG-
28 ATIVE_NORMAL, IEEE_NEGATIVE_ZERO, IEEE_POSITIVE_ZERO or IEEE_POSITIVE_-
29 NORMAL; otherwise, the result has the value false.

30 **Example.** IEEE_IS_NORMAL(SQRT(-1.0)) has the value false if IEEE_SUPPORT_SQRT(1.0)
31 has the value true.

32 **14.10.12 IEEE_LOGB (X)**

33 **Description.** Unbiased exponent in the IEEE floating point format.

34 **Class.** Elemental function.

35 **Argument.** X shall be of type real.

1 **Restriction.** IEEE_LOGB(X) shall not be invoked if IEEE_SUPPORT_DATATYPE(X) has the
2 value false.

3 **Result Characteristics.** Same as X.

4 **Result Value.**

5 *Case (i):* If the value of X is neither zero, infinity, nor NaN, the result has the value of the
6 unbiased exponent of X. Note: this value is equal to EXPONENT(X)-1.

7 *Case (ii):* If X==0, the result is -infinity if IEEE_SUPPORT_INF(X) is true and -HUGE(X)
8 otherwise; IEEE_DIVIDE_BY_ZERO signals.

9 **Example.** IEEE_LOGB(-1.1) has the value 0.0.

10 14.10.13 IEEE_NEXT_AFTER (X, Y)

11 **Description.** Returns the next representable neighbor of X in the direction toward Y.

12 **Class.** Elemental function.

13 **Arguments.** The arguments shall be of type real.

14 **Restriction.** IEEE_NEXT_AFTER(X,Y) shall not be invoked if IEEE_SUPPORT_DATA-
15 TYPE(X) or IEEE_SUPPORT_DATATYPE(Y) has the value false.

16 **Result Characteristics.** Same as X.

17 **Result Value.**

18 *Case (i):* If X == Y, the result is X and no exception is signaled.

19 *Case (ii):* If X /= Y, the result has the value of the next representable neighbor of X in the
20 direction of Y. The neighbors of zero (of either sign) are both nonzero. IEEE-
21 OVERFLOW is signaled when X is finite but IEEE_NEXT_AFTER(X,Y) is infi-
22 nite; IEEE_UNDERFLOW is signaled when IEEE_NEXT_AFTER(X,Y) is denor-
23 malized; in both cases, IEEE_INEXACT signals.

24 **Example.** The value of IEEE_NEXT_AFTER(1.0,2.0) is 1.0+EPSILON(X).

25 14.10.14 IEEE_REM (X, Y)

26 **Description.** IEEE REM function.

27 **Class.** Elemental function.

28 **Arguments.** The arguments shall be of type real.

29 **Restriction.** IEEE_REM(X,Y) shall not be invoked if IEEE_SUPPORT_DATATYPE(X) or
30 IEEE_SUPPORT_DATATYPE(Y) has the value false.

31 **Result Characteristics.** Real with the kind type parameter of whichever argument has the
32 greater precision.

33 **Result Value.** The result value, regardless of the rounding mode, shall be exactly $X - Y * N$,
34 where N is the integer nearest to the exact value X/Y ; whenever $|N - X/Y| = 1/2$, N shall be
35 even. If the result value is zero, the sign shall be that of X.

36 **Examples.** The value of IEEE_REM(4.0,3.0) is 1.0, the value of IEEE_REM(3.0,2.0) is -1.0, and
37 the value of IEEE_REM(5.0,2.0) is 1.0.

1 14.10.15 IEEE_RINT (X)

2 **Description.** Round to an integer value according to the current rounding mode.

3 **Class.** Elemental function.

4 **Argument.** X shall be of type real.

5 **Restriction.** IEEE_RINT(X) shall not be invoked if IEEE_SUPPORT_DATATYPE(X) has the
6 value false.

7 **Result Characteristics.** Same as X.

8 **Result Value.** The value of the result is the value of X rounded to an integer according to the
9 current rounding mode. If the result has the value zero, the sign is that of X.

10 **Examples.** If the current rounding mode is round to nearest, the value of IEEE_RINT(1.1) is
11 1.0. If the current rounding mode is round up, the value of IEEE_RINT(1.1) is 2.0.

12 14.10.16 IEEE_SCALB (X, I)

13 **Description.** Returns $X \times 2^I$.

14 **Class.** Elemental function.

15 **Arguments.**

16 X shall be of type real.

17 I shall be of type integer.

18 **Restriction.** IEEE_SCALB(X) shall not be invoked if IEEE_SUPPORT_DATATYPE(X) has
19 the value false.

20 **Result Characteristics.** Same as X.

21 **Result Value.**

22 *Case (i):* If $X \times 2^I$ is representable as a normal number, the result has this value.

23 *Case (ii):* If X is finite and $X \times 2^I$ is too large, the IEEE_OVERFLOW exception shall occur.
24 If IEEE_SUPPORT_INF(X) is true, the result value is infinity with the sign of X;
25 otherwise, the result value is SIGN(HUGE(X),X).

26 *Case (iii):* If $X \times 2^I$ is too small and there is loss of accuracy, the IEEE_UNDERFLOW
27 exception shall occur. The result is the representable number having a magnitude
28 nearest to $|2^I|$ and the same sign as X.

29 *Case (iv):* If X is infinite, the result is the same as X; no exception signals.

30 **Example.** The value of IEEE_SCALB(1.0,2) is 4.0.

31 14.10.17 IEEE_SELECTED_REAL_KIND ([P, R, RADIX])

32 **Description.** Returns a value of the kind type parameter of an IEEE real data type with decimal
33 precision of at least P digits, a decimal exponent range of at least R, and a radix of RADIX. For
34 data objects of such a type, IEEE_SUPPORT_DATATYPE(X) has the value true.

35 **Class.** Transformational function.

36 **Arguments.** At least one argument shall be present.

1 P (optional) shall be scalar and of type integer.

2 R (optional) shall be scalar and of type integer.

3 RADIX (optional) shall be scalar and of type integer.

4 **Result Characteristics.** Default integer scalar.

5 **Result Value.** If P or R is absent, the result value is the same as if it were present with the value
6 zero. If RADIX is absent, there is no requirement on the radix of the selected kind. The result has
7 a value equal to a value of the kind type parameter of an IEEE real type with decimal precision, as
8 returned by the function PRECISION, of at least P digits, a decimal exponent range, as returned
9 by the function RANGE, of at least R, and a radix, as returned by the function RADIX, of
10 RADIX, if such a kind type parameter is available on the processor.

11 Otherwise, the result is -1 if the processor supports an IEEE real type with radix RADIX and exponent
12 range of at least R but not with precision of at least P, -2 if the processor supports an IEEE real
13 type with radix RADIX and precision of at least P but not with exponent range of at least R, -3 if
14 the processor supports an IEEE real type with radix RADIX but with neither precision of at least P
15 nor exponent range of at least R, -4 if the processor supports an IEEE real type with radix RADIX
16 and either precision of at least P or exponent range of at least R but not both together, and -5 if the
17 processor supports no IEEE real type with radix RADIX.

18 If more than one kind type parameter value meets the criteria, the value returned is the one with the
19 smallest decimal precision, unless there are several such values, in which case the smallest of these kind
20 values is returned.

21 **Example.** IEEE_SELECTED_REAL_KIND(6,30) has the value KIND(0.0) on a machine that
22 supports IEEE single precision arithmetic for its default real approximation method.

23 14.10.18 IEEE_SET_FLAG (FLAG, FLAG_VALUE)

24 **Description.** Assign a value to an exception flag.

25 **Class.** Pure subroutine.

26 **Arguments.**

FLAG shall be a scalar or array of type TYPE(IEEE_FLAG_TYPE). If a value
of FLAG is IEEE_INVALID, IEEE_OVERFLOW, IEEE_DIVIDE_BY_ZERO,
IEEE_UNDERFLOW, or IEEE_INEXACT, the corresponding exception flag is
27 assigned a value. No two elements of FLAG shall have the same value.

FLAG_VALUE shall be a scalar or array of type default logical. It shall be conformable with
FLAG. If an element has the value true, the corresponding flag is set to be
28 signaling; otherwise, the flag is set to be quiet.

29 **Example.** CALL IEEE_SET_FLAG(IEEE_OVERFLOW, .TRUE.) sets the IEEE_OVERFLOW
30 flag to be signaling.

31 14.10.19 IEEE_SET_HALTING_MODE (FLAG, HALTING)

32 **Description.** Controls continuation or halting after an exception.

33 **Class.** Pure subroutine.

34 **Arguments.**

1 FLAG shall be a scalar or array of type TYPE(IEEE_FLAG_TYPE). It shall have only the values IEEE_INVALID, IEEE_OVERFLOW, IEEE_DIVIDE_BY_ZERO, IEEE_UNDERFLOW, or IEEE_INEXACT. No two elements of FLAG shall have the same value.

2 HALTING shall be a scalar or array of type default logical. It shall be conformable with FLAG. If an element has the value is true, the corresponding exception specified by FLAG will cause halting. Otherwise, execution will continue after this exception.

3 **Restriction.** IEEE_SET_HALTING_MODE(FLAG,HALTING) shall not be invoked if IEEE_SUPPORT_HALTING(FLAG) has the value false.

4 **Example.** CALL IEEE_SET_HALTING_MODE(IEEE_DIVIDE_BY_ZERO,.TRUE.) causes halting after a divide_by_zero exception.

NOTE 14.10

The initial halting mode is processor dependent. Halting is not precise and may occur some time after the exception has occurred.

7 14.10.20 IEEE_SET_ROUNDING_MODE (ROUND_VALUE)

8 **Description.** Set the current IEEE rounding mode.

9 **Class.** Subroutine.

10 **Argument.** ROUND_VALUE shall be scalar and of type TYPE(IEEE_ROUND_TYPE). It specifies the mode to be set.

11 **Restriction.** IEEE_SET_ROUNDING_MODE(ROUND_VALUE) shall not be invoked unless IEEE_SUPPORT_ROUNDING(ROUND_VALUE,X) is true for some X such that IEEE_SUPPORT_DATATYPE(X) is true.

12 **Example.** To store the rounding mode, do a calculation with round to nearest, and restore the rounding mode later:

```
17 USE, INTRINSIC :: IEEE_ARITHMETIC
18 TYPE(IEEE_ROUND_TYPE) ROUND_VALUE
19 ...
20 CALL IEEE_GET_ROUNDING_MODE(ROUND_VALUE) ! Store the rounding mode
21 CALL IEEE_SET_ROUNDING_MODE(IEEE_NEAREST)
22 : ! calculation with round to nearest
23 CALL IEEE_SET_ROUNDING_MODE(ROUND_VALUE) ! Restore the rounding mode
```

24 14.10.21 IEEE_SET_STATUS (STATUS_VALUE)

25 **Description.** Restore the value of the floating point status (14.6).

26 **Class.** Subroutine.

27 **Argument.** STATUS_VALUE shall be scalar and of type TYPE(IEEE_STATUS_TYPE). Its value shall have been set in a previous invocation of IEEE_GET_STATUS.

28 **Example.** To store all the exceptions flags, do a calculation involving exception handling, and restore them later:

```

1      USE, INTRINSIC :: IEEE_EXCEPTIONS
2      TYPE(IEEE_STATUS_TYPE) STATUS_VALUE
3      ...
4      CALL IEEE_GET_STATUS(STATUS_VALUE) ! Store the flags
5      CALL IEEE_SET_FLAGS(IEEE_ALL,.FALSE.) ! Set them quiet
6      ... ! calculation involving exception handling
7      CALL IEEE_SET_STATUS(STATUS_VALUE) ! Restore the flags

```

NOTE 14.11

On some processors this may be a very time consuming process.

8 14.10.22 IEEE_SET_UNDERFLOW_MODE (GRADUAL)

9 **Description.** Set the current underflow mode.

10 **Class.** Subroutine.

11 **Argument.** GRADUAL shall be scalar and of type default logical. If it is true, the current
 12 underflow mode is set to gradual underflow. If it is false, the current underflow mode is set to
 13 abrupt underflow.

14 **Restriction.** IEEE_SET_UNDERFLOW_MODE shall not be invoked unless IEEE_SUPPORT_
 15 UNDERFLOW_CONTROL(X) is true for some X.

16 **Example.** To perform some calculations with abrupt underflow and then restore the previous
 17 mode:

```

18 USE,INTRINSIC :: IEEE_ARITHMETIC
19 LOGICAL SAVE_UNDERFLOW_MODE
20 ...
21 CALL IEEE_GET_UNDERFLOW_MODE(SAVE_UNDERFLOW_MODE)
22 CALL IEEE_SET_UNDERFLOW_MODE(GRADUAL=.FALSE.)
23 ... ! Perform some calculations with abrupt underflow
24 CALL IEEE_SET_UNDERFLOW_MODE(SAVE_UNDERFLOW_MODE)

```

25 14.10.23 IEEE_SUPPORT_DATATYPE () or IEEE_SUPPORT_DATATYPE (X)

26 **Description.** Inquire whether the processor supports IEEE arithmetic.

27 **Class.** Inquiry function.

28 **Argument.** X shall be of type real. It may be a scalar or an array.

29 **Result Characteristics.** Default logical scalar.

30 **Result Value.** The result has the value true if the processor supports IEEE arithmetic for all
 31 reals (X absent) or for real variables of the same kind type parameter as X; otherwise, it has the
 32 value false. Here, support is as defined in the first paragraph of 14.8.

33 **Example.** If default real type conforms to the IEEE International Standard except that underflow
 34 values flush to zero instead of being denormal, IEEE_SUPPORT_DATATYPE(1.0) has the value
 35 true.

1 14.10.24 IEEE_SUPPORT_DENORMAL () or IEEE_SUPPORT_DENORMAL (X)

2 **Description.** Inquire whether the processor supports IEEE denormalized numbers.

3 **Class.** Inquiry function.

4 **Argument.** X shall be of type real. It may be a scalar or an array.

5 **Result Characteristics.** Default logical scalar.

6 **Result Value.**

7 *Case (i):* IEEE_SUPPORT_DENORMAL(X) has the value true if IEEE_SUPPORT_
8 DATATYPE(X) has the value true and the processor supports arithmetic oper-
9 ations and assignments with denormalized numbers (biased exponent $e = 0$ and
10 fraction $f \neq 0$, see subclause 3.2 of the IEEE International Standard) for real
11 variables of the same kind type parameter as X; otherwise, it has the value false.

12 *Case (ii):* IEEE_SUPPORT_DENORMAL() has the value true if and only if IEEE_SUP-
13 PORT_DENORMAL(X) has the value true for all real X.

14 **Example.** IEEE_SUPPORT_DENORMAL(X) has the value true if the processor supports de-
15 normalized numbers for X.

NOTE 14.12

The denormalized numbers are not included in the 13.4 model for real numbers; they satisfy the inequality $ABS(X) < TINY(X)$. They usually occur as a result of an arithmetic operation whose exact result is less than $TINY(X)$. Such an operation causes IEEE_UNDERFLOW to signal unless the result is exact. IEEE_SUPPORT_DENORMAL(X) is false if the processor never returns a denormalized number as the result of an arithmetic operation.

16 14.10.25 IEEE_SUPPORT_DIVIDE () or IEEE_SUPPORT_DIVIDE (X)

17 **Description.** Inquire whether the processor supports divide with the accuracy specified by the
18 IEEE International Standard.

19 **Class.** Inquiry function.

20 **Argument.** X shall be of type real. It may be a scalar or an array.

21 **Result Characteristics.** Default logical scalar.

22 **Result Value.**

23 *Case (i):* IEEE_SUPPORT_DIVIDE(X) has the value true if the processor supports divide
24 with the accuracy specified by the IEEE International Standard for real variables
25 of the same kind type parameter as X; otherwise, it has the value false.

26 *Case (ii):* IEEE_SUPPORT_DIVIDE() has the value true if and only if IEEE_SUPPORT_
27 DIVIDE(X) has the value true for all real X.

28 **Example.** IEEE_SUPPORT_DIVIDE(X) has the value true if the processor supports IEEE divide
29 for X.

30 14.10.26 IEEE_SUPPORT_FLAG (FLAG) or IEEE_SUPPORT_FLAG (FLAG, X)

31 **Description.** Inquire whether the processor supports an exception.

32 **Class.** Inquiry function.

1 **Arguments.**

2 FLAG shall be scalar and of type TYPE(IEEE_FLAG_TYPE). Its value shall be one of
 3 IEEE_INVALID, IEEE_OVERFLOW, IEEE_DIVIDE_BY_ZERO, IEEE_UN-
 4 DERFLOW, or IEEE_INEXACT.

5 X shall be of type real. It may be a scalar or an array.

6 **Result Characteristics.** Default logical scalar.

7 **Result Value.**

8 *Case (i):* IEEE_SUPPORT_FLAG(FLAGS, X) has the value true if the processor supports de-
 9 tection of the specified exception for real variables of the same kind type parameter
 10 as X; otherwise, it has the value false.

11 *Case (ii):* IEEE_SUPPORT_FLAG(FLAGS) has the value true if and only if IEEE_SUPPORT-
 12 _FLAG(FLAGS, X) has the value true for all real X.

13 **Example.** IEEE_SUPPORT_FLAG(IEEE_INEXACT) has the value true if the processor sup-
 14 ports the inexact exception.

15 **14.10.27 IEEE_SUPPORT_HALTING (FLAG)**

16 **Description.** Inquire whether the processor supports the ability to control during program
 17 execution whether to abort or continue execution after an exception.

18 **Class.** Inquiry function.

19 **Argument.** FLAG shall be scalar and of type TYPE(IEEE_FLAG_TYPE). Its value shall be
 20 one of IEEE_INVALID, IEEE_OVERFLOW, IEEE_DIVIDE_BY_ZERO, IEEE_UNDERFLOW,
 21 or IEEE_INEXACT.

22 **Result Characteristics.** Default logical scalar.

23 **Result Value.** The result has the value true if the processor supports the ability to control
 24 during program execution whether to abort or continue execution after the exception specified
 25 by FLAG; otherwise, it has the value false. Support includes the ability to change the mode by
 26 CALL IEEE_SET_HALTING(FLAGS).

27 **Example.** IEEE_SUPPORT_HALTING(IEEE_OVERFLOW) has the value true if the processor
 28 supports control of halting after an overflow.

29 **14.10.28 IEEE_SUPPORT_INF () or IEEE_SUPPORT_INF (X)**

30 **Description.** Inquire whether the processor supports the IEEE infinity facility.

31 **Class.** Inquiry function.

32 **Argument.** X shall be of type real. It may be a scalar or an array.

33 **Result Characteristics.** Default logical scalar.

34 **Result Value.**

35 *Case (i):* IEEE_SUPPORT_INF(X) has the value true if the processor supports IEEE infini-
 36 ties (positive and negative) for real variables of the same kind type parameter as
 37 X; otherwise, it has the value false.

38 *Case (ii):* IEEE_SUPPORT_INF() has the value true if and only if IEEE_SUPPORT_INF(X)
 39 has the value true for all real X.

1 **Example.** IEEE_SUPPORT_INF(X) has the value true if the processor supports IEEE infinities
2 for X.

3 **14.10.29 IEEE_SUPPORT_IO () or IEEE_SUPPORT_IO (X)**

4 **Description.** Inquire whether the processor supports IEEE base conversion rounding during
5 formatted input/output (9.4.5.14, 9.5.2.13, 10.7.2.3.7).

6 **Class.** Inquiry function.

7 **Argument.** X shall be of type real. It may be a scalar or an array.

8 **Result Characteristics.** Default logical scalar.

9 **Result Value.**

10 *Case (i):* IEEE_SUPPORT_IO(X) has the value true if the processor supports IEEE base con-
11 version during formatted input/output (9.4.5.14, 9.5.2.13, 10.7.2.3.7) as described
12 in the IEEE International Standard for the modes UP, DOWN, ZERO, and NEAR-
13 EST for real variables of the same kind type parameter as X; otherwise, it has the
14 value false.

15 *Case (ii):* IEEE_SUPPORT_IO() has the value true if and only if IEEE_SUPPORT_IO(X)
16 has the value true for all real X.

17 **Example.** IEEE_SUPPORT_IO(X) has the value true if the processor supports IEEE base con-
18 version for X.

19 **14.10.30 IEEE_SUPPORT_NAN () or IEEE_SUPPORT_NAN (X)**

20 **Description.** Inquire whether the processor supports the IEEE Not-a-Number facility.

21 **Class.** Inquiry function.

22 **Argument.** X shall be of type real. It may be a scalar or an array.

23 **Result Characteristics.** Default logical scalar.

24 **Result Value.**

25 *Case (i):* IEEE_SUPPORT_NAN(X) has the value true if the processor supports IEEE NaNs
26 for real variables of the same kind type parameter as X; otherwise, it has the value
27 false.

28 *Case (ii):* IEEE_SUPPORT_NAN() has the value true if and only if IEEE_SUPPORT_-
29 NAN(X) has the value true for all real X.

30 **Example.** IEEE_SUPPORT_NAN(X) has the value true if the processor supports IEEE NaNs
31 for X.

32 **14.10.31 IEEE_SUPPORT_ROUNDING (ROUND_VALUE) or** 33 **IEEE_SUPPORT_ROUNDING (ROUND_VALUE, X)**

34 **Description.** Inquire whether the processor supports a particular IEEE rounding mode.

35 **Class.** Inquiry function.

36 **Arguments.**

ROUND_VALUE shall be of type TYPE(IEEE_ROUND_TYPE).

1

2 X shall be of type real. It may be a scalar or an array.

3 **Result Characteristics.** Default logical scalar.4 **Result Value.**5 *Case (i):* IEEE_SUPPORT_ROUNDING(ROUND_VALUE, X) has the value true if the pro-
6 cessor supports the rounding mode defined by ROUND_VALUE for real variables
7 of the same kind type parameter as X; otherwise, it has the value false. Sup-
8 port includes the ability to change the mode by CALL IEEE_SET_ROUNDING_
9 MODE(ROUND_VALUE).10 *Case (ii):* IEEE_SUPPORT_ROUNDING(ROUND_VALUE) has the value true if and only if
11 IEEE_SUPPORT_ROUNDING(ROUND_VALUE, X) has the value true for all real
12 X.13 **Example.** IEEE_SUPPORT_ROUNDING(IEEE_TO_ZERO) has the value true if the processor
14 supports rounding to zero for all reals.15 **14.10.32 IEEE_SUPPORT_SQRT () or IEEE_SUPPORT_SQRT (X)**16 **Description.** Inquire whether the intrinsic function SQRT conforms to the IEEE International
17 Standard.18 **Class.** Inquiry function.19 **Argument.** X shall be of type real. It may be a scalar or an array.20 **Result Characteristics.** Default logical scalar.21 **Result Value.**22 *Case (i):* IEEE_SUPPORT_SQRT(X) has the value true if the intrinsic function SQRT con-
23 forms to the IEEE International Standard for real variables of the same kind type
24 parameter as X; otherwise, it has the value false.25 *Case (ii):* IEEE_SUPPORT_SQRT() has the value true if and only if IEEE_SUPPORT_
26 SQRT(X) has the value true for all real X.27 **Example.** If IEEE_SUPPORT_SQRT(1.0) has the value true, SQRT(-0.0) will have the value
28 -0.0.29 **14.10.33 IEEE_SUPPORT_STANDARD () or IEEE_SUPPORT_STANDARD (X)**30 **Description.** Inquire whether the processor supports all the IEEE facilities defined in this
31 standard.32 **Class.** Inquiry function.33 **Argument.** X shall be of type real. It may be a scalar or an array.34 **Result Characteristics.** Default logical scalar.35 **Result Value.**36 *Case (i):* IEEE_SUPPORT_STANDARD(X) has the value true if the results of all
37 the functions IEEE_SUPPORT_DATATYPE(X), IEEE_SUPPORT_DENOR-
38 MAL(X), IEEE_SUPPORT_DIVIDE(X), IEEE_SUPPORT_FLAG(FLAG,X)
39 for valid FLAG, IEEE_SUPPORT_HALTING(FLAG) for valid FLAG, IEEE-

1 `__SUPPORT__INF(X)`, `IEEE_SUPPORT_NAN(X)`, `IEEE_SUPPORT_ROUND-`
 2 `ING (ROUND_VALUE,X)` for valid `ROUND_VALUE`, and `IEEE_SUPPORT__`
 3 `SQRT (X)` are all true; otherwise, the result has the value false.

4 *Case (ii):* `IEEE_SUPPORT_STANDARD()` has the value true if and only if `IEEE_SUPPORT-`
 5 `STANDARD(X)` has the value true for all real `X`.

6 **Example.** `IEEE_SUPPORT_STANDARD()` has the value false if the processor supports both
 7 IEEE and non-IEEE kinds of reals.

8 **14.10.34 IEEE_SUPPORT_UNDERFLOW_CONTROL() or** 9 **IEEE_SUPPORT_UNDERFLOW_CONTROL(X)**

10 **Description.** Inquire whether the procedure supports the ability to control the underflow mode
 11 during program execution.

12 **Class.** Inquiry function.

13 **Argument.** `X` shall be of type real. It may be a scalar or an array.

14 **Result Characteristics.** Default logical scalar.

15 **Result Value.**

16 *Case (i):* `IEEE_SUPPORT_UNDERFLOW_CONTROL(X)` has the value true if the proces-
 17 sor supports control of the underflow mode for floating-point calculations with the
 18 same type as `X`, and false otherwise.

19 *Case (ii):* `IEEE_SUPPORT_UNDERFLOW_CONTROL()` has the value true if the processor
 20 supports control of the underflow mode for all floating-point calculations, and false
 21 otherwise.

22 **Example.** `IEEE_SUPPORT_UNDERFLOW_CONTROL(2.5)` has the value true if the processor
 23 supports underflow mode control for calculations of type default real.

24 **14.10.35 IEEE_UNORDERED (X, Y)**

25 **Description.** IEEE unordered function. True if `X` or `Y` is a NaN, and false otherwise.

26 **Class.** Elemental function.

27 **Arguments.** The arguments shall be of type real.

28 **Restriction.** `IEEE_UNORDERED(X,Y)` shall not be invoked if `IEEE_SUPPORT_DATA-`
 29 `TYPE(X)` or `IEEE_SUPPORT_DATATYPE(Y)` has the value false.

30 **Result Characteristics.** Default logical.

31 **Result Value.** The result has the value true if `X` or `Y` is a NaN or both are NaNs; otherwise, it
 32 has the value false.

33 **Example.** `IEEE_UNORDERED(0.0,SQRT(-1.0))` has the value true if `IEEE_SUPPORT_-`
 34 `SQRT(1.0)` has the value true.

35 **14.10.36 IEEE_VALUE (X, CLASS)**

36 **Description.** Generate an IEEE value.

37 **Class.** Elemental function.

1 **Arguments.**

2 X shall be of type real.

3 CLASS shall be of type TYPE(IEEE_CLASS_TYPE). The value is permitted to be:
 IEEE_SIGNALING_NAN or IEEE_QUIET_NAN if IEEE_SUPPORT_NAN(X)
 has the value true, IEEE_NEGATIVE_INF or IEEE_POSITIVE_INF if IEEE_
 SUPPORT_INF(X) has the value true, IEEE_NEGATIVE_DENORMAL or
 IEEE_POSITIVE_DENORMAL if IEEE_SUPPORT_DENORMAL(X) has the
 value true, IEEE_NEGATIVE_NORMAL, IEEE_NEGATIVE_ZERO, IEEE_
 POSITIVE_ZERO or IEEE_POSITIVE_NORMAL.

4 **Restriction.** IEEE_VALUE(X,CLASS) shall not be invoked if IEEE_SUPPORT_DATATYPE(X)
 5 has the value false.

6 **Result Characteristics.** Same as X.

7 **Result Value.** The result value is an IEEE value as specified by CLASS. Although in most cases
 8 the value is processor dependent, the value shall not vary between invocations for any particular
 9 X kind type parameter and CLASS value.

10 **Example.** IEEE_VALUE(1.0,IEEE_NEGATIVE_INF) has the value -infinity.11 **14.11 Examples****NOTE 14.13**

```

MODULE DOT
! Module for dot product of two real arrays of rank 1.
! The caller needs to ensure that exceptions do not cause halting.
  USE, INTRINSIC :: IEEE_EXCEPTIONS
  LOGICAL :: MATRIX_ERROR = .FALSE.
  INTERFACE OPERATOR(.dot.)
    MODULE PROCEDURE MULT
  END INTERFACE
CONTAINS
  REAL FUNCTION MULT(A,B)
    REAL, INTENT(IN) :: A(:),B(:)
    INTEGER I
    LOGICAL OVERFLOW
    IF (SIZE(A)/=SIZE(B)) THEN
      MATRIX_ERROR = .TRUE.
      RETURN
    END IF
! The processor ensures that IEEE_OVERFLOW is quiet
    MULT = 0.0
    DO I = 1, SIZE(A)
      MULT = MULT + A(I)*B(I)
    END DO
    CALL IEEE_GET_FLAG(IEEE_OVERFLOW,OVERFLOW)

```


NOTE 14.13 (cont.)

```

      IF (OVERFLOW) MATRIX_ERROR = .TRUE.
    END FUNCTION MULT
END MODULE DOT

```

This module provides the dot product of two real arrays of rank 1. If the sizes of the arrays are different, an immediate return occurs with `MATRIX_ERROR` true. If overflow occurs during the actual calculation, the `IEEE_OVERFLOW` flag will signal and `MATRIX_ERROR` will be true.

NOTE 14.14

```

USE, INTRINSIC :: IEEE_EXCEPTIONS
USE, INTRINSIC :: IEEE_FEATURES, ONLY: IEEE_INVALID_FLAG
! The other exceptions of IEEE_USUAL (IEEE_OVERFLOW and
! IEEE_DIVIDE_BY_ZERO) are always available with IEEE_EXCEPTIONS
TYPE(IEEE_STATUS_TYPE) STATUS_VALUE
LOGICAL, DIMENSION(3) :: FLAG_VALUE
...
CALL IEEE_GET_STATUS(STATUS_VALUE)
CALL IEEE_SET_HALTING_MODE(IEEE_USUAL,.FALSE.) ! Needed in case the
!           default on the processor is to halt on exceptions
CALL IEEE_SET_FLAG(IEEE_USUAL,.FALSE.)
! First try the "fast" algorithm for inverting a matrix:
MATRIX1 = FAST_INV(MATRIX) ! This shall not alter MATRIX.
CALL IEEE_GET_FLAG(IEEE_USUAL,FLAG_VALUE)
IF (ANY(FLAG_VALUE)) THEN
! "Fast" algorithm failed; try "slow" one:
  CALL IEEE_SET_FLAG(IEEE_USUAL,.FALSE.)
  MATRIX1 = SLOW_INV(MATRIX)
  CALL IEEE_GET_FLAG(IEEE_USUAL,FLAG_VALUE)
  IF (ANY(FLAG_VALUE)) THEN
    WRITE (*, *) 'Cannot invert matrix'
    STOP
  END IF
END IF
CALL IEEE_SET_STATUS(STATUS_VALUE)

```

In this example, the function `FAST_INV` may cause a condition to signal. If it does, another try is made with `SLOW_INV`. If this still fails, a message is printed and the program stops. Note, also, that it is important to set the flags quiet before the second try. The state of all the flags is stored and restored.

NOTE 14.15

```

USE, INTRINSIC :: IEEE_EXCEPTIONS
LOGICAL FLAG_VALUE

```

NOTE 14.15 (cont.)

```

...
CALL IEEE_SET_HALTING_MODE(IEEE_OVERFLOW,.FALSE.)
! First try a fast algorithm for inverting a matrix.
CALL IEEE_SET_FLAG(IEEE_OVERFLOW,.FALSE.)
DO K = 1, N
...
    CALL IEEE_GET_FLAG(IEEE_OVERFLOW,FLAG_VALUE)
    IF (FLAG_VALUE) EXIT
END DO
IF (FLAG_VALUE) THEN
! Alternative code which knows that K-1 steps have executed normally.
...
END IF

```

Here the code for matrix inversion is in line and the transfer is made more precise by adding extra tests of the flag.

NOTE 14.16

```

REAL FUNCTION HYPOT(X, Y)
! In rare circumstances this may lead to the signaling of IEEE_OVERFLOW
! The caller needs to ensure that exceptions do not cause halting.
    USE, INTRINSIC :: IEEE_ARITHMETIC
    USE, INTRINSIC :: IEEE_FEATURES, ONLY: IEEE_UNDERFLOW_FLAG
! IEEE_OVERFLOW is always available with IEEE_ARITHMETIC
    REAL X, Y
    REAL SCALED_X, SCALED_Y, SCALED_RESULT
    LOGICAL, DIMENSION(2) :: FLAGS
    TYPE(IEEE_FLAG_TYPE), PARAMETER, DIMENSION(2) :: &
        OUT_OF_RANGE = (/ IEEE_OVERFLOW, IEEE_UNDERFLOW /)
    INTRINSIC SQRT, ABS, EXPONENT, MAX, DIGITS, SCALE
! The processor clears the flags on entry
! Try a fast algorithm first
    HYPOT = SQRT( X**2 + Y**2 )
    CALL IEEE_GET_FLAG(OUT_OF_RANGE,FLAGS)
    IF ( ANY(FLAGS) ) THEN
        CALL IEEE_SET_FLAG(OUT_OF_RANGE,.FALSE.)
        IF ( X==0.0 .OR. Y==0.0 ) THEN
            HYPOT = ABS(X) + ABS(Y)
        ELSE IF ( 2*ABS(EXPONENT(X)-EXPONENT(Y)) > DIGITS(X)+1 ) THEN
            HYPOT = MAX( ABS(X), ABS(Y) )! one of X and Y can be ignored
        ELSE ! scale so that ABS(X) is near 1
            SCALED_X = SCALE( X, -EXPONENT(X) )
            SCALED_Y = SCALE( Y, -EXPONENT(X) )

```

NOTE 14.16 (cont.)

```
        SCALED_RESULT = SQRT( SCALED_X**2 + SCALED_Y**2 )
        HYPOT = SCALE( SCALED_RESULT, EXPONENT(X) ) ! may cause overflow
    END IF
END IF
! The processor resets any flag that was signaling on entry
END FUNCTION HYPOT
```

An attempt is made to evaluate this function directly in the fastest possible way. This will work almost every time, but if an exception occurs during this fast computation, a safe but slower way evaluates the function. This slower evaluation might involve scaling and unscaling, and in (very rare) extreme cases this unscaling can cause overflow (after all, the true result might overflow if X and Y are both near the overflow limit). If the IEEE_OVERFLOW or IEEE_UNDERFLOW flag is signaling on entry, it is reset on return by the processor, so that earlier exceptions are not lost.

1 15 Interoperability with C

2 15.1 General

3 Fortran provides a means of referencing procedures that are defined by means of the C programming
4 language or procedures that can be described by C prototypes as defined in 6.7.5.3 of the C International
5 Standard, even if they are not actually defined by means of C. Conversely, there is a means of specifying
6 that a procedure defined by a Fortran subprogram can be referenced from a function defined by means
7 of C. In addition, there is a means of declaring global variables that are associated with C variables that
8 have external linkage as defined in 6.2.2 of the C International Standard.

9 The ISO_C_BINDING module provides access to named constants that represent kind type parameters
10 of data representations compatible with C types. Fortran also provides facilities for defining derived
11 types (4.5) and enumerations (4.6) that correspond to C types.

12 15.2 The ISO_C_BINDING intrinsic module

13 15.2.1 Summary of contents

14 The processor shall provide the intrinsic module ISO_C_BINDING. This module shall make accessible
15 the following entities: named constants with the names listed in the second column of Table 15.2 and
16 the first column of Table 15.1, the procedures specified in 15.2.3, C_PTR, C_FUNPTR, C_NULL_PTR,
17 and C_NULL_FUNPTR. A processor may provide other public entities in the ISO_C_BINDING intrinsic
18 module in addition to those listed here.

NOTE 15.1

To avoid potential name conflicts with program entities, it is recommended that a program use the ONLY option in any USE statement that references the ISO_C_BINDING intrinsic module.
--

19 15.2.2 Named constants and derived types in the module

20 The entities listed in the second column of Table 15.2, shall be named constants of type default integer.

21 The value of C_INT shall be a valid value for an integer kind parameter on the processor. The values of C_
22 SHORT, C_LONG, C_LONG_LONG, C_SIGNED_CHAR, C_SIZE_T, C_INT8_T, C_INT16_T, C_INT32_
23 T, C_INT64_T, C_INT_LEAST8_T, C_INT_LEAST16_T, C_INT_LEAST32_T, C_INT_LEAST64_T, C_
24 INT_FAST8_T, C_INT_FAST16_T, C_INT_FAST32_T, C_INT_FAST64_T, C_INTMAX_T, and C_INT_
25 PTR_T shall each be a valid value for an integer kind type parameter on the processor or shall be -1
26 if the companion C processor defines the corresponding C type and there is no interoperating Fortran
27 processor kind or -2 if the C processor does not define the corresponding C type.

28 The values of C_FLOAT, C_DOUBLE, and C_LONG_DOUBLE shall each be a valid value for a real
29 kind type parameter on the processor or shall be -1 if the C processor's type does not have a precision
30 equal to the precision of any of the Fortran processor's real kinds, -2 if the C processor's type does not
31 have a range equal to the range of any of the Fortran processor's real kinds, -3 if the C processor's type
32 has neither the precision nor range of any of the Fortran processor's real kinds, and equal to -4 if there
33 is no interoperating Fortran processor kind for other reasons. The values of C_FLOAT_COMPLEX,
34 C_DOUBLE_COMPLEX, and C_LONG_DOUBLE_COMPLEX shall be the same as those of C_FLOAT,

1 C_DOUBLE, and C_LONG_DOUBLE, respectively.

NOTE 15.2

If the C processor supports more than one variety of float, double or long double, the Fortran processor might find it helpful to select from among more than one ISO_C_BINDING module by a processor dependent means.

2 The value of C_BOOL shall be a valid value for a logical kind parameter on the processor or shall be -1.

3 The value of C_INT_BITS shall be a valid value for a bits kind type parameter of the processor.

4 The values of C_SHORT_BITS, C_LONG_BITS, C_LONG_LONG_BITS, C_SIGNED_CHAR_BITS, C_-
 5 INT8_T_BITS, C_INT16_T_BITS, C_INT32_T_BITS, C_INT64_T_BITS, C_INT_LEAST8_T_BITS, C_-
 6 INT_LEAST16_T_BITS, C_INT_LEAST32_T_BITS, C_INT_LEAST64_T_BITS, C_INT_FAST8_T_BITS,
 7 C_INT_FAST16_T_BITS, C_INT_FAST32_T_BITS, C_INT_FAST64_T_BITS, C_INTMAX_T_BITS, C_-
 8 INTPTR_T_BITS, and C_BOOL_BITS shall each be a valid value for a bits kind type parameter on the
 9 processor, -1 if the companion C processor defines the corresponding C type and there is no interoper-
 10 ating Fortran processor kind, or -2 if the C processor does not define the corresponding C type.

11 The value of C_CHAR shall be a valid value for a character kind type parameter on the processor or
 12 shall be -1. The value of C_CHAR is known as the **C character kind**.

13 The following entities shall be named constants of type character with a length parameter of one. The
 14 kind parameter value shall be equal to the value of C_CHAR unless C_CHAR = -1, in which case the
 15 kind parameter value shall be the same as for default kind. The values of these constants are specified
 16 in Table 15.1. In the case that C_CHAR \neq -1 the value is specified using C syntax. The semantics of
 17 these values are explained in 5.2.1 and 5.2.2 of the C International Standard.

Table 15.1: Names of C characters with special semantics

Name	C definition	Value	
		C_CHAR = -1	C_CHAR \neq -1
C_NULL_CHAR	null character	CHAR(0)	'\0'
C_ALERT	alert	ACHAR(7)	'\a'
C_BACKSPACE	backspace	ACHAR(8)	'\b'
C_FORM_FEED	form feed	ACHAR(12)	'\f'
C_NEW_LINE	new line	ACHAR(10)	'\n'
C_CARRIAGE_RETURN	carriage return	ACHAR(13)	'\r'
C_HORIZONTAL_TAB	horizontal tab	ACHAR(9)	'\t'
C_VERTICAL_TAB	vertical tab	ACHAR(11)	'\v'

NOTE 15.3

The value of NEW_LINE(C_NEW_LINE) is C_NEW_LINE (13.7.127).

18 The entities C_PTR and C_FUNPTR are described in 15.3.3.

19 The entity C_NULL_PTR shall be a named constant of type C_PTR. The value of C_NULL_PTR shall
 20 be the same as the value NULL in C. The entity C_NULL_FUNPTR shall be a named constant of type
 21 C_FUNPTR. The value of C_NULL_FUNPTR shall be that of a null pointer to a function in C.

22 **15.2.3 Procedures in the module**

23 In the detailed descriptions below, procedure names are generic and not specific.

24 A C procedure argument is often defined in terms of a **C address**. The C_LOC and C_FUNLOC

1 functions are provided so that Fortran applications can determine the appropriate value to use with C
 2 facilities. The C_ASSOCIATED function is provided so that Fortran programs can compare C addresses.
 3 The C_F_POINTER and C_F_PROCPOINTER subroutines provide a means of associating a Fortran
 4 pointer with the target of a C pointer.

5 15.2.3.1 C_ASSOCIATED (C_PTR_1 [, C_PTR_2])

6 **Description.** Indicates the association status of C_PTR_1 or indicates whether C_PTR_1 and
 7 C_PTR_2 are associated with the same entity.

8 **Class.** Inquiry function.

9 **Arguments.**

10 C_PTR_1 shall be a scalar of type C_PTR or C_FUNPTR.

11 C_PTR_2 shall be a scalar of the same type as C_PTR_1.
 (optional)

12 **Result Characteristics.** Default logical scalar.

13 **Result Value.**

14 *Case (i):* If C_PTR_2 is absent, the result is false if C_PTR_1 is a C null pointer and true
 15 otherwise.

16 *Case (ii):* If C_PTR_2 is present, the result is false if C_PTR_1 is a C null pointer. Otherwise,
 17 the result is true if C_PTR_1 compares equal to C_PTR_2 in the sense of 6.3.2.3
 18 and 6.5.9 of the C International Standard, and false otherwise.

NOTE 15.4

The following example illustrates the use of C_LOC and C_ASSOCIATED.

```

USE, INTRINSIC :: ISO_C_BINDING, ONLY: C_PTR, C_FLOAT, C_ASSOCIATED, C_LOC
INTERFACE
  SUBROUTINE FOO(GAMMA) BIND(C)
    IMPORT C_PTR
    TYPE(C_PTR), VALUE :: GAMMA
  END SUBROUTINE FOO
END INTERFACE
REAL(C_FLOAT), TARGET, DIMENSION(100) :: ALPHA
TYPE(C_PTR) :: BETA
...
IF (.NOT. C_ASSOCIATED(BETA)) THEN
  BETA = C_LOC(ALPHA)
ENDIF
CALL FOO(BETA)

```

19 15.2.3.2 C_F_POINTER (CPTR, FPTR [, SHAPE])

20 **Description.** Associates a data pointer with the target of a C pointer and specifies its shape.

21 **Class.** Subroutine.

1 **Arguments.**

CPTR shall be a scalar of type C_PTR. It is an INTENT(IN) argument. Its value shall be

- (1) the C address of an interoperable data entity, or
- (2) the result of a reference to C_LOC with a noninteroperable argument.

2

The value of CPTR shall not be the C address of a Fortran variable that does not have the TARGET attribute.

3

FPTR shall be a pointer, and shall not be a co-indexed object. It is an INTENT(OUT) argument.

If the value of CPTR is the C address of an interoperable data entity, FPTR shall be a data pointer with type and type parameters interoperable with the type of the entity. In this case, FPTR becomes pointer associated with the target of CPTR. If FPTR is an array, its shape is specified by SHAPE and each lower bound is 1.

If the value of CPTR is the result of a reference to C_LOC with a noninteroperable argument X, FPTR shall be a nonpolymorphic scalar pointer with the same type and type parameters as X. In this case, X or its target if it is a pointer shall not have been deallocated or have become undefined due to execution of a RETURN or END statement since the reference. FPTR becomes pointer associated with X or its target.

4

SHAPE
(optional) shall be of type integer and rank one. It is an INTENT(IN) argument. SHAPE shall be present if and only if FPTR is an array; its size shall be equal to the rank of FPTR.

5

6 **15.2.3.3 C_F_PROCPOINTER (CPTR, FPTR)**

7 **Description.** Associates a procedure pointer with the target of a C function pointer.

8 **Class.** Subroutine.

9 **Arguments.**

CPTR shall be a scalar of type C_FUNPTR. It is an INTENT(IN) argument. Its value shall be the C address of a procedure that is interoperable.

10

FPTR shall be a procedure pointer, and shall not be a co-indexed object. It is an INTENT(OUT) argument. The interface for FPTR shall be interoperable with the prototype that describes the target of CPTR. FPTR becomes pointer associated with the target of CPTR.

11

NOTE 15.5

The term “target” in the descriptions of C_F_POINTER and C_F_PROCPOINTER denotes the entity referenced by a C pointer, as described in 6.2.5 of the C International Standard.

12 **15.2.3.4 C_FUNLOC (X)**

13 **Description.** Returns the C address of the argument.

14 **Class.** Inquiry function.

15 **Argument.** X shall either be a procedure that is interoperable, or a procedure pointer associated with an interoperable procedure. It shall not be a co-indexed object.

16

1 **Result Characteristics.** Scalar of type C_FUNPTR.

2 **Result Value.**

3 The result value is described using the result name CPTR. The result is determined as if C_-
4 FUNPTR were a derived type containing an implicit-interface procedure pointer component PX
5 and the pointer assignment CPTR%PX => X were executed.

6 The result is a value that can be used as an actual CPTR argument in a call to C_F_PROC-
7 POINTER where FPTR has attributes that would allow the pointer assignment FPTR => X.
8 Such a call to C_F_PROCPOINTER shall have the effect of the pointer assignment FPTR =>
9 X.

10 **15.2.3.5 C_LOC (X)**

11 **Description.** Returns the C address of the argument.

12 **Class.** Inquiry function.

13 **Argument.** X shall have either the POINTER or TARGET attribute. It shall not be a co-indexed
14 object. It shall either be a contiguous variable with interoperable type and type parameters, or
15 be a scalar, nonpolymorphic variable with no length type parameters. If it is allocatable, it shall
16 be allocated. If it is a pointer, it shall be associated. If it is an array, it shall have nonzero size.

17 **Result Characteristics.** Scalar of type C_PTR.

18 **Result Value.**

19 The result value is described using the result name CPTR.

20 If X is a scalar data entity, the result is determined as if C_PTR were a derived type containing a scalar
21 pointer component PX of the type and type parameters of X and the pointer assignment CPTR%PX
22 => X were executed.

23 If X is an array data entity, the result is determined as if C_PTR were a derived type containing a scalar
24 pointer component PX of the type and type parameters of X and the pointer assignment of CPTR%PX
25 to the first element of X were executed.

26 If X is a data entity that is interoperable or has interoperable type and type parameters, the
27 result is the value that the C processor returns as the result of applying the unary “&” operator
28 (as defined in the C International Standard, 6.5.3.2) to the target of CPTR%PX.

29 The result is a value that can be used as an actual CPTR argument in a call to C_F_POINTER
30 where FPTR has attributes that would allow the pointer assignment FPTR => X. Such a call
31 to C_F_POINTER shall have the effect of the pointer assignment FPTR => X.

NOTE 15.6

Where the actual argument is of noninteroperable type or type parameters, the result of C_LOC provides an opaque “handle” for it. In an actual implementation, this handle might be the C address of the argument; however, portable C functions should treat it as a void (generic) C pointer that cannot be dereferenced (6.5.3.2 in the C International Standard).

32 **15.2.3.6 C_SIZEOF (X)**

33 **Description.** Returns the size of X in bytes.

Class. Inquiry function.

1 **Argument.** X shall be an interoperable data entity that is not an assumed-size array.

2 **Result Characteristics.** Scalar integer of kind `C_SIZE_T` (15.3.2).

3 **Result Value.**

4 If X is scalar, the result value is the value that the C processor returns as the result of applying
5 the `sizeof` operator (C International Standard, subclause 6.5.3.4) to an object of a type that
6 interoperates with the type and type parameters of X.

7 If X is an array, the result value is the value that the C processor returns as the result of applying
8 the `sizeof` operator to an object of a type that interoperates with the type and type parameters of X,
9 multiplied by the number of elements in X.

10 15.3 Interoperability between Fortran and C entities

11 15.3.1 General

12 Subclause 15.3 defines the conditions under which a Fortran entity is interoperable. If a Fortran entity is
13 interoperable, an equivalent entity may be defined by means of C and the Fortran entity **interoperates**
14 with the C entity. There does not have to be such an interoperating C entity.

NOTE 15.7

A Fortran entity can be interoperable with more than one C entity.

15 15.3.2 Interoperability of intrinsic types

16 Table 15.2 shows the interoperability between Fortran intrinsic types and C types. A Fortran intrinsic
17 type with particular type parameter values is interoperable with a C type if the type and kind type
18 parameter value are listed in the table on the same row as that C type; if the type is character, inter-
19 operability also requires that the length type parameter be omitted or be specified by an initialization
20 expression whose value is one. A combination of Fortran type and type parameters that is interoperable
21 with a C type listed in the table is also interoperable with any unqualified C type that is compatible
22 with the listed C type.

23 The second column of the table refers to the named constants made accessible by the `ISO_C_BINDING`
24 intrinsic module. If the value of any of these named constants is negative, there is no combination of
25 Fortran type and type parameters interoperable with the C type shown in that row.

26 A combination of intrinsic type and type parameters is **interoperable** if it is interoperable with a C
27 type.

Table 15.2: Interoperability between Fortran and C types

Fortran type	Named constant from the <code>ISO_C_BINDING</code> module (kind type parameter if value is positive)	C type
	<code>C_INT</code>	int
	<code>C_SHORT</code>	short int
	<code>C_LONG</code>	long int
	<code>C_LONG_LONG</code>	long long int
	<code>C_SIGNED_CHAR</code>	signed char unsigned char
	<code>C_SIZE_T</code>	size_t

Interoperability between Fortran and C types

(cont.)

Fortran type	Named constant from the ISO_C_BINDING module (kind type parameter if value is positive)	C type
INTEGER	C_INT8_T	int8_t
	C_INT16_T	int16_t
	C_INT32_T	int32_t
	C_INT64_T	int64_t
	C_INT_LEAST8_T	int_least8_t
	C_INT_LEAST16_T	int_least16_t
	C_INT_LEAST32_T	int_least32_t
	C_INT_LEAST64_T	int_least64_t
	C_INT_FAST8_T	int_fast8_t
	C_INT_FAST16_T	int_fast16_t
	C_INT_FAST32_T	int_fast32_t
	C_INT_FAST64_T	int_fast64_t
	C_INTMAX_T	intmax_t
	C_INTPTR_T	intptr_t
REAL	C_FLOAT	float
	C_DOUBLE	double
	C_LONG_DOUBLE	long double
COMPLEX	C_FLOAT_COMPLEX	float _Complex
	C_DOUBLE_COMPLEX	double _Complex
	C_LONG_DOUBLE_COMPLEX	long double _Complex
LOGICAL	C_BOOL	_Bool
BITS	C_INT_BITS	unsigned int or int
	C_SHORT_BITS	unsigned short or short
	C_LONG_BITS	unsigned long or long
	C_LONG_LONG_BITS	unsigned long long long long
	C_SIGNED_CHAR_BITS	unsigned char signed char
	C_INT8_T_BITS	uint8_t or int8_t
	C_INT16_T_BITS	uint16_t or int16_t
	C_INT32_T_BITS	uint32_t or int32_t
	C_INT64_T_BITS	uint64_t or int64_t
	C_INT_LEAST8_T_BITS	uint_least8_t int_least8_t
	C_INT_LEAST16_T_BITS	uint_least16_t int_least16_t
	C_INT_LEAST32_T_BITS	uint_least32_t int_least32_t
	C_INT_LEAST64_T_BITS	uint_least64_t int_least64_t
	C_INT_FAST8_T_BITS	uint_fast8_t or int_fast8_t

Interoperability between Fortran and C types

(cont.)

Fortran type	Named constant from the ISO_C_BINDING module (kind type parameter if value is positive)	C type
	C_INT_FAST16_T_BITS	uint_fast16_t int_fast16_t
	C_INT_FAST32_T_BITS	uint_fast32_t int_fast32_t
	C_INT_FAST64_T_BITS	uint_fast64_t int_fast64_t
	C_INTMAX_T_BITS	uintmax_t or intmax_t
	C_INTPTR_T_BITS	uintptr_t or intptr_t
	C_BOOL_BITS	_Bool
CHARACTER	C_CHAR	char
The above mentioned C types are defined in the C International Standard, subclauses 6.2.5, 7.17, and 7.18.1.		

NOTE 15.8

For example, the type integer with a kind type parameter of C_SHORT is interoperable with the C type short or any C type derived (via typedef) from short.

NOTE 15.9

The C International Standard specifies that the representations for nonnegative signed integers are the same as the corresponding values of unsigned integers. A user can use the signed kinds of integers to interoperate with the unsigned types and all their qualified versions as well. This has the potentially surprising side effect that the C type unsigned char is interoperable with the type integer with a kind type parameter of C_SIGNED_CHAR.

NOTE 15.10

If a variable of type bits is the actual argument corresponding to an unsigned integer parameter of a C function and is interoperable with that parameter, or the unsigned integer result of a C function is assigned to a variable of type bits that is interoperable with the function result, the I format can be used to output the correct form of the unsigned integer value.

1 15.3.3 Interoperability with C pointer types

- 2 C_PTR and C_FUNPTR shall be derived types with only private components. C_PTR is interoperable
3 with any C object pointer type. C_FUNPTR is interoperable with any C function pointer type.

NOTE 15.11

This implies that a C processor is required to have the same representation method for all C object pointer types and the same representation method for all C function pointer types if the C processor is to be the target of interoperability of a Fortran processor. The C International Standard does not impose this requirement.

NOTE 15.12

The function C_LOC can be used to return a value of type C_PTR that is the C address of an allocated allocatable variable. The function C_FUNLOC can be used to return a value of type

NOTE 15.12 (cont.)

C_FUNPTR that is the C address of a procedure. For C_LOC and C_FUNLOC the returned value is of an interoperable type and thus may be used in contexts where the procedure or allocatable variable is not directly allowed. For example, it could be passed as an actual argument to a C function.

Similarly, type C_FUNPTR or C_PTR can be used in a dummy argument or structure component and can have a value that is the C address of a procedure or allocatable variable, even in contexts where a procedure or allocatable variable is not directly allowed.

1 15.3.4 Interoperability of derived types and C struct types

2 A Fortran derived type is **interoperable** if it has the BIND attribute.

3 C1501 (R430) A derived type with the BIND attribute shall not be a SEQUENCE type.

4 C1502 (R430) A derived type with the BIND attribute shall not have type parameters.

5 C1503 (R430) A derived type with the BIND attribute shall not have the EXTENDS attribute.

6 C1504 (R430) A derived type with the BIND attribute shall not have a *type-bound-procedure-part*.

7 C1505 (R430) Each component of a derived type with the BIND attribute shall be a nonpointer,
8 nonallocatable data component with interoperable type and type parameters.

NOTE 15.13

The syntax rules and their constraints require that a derived type that is interoperable have components that are all data objects that are interoperable. No component is permitted to be a procedure or allocatable, but a component of type C_FUNPTR or C_PTR may hold the C address of such an entity.

9 A Fortran derived type is interoperable with a C struct type if the derived-type definition of the Fortran
10 type specifies BIND(C) (4.5.2), the Fortran derived type and the C struct type have the same number
11 of components, and the components of the Fortran derived type have types and type parameters that
12 are interoperable with the types of the corresponding components of the struct type. A component of
13 a Fortran derived type and a component of a C struct type correspond if they are declared in the same
14 relative position in their respective type definitions.

NOTE 15.14

The names of the corresponding components of the derived type and the C struct type need not be the same.

15 There is no Fortran type that is interoperable with a C struct type that contains a bit field or that
16 contains a flexible array member. There is no Fortran type that is interoperable with a C union type.

NOTE 15.15

For example, the C type myctype, declared below, is interoperable with the Fortran type myftype, declared below.

```
typedef struct {
    int m, n;
    float r;
```

NOTE 15.15 (cont.)

```
} myctype
```

```
USE, INTRINSIC :: ISO_C_BINDING
TYPE, BIND(C) :: MYFTYPE
  INTEGER(C_INT) :: I, J
  REAL(C_FLOAT) :: S
END TYPE MYFTYPE
```

The names of the types and the names of the components are not significant for the purposes of determining whether a Fortran derived type is interoperable with a C struct type.

NOTE 15.16

The C International Standard requires the names and component names to be the same in order for the types to be compatible (C International Standard, subclause 6.2.7). This is similar to Fortran's rule describing when different derived type definitions describe the same sequence type. This rule was not extended to determine whether a Fortran derived type is interoperable with a C struct type because the case of identifiers is significant in C but not in Fortran.

1 **15.3.5 Interoperability of scalar variables**

2 A scalar Fortran variable is **interoperable** if its type and type parameters are interoperable and it has
3 neither the pointer nor the allocatable attribute.

4 An interoperable scalar Fortran variable is interoperable with a scalar C entity if their types and type
5 parameters are interoperable.

6 **15.3.6 Interoperability of array variables**

7 An array Fortran variable is **interoperable** if its type and type parameters are interoperable and it is
8 of explicit shape or assumed size.

9 An explicit-shape or assumed-size array of rank r , with a shape of $[e_1 \ \dots \ e_r]$ is interoperable with
10 a C array if its size is nonzero and

11 (1) either

12 (a) the array is assumed size, and the C array does not specify a size, or

13 (b) the array is an explicit shape array, and the extent of the last dimension (e_r) is the
14 same as the size of the C array, and

15 (2) either

16 (a) r is equal to one, and an element of the array is interoperable with an element of the
17 C array, or

18 (b) r is greater than one, and an explicit-shape array with shape of $[e_1 \ \dots \ e_{r-1}]$,
19 with the same type and type parameters as the original array, is interoperable with a
20 C array of a type equal to the element type of the original C array.

NOTE 15.17

An element of a multi-dimensional C array is an array type, so a Fortran array of rank one is not interoperable with a multidimensional C array.

NOTE 15.18

A polymorphic, allocatable, or pointer array is never interoperable. Such arrays are not explicit shape or assumed size.

NOTE 15.19

For example, a Fortran array declared as

```
INTEGER :: A(18, 3:7, *)
```

is interoperable with a C array declared as

```
int b[][5][18]
```

NOTE 15.20

The C programming language defines null-terminated strings, which are actually arrays of the C type char that have a C null character in them to indicate the last valid element. A Fortran array of type character with a kind type parameter equal to C.CHAR is interoperable with a C string.

Fortran's rules of sequence association (12.5.2.12) permit a character scalar actual argument to be associated with a dummy argument array. This makes it possible to argument associate a Fortran character string with a C string.

Note 15.24 has an example of interoperation between Fortran and C strings.

1 15.3.7 Interoperability of procedures and procedure interfaces

2 A Fortran procedure is **interoperable** if it has the BIND attribute, that is, if its interface is specified
3 with a *proc-language-binding-spec*.

4 A Fortran procedure interface is interoperable with a C function prototype if

- 5 (1) the interface has the BIND attribute,
- 6 (2) either
 - 7 (a) the interface describes a function whose result variable is a scalar that is interoperable
8 with the result of the prototype or
 - 9 (b) the interface describes a subroutine and the prototype has a result type of void,
- 10 (3) the number of dummy arguments of the interface is equal to the number of formal parameters
11 of the prototype,
- 12 (4) any dummy argument with the VALUE attribute is interoperable with the corresponding
13 formal parameter of the prototype,
- 14 (5) any dummy argument without the VALUE attribute corresponds to a formal parameter
15 of the prototype that is of a pointer type, and the dummy argument is interoperable with
16 an entity of the referenced type (C International Standard, 6.2.5, 7.17, and 7.18.1) of the
17 formal parameter, and
- 18 (6) the prototype does not have variable arguments as denoted by the ellipsis (...).

NOTE 15.21

The **referenced type** of a C pointer type is the C type of the object that the C pointer type points to. For example, the referenced type of the pointer type `int *` is `int`.

NOTE 15.22

The C language allows specification of a C function that can take a variable number of arguments (C International Standard, 7.15). This standard does not provide a mechanism for Fortran procedures to interoperate with such C functions.

- 1 A formal parameter of a C function prototype corresponds to a dummy argument of a Fortran interface if
- 2 they are in the same relative positions in the C parameter list and the dummy argument list, respectively.

NOTE 15.23

For example, a Fortran procedure interface described by

```
INTERFACE
  FUNCTION FUNC(I, J, K, L, M) BIND(C)
    USE, INTRINSIC :: ISO_C_BINDING
    INTEGER(C_SHORT) :: FUNC
    INTEGER(C_INT), VALUE :: I
    REAL(C_DOUBLE) :: J
    INTEGER(C_INT) :: K, L(10)
    TYPE(C_PTR), VALUE :: M
  END FUNCTION FUNC
END INTERFACE
```

is interoperable with the C function prototype

```
short func(int i, double *j, int *k, int l[10], void *m)
```

A C pointer may correspond to a Fortran dummy argument of type C_PTR with the VALUE attribute or to a Fortran scalar that does not have the VALUE attribute. In the above example, the C pointers j and k correspond to the Fortran scalars J and K, respectively, and the C pointer m corresponds to the Fortran dummy argument M of type C_PTR.

NOTE 15.24

The interoperability of Fortran procedure interfaces with C function prototypes is only one part of invocation of a C function from Fortran. There are four pieces to consider in such an invocation: the procedure reference, the Fortran procedure interface, the C function prototype, and the C function. Conversely, the invocation of a Fortran procedure from C involves the function reference, the C function prototype, the Fortran procedure interface, and the Fortran procedure. In order to determine whether a reference is allowed, it is necessary to consider all four pieces.

For example, consider a C function that can be described by the C function prototype

```
void copy(char in[], char out[]);
```

Such a function may be invoked from Fortran as follows:

```
USE, INTRINSIC :: ISO_C_BINDING, ONLY: C_CHAR, C_NULL_CHAR
INTERFACE
  SUBROUTINE COPY(IN, OUT) BIND(C)
    IMPORT C_CHAR
```


NOTE 15.24 (cont.)

```

    CHARACTER(KIND=C_CHAR), DIMENSION(*) :: IN, OUT
    END SUBROUTINE COPY
    END INTERFACE

    CHARACTER(LEN=10, KIND=C_CHAR) :: &
&    DIGIT_STRING = C_CHAR_'123456789' // C_NULL_CHAR
    CHARACTER(KIND=C_CHAR) :: DIGIT_ARR(10)

    CALL COPY(DIGIT_STRING, DIGIT_ARR)
    PRINT '(1X, A1)', DIGIT_ARR(1:9)
    END

```

The procedure reference has character string actual arguments. These correspond to character array dummy arguments in the procedure interface body as allowed by Fortran's rules of sequence association (12.5.2.12). Those array dummy arguments in the procedure interface are interoperable with the formal parameters of the C function prototype. The C function is not shown here, but is assumed to be compatible with the C function prototype.

1 15.4 Interoperation with C global variables

2 15.4.1 General

3 A C variable with external linkage may interoperate with a common block or with a variable declared
4 in the scope of a module. The common block or variable shall be specified to have the BIND attribute.

5 At most one variable that is associated with a particular C variable with external linkage is permitted
6 to be declared within a program. A variable shall not be initially defined by more than one processor.

7 If a common block is specified in a BIND statement, it shall be specified in a BIND statement with
8 the same binding label in each scoping unit in which it is declared. A C variable with external linkage
9 interoperates with a common block that has been specified in a BIND statement

10 (1) if the C variable is of a struct type and the variables that are members of the common block
11 are interoperable with corresponding components of the struct type, or

12 (2) if the common block contains a single variable, and the variable is interoperable with the C
13 variable.

14 There does not have to be an associated C entity for a Fortran entity with the BIND attribute.

NOTE 15.25

The following are examples of the usage of the BIND attribute for variables and for a common block. The Fortran variables, C_EXTERN and C2, interoperate with the C variables, c_extern and myVariable, respectively. The Fortran common blocks, COM and SINGLE, interoperate with the C variables, com and single, respectively.

```

MODULE LINK_TO_C_VARS
    USE, INTRINSIC :: ISO_C_BINDING
    INTEGER(C_INT), BIND(C) :: C_EXTERN
    INTEGER(C_LONG) :: C2

```

NOTE 15.25 (cont.)

```

BIND(C, NAME='myVariable') :: C2

COMMON /COM/ R, S
REAL(C_FLOAT) :: R, S, T
BIND(C) :: /COM/, /SINGLE/
COMMON /SINGLE/ T
END MODULE LINK_TO_C_VARS
int c_extern;
long myVariable;
struct {float r, s;} com;
float single;

```

1 15.4.2 Binding labels for common blocks and variables

2 The **binding label** of a variable or common block is a value of type default character that specifies the
3 name by which the variable or common block is known to the companion processor.

4 If a variable or common block has the BIND attribute with the NAME= specifier and the value of its
5 expression, after discarding leading and trailing blanks, has nonzero length, the variable or common
6 block has this as its binding label. The case of letters in the binding label is significant. If a variable
7 or common block has the BIND attribute specified without a NAME= specifier, the binding label is the
8 same as the name of the entity using lower case letters. Otherwise, the variable or common block has
9 no binding label.

10 The binding label of a C variable with external linkage is the same as the name of the C variable. A
11 Fortran variable or common block with the BIND attribute that has the same binding label as a C
12 variable with external linkage is linkage associated (16.5.1.5) with that variable.

13 15.5 Interoperation with C functions**14 15.5.1 Definition and reference of interoperable procedures**

15 A procedure that is interoperable may be defined either by means other than Fortran or by means of a
16 Fortran subprogram, but not both.

17 If the procedure is defined by means other than Fortran, it shall

- 18 (1) be describable by a C prototype that is interoperable with the interface,
- 19 (2) have external linkage as defined by 6.2.2 of the C International Standard, and
- 20 (3) have the same binding label as the interface.

21 A reference to such a procedure causes the function described by the C prototype to be called as specified
22 in the C International Standard.

23 A reference in C to a procedure that has the BIND attribute, has the same binding label, and is defined
24 by means of Fortran, causes the Fortran procedure to be invoked.

25 A procedure defined by means of Fortran shall not invoke setjmp or longjmp (C International Standard,
26 7.13). If a procedure defined by means other than Fortran invokes setjmp or longjmp, that procedure
27 shall not cause any procedure defined by means of Fortran to be invoked. A procedure defined by means

- 1 of Fortran shall not be invoked as a signal handler (C International Standard, 7.14.1).
 2 If a procedure defined by means of Fortran and a procedure defined by means other than Fortran perform
 3 input/output operations on the same external file, the results are processor dependent (9.4.3).

4 **15.5.2 Binding labels for procedures**

- 5 The **binding label** of a procedure is a value of type default character that specifies the name by which
 6 a procedure with the BIND attribute is known to the companion processor.

- 7 If a procedure has the BIND attribute with the NAME= specifier and the value of its expression, after
 8 discarding leading and trailing blanks, has nonzero length, the procedure has this as its binding label.
 9 The case of letters in the binding label is significant. If a procedure has the BIND attribute with no
 10 NAME= specifier, and the procedure is not a dummy procedure or procedure pointer, then the binding
 11 label of the procedure is the same as the name of the procedure using lower case letters. Otherwise, the
 12 procedure has no binding label.

- 13 C1506 A procedure defined in a submodule shall not have a binding label unless its interface is declared
 14 in the ancestor module.

- 15 The binding label for a C function with external linkage is the same as the C function name.

NOTE 15.26

In the following sample, the binding label of C_SUB is "c_sub", and the binding label of C_FUNC is "C_func".

```

SUBROUTINE C_SUB() BIND(C)
  ...
END SUBROUTINE C_SUB

INTEGER(C_INT) FUNCTION C_FUNC() BIND(C, NAME="C_func")
  USE, INTRINSIC :: ISO_C_BINDING
  ...
END FUNCTION C_FUNC

```

The C International Standard permits functions to have names that are not permitted as Fortran names; it also distinguishes between names that would be considered as the same name in Fortran. For example, a C name may begin with an underscore, and C names that differ in case are distinct names.

The specification of a binding label allows a program to use a Fortran name to refer to a procedure defined by a companion processor.

16 **15.5.3 Exceptions and IEEE arithmetic procedures**

- 17 A procedure defined by means other than Fortran shall not use signal (C International Standard, 7.14.1)
 18 to change the handling of any exception that is being handled by the Fortran processor.

- 19 A procedure defined by means other than Fortran shall not alter the floating point status (14.6) other
 20 than by setting an exception flag to signaling.

- 21 The values of the floating point exception flags on entry to a procedure defined by means other than
 22 Fortran are processor-dependent.

1 16 Scope, association, and definition

2 16.1 Identifiers and entities

3 Entities are identified by identifiers within a **scope** that is a program, a scoping unit, a construct, a
4 single statement, or part of a statement.

- 5 • A **global identifier** has a scope of a program (2.2.2);
- 6 • A **local identifier** has a scope of a scoping unit (2.2);
- 7 • An identifier of a **construct entity** has a scope of a construct (7.4.3, 7.4.4, 8.1);
- 8 • An identifier of a **statement entity** has a scope of a statement or part of a statement (3.3).

9 An entity may be identified by

- 10 (1) an image index (2.3.2),
- 11 (2) a name (3.2.1),
- 12 (3) a statement label (3.2.4),
- 13 (4) an external input/output unit number (9.4),
- 14 (5) an identifier of a pending data transfer operation (9.5.2.9, 9.6),
- 15 (6) a submodule identifier (11.2.3),
- 16 (7) a generic identifier (12.4.3.2), or
- 17 (8) a binding label (15.5.2, 15.4.2).

18 By means of association, an entity may be referred to by the same identifier or a different identifier in
19 a different scoping unit, or by a different identifier in the same scoping unit.

20 16.2 Scope of global identifiers

21 Program units, common blocks, external procedures, entities with binding labels, and images are **global**
22 **entities** of a program. The name of a non-submodule program unit, common block, or external pro-
23 cedure is a global identifier and shall not be the same as the name of any other such global entity in
24 the same program, except that an intrinsic module may have the same name as another program unit,
25 common block, or external procedure in the same program. The submodule identifier of a submodule
26 is a global identifier and shall not be the same as the submodule identifier of any other submodule. A
27 binding label of an entity of the program is a global identifier and shall not be the same as the binding
28 label of any other entity of the program; nor shall it be the same as the name of any other global entity
29 of the program that is not an intrinsic module, ignoring differences in case. An entity of the program
30 shall not be identified by more than one binding label.

NOTE 16.1

The name of a global entity may be the same as a binding label that identifies the same global entity.

NOTE 16.2

Of the various types of procedures, only external procedures have global names. An implementation may wish to assign global names to other entities in the Fortran program such as internal procedures, intrinsic procedures, procedures implementing intrinsic operators, procedures implementing input/output operations, etc. If this is done, it is the responsibility of the processor to ensure that none of these names conflicts with any of the names of the external procedures, with other globally named entities in a standard-conforming program, or with each other. For example, this might be done by including in each such added name a character that is not allowed in a standard-conforming name or by using such a character to combine a local designation with the global name of the program unit in which it appears.

NOTE 16.3

Submodule identifiers are global identifiers, but because they consist of a module name and a descendant submodule name, the name of a submodule can be the same as the name of another submodule so long as they do not have the same ancestor module.

1 External input/output units and pending data transfer operations are global entities.

2 **16.3 Scope of local identifiers**

3 **16.3.1 Classes of local identifiers**

4 Within a scoping unit, identifiers of entities in the classes

- 5 (1) named variables that are not statement or construct entities (16.4), named constants, named
- 6 constructs, statement functions, internal procedures, module procedures, dummy procedures,
- 7 intrinsic procedures, abstract interfaces, module procedure interfaces, generic interfaces, de-
- 8 rived types, namelist groups, external procedures accessed via USE, macros, and statement
- 9 labels,
- 10 (2) type parameters, components, and type-bound procedure bindings, in a separate class for
- 11 each type, and
- 12 (3) argument keywords, in a separate class for each procedure with an explicit interface

13 are local identifiers in that scoping unit.

14 Within a scoping unit, a local identifier of an entity of class (1) shall not be the same as a global identifier
15 used in that scoping unit unless the global identifier

- 16 (1) is used only as the *use-name* of a *rename* in a USE statement,
- 17 (2) is a common block name (16.3.2),
- 18 (3) is an external procedure name that is also a generic name, or
- 19 (4) is an external function name and the scoping unit is its defining subprogram (16.3.3).

20 Within a scoping unit, a local identifier of one class shall not be the same as another local identifier of
21 the same class, except that a generic name may be the same as the name of a procedure as explained
22 in 12.4.3.2 or the same as the name of a derived type (4.5.10), and a separate module procedure shall
23 have the same name as its corresponding module procedure interface body (12.6.2.4). A local identifier
24 of one class may be the same as a local identifier of another class.

NOTE 16.4

An intrinsic procedure is inaccessible by its own name in a scoping unit that uses the same name as a local identifier of class (1) for a different entity. For example, in the program fragment

NOTE 16.4 (cont.)

```
SUBROUTINE SUB
```

```
...
```

```
A = SIN (K)
```

```
...
```

```
CONTAINS
```

```
FUNCTION SIN (X)
```

```
...
```

```
END FUNCTION SIN
```

```
END SUBROUTINE SUB
```

any reference to function SIN in subroutine SUB refers to the internal function SIN, not to the intrinsic function of the same name.

1 A local identifier identifies an entity in a scoping unit and may be used to identify an entity in another
2 scoping unit except in the following cases.

3 (1) The name that appears as a *subroutine-name* in a *subroutine-stmt* has limited use within
4 the scope established by the *subroutine-stmt*. It can be used to identify recursive references
5 of the subroutine or to identify a common block (the latter is possible only for internal and
6 module subroutines).

7 (2) The name that appears as a *function-name* in a *function-stmt* has limited use within the
8 scope established by that *function-stmt*. It can be used to identify the result variable, to
9 identify recursive references of the function, or to identify a common block (the latter is
10 possible only for internal and module functions).

11 (3) The name that appears as an *entry-name* in an *entry-stmt* has limited use within the scope
12 of the subprogram in which the *entry-stmt* appears. It can be used to identify the result
13 variable if the subprogram is a function, to identify recursive references, or to identify a
14 common block (the latter is possible only if the *entry-stmt* is in a module subprogram).

15 16.3.2 Local identifiers that are the same as common block names

16 A name that identifies a common block in a scoping unit shall not be used to identify a constant or an
17 intrinsic procedure in that scoping unit. If a local identifier is also the name of a common block, the
18 appearance of that name in any context other than as a common block name in a COMMON or SAVE
19 statement is an appearance of the local identifier.

NOTE 16.5

An intrinsic procedure name may be a common block name in a scoping unit that does not reference
the intrinsic procedure.

20 16.3.3 Function results

21 For each FUNCTION statement or ENTRY statement in a function subprogram, there is a result
22 variable. If there is no RESULT clause, the result variable has the same name as the function being
23 defined; otherwise, the result variable has the name specified in the RESULT clause.

24 16.3.4 Components, type parameters, and bindings

25 A component name has the scope of its derived-type definition. Outside the type definition, it may
26 appear only within a designator of a component of a structure of that type or as a component keyword

1 in a structure constructor for that type.

2 A type parameter name has the scope of its derived-type definition. Outside the derived-type definition,
3 it may appear only as a type parameter keyword in a *derived-type-spec* for the type or as the *type-param-*
4 *name* of a *type-param-inquiry*.

5 The binding name (4.5.5) of a type-bound procedure has the scope of its derived-type definition. Outside
6 of the derived-type definition, it may appear only as the *binding-name* in a procedure reference.

7 A generic binding for which the *generic-spec* is not a *generic-name* has a scope that consists of all scoping
8 units in which an entity of the type is accessible.

9 A component name or binding name may appear only in scoping units in which it is accessible.

10 The accessibility of components and bindings is specified in 4.5.4.7 and 4.5.5.

11 **16.3.5 Argument keywords**

12 As an argument keyword, a dummy argument name in an internal procedure, module procedure, or
13 an interface body has a scope of the scoping unit of the host of the procedure or interface body. It
14 may appear only in a procedure reference for the procedure of which it is a dummy argument. If the
15 procedure or interface body is accessible in another scoping unit by use association or host association
16 (16.5.1.3, 16.5.1.4), the argument keyword is accessible for procedure references for that procedure in
17 that scoping unit.

18 A dummy argument name in an intrinsic procedure has a scope as an argument keyword of the scoping
19 unit in which the reference to the procedure occurs. As an argument keyword, it may appear only in a
20 procedure reference for the procedure of which it is a dummy argument.

21 **16.4 Statement and construct entities**

22 A variable that appears as a *data-i-do-variable* in a DATA statement or an *ac-do-variable* in an array
23 constructor, as a dummy argument in a statement function statement, or as an *index-name* in a FORALL
24 statement is a statement entity. A variable that appears as an *index-name* in a FORALL or DO
25 CONCURRENT construct or as an *associate-name* in a SELECT TYPE or ASSOCIATE construct is
26 a construct entity. A macro local variable is a construct entity. An entity that is declared in a BLOCK
27 construct and is not accessed by use association is a construct entity.

28 If a global or local identifier is the same as that of a construct entity, the identifier is interpreted within
29 the construct as that of the construct entity. Elsewhere in the scoping unit, the identifier is interpreted
30 as the global or local identifier.

31 If a global or local identifier accessible in the scoping unit of a statement is the same as the name of a
32 statement entity in that statement, the name is interpreted within the scope of the statement entity as
33 that of the statement entity. Elsewhere in the scoping unit, including parts of the statement outside the
34 scope of the statement entity, the name is interpreted as the global or local identifier.

35 If the name of a statement entity is the same as the name of a construct entity and the statement is
36 within the scope of the construct entity, the name is interpreted within the scope of the statement entity
37 as that of the statement entity. Elsewhere in the construct, including parts of the statement outside the
38 scope of the statement entity, the name is interpreted as that of the construct entity.

39 Except for a common block name or a scalar variable name, a global identifier or a local identifier of
40 class (1) (16.3) in the scoping unit that contains a statement shall not be the name of a statement entity
41 of that statement. Within the scope of a statement entity, another statement entity shall not have the

1 same name.

2 The name of a *data-i-do-variable* in a DATA statement or an *ac-do-variable* in an array constructor
3 has a scope of its *data-implied-do* or *ac-implied-do*. It is a scalar variable that has the type and type
4 parameters that it would have if it were the name of a variable in the scoping unit that includes the
5 DATA statement or array constructor, and this type shall be integer type; it has no other attributes.
6 The appearance of a name as a *data-i-do-variable* of an implied DO in a DATA statement or an *ac-do-*
7 *variable* in an array constructor is not an implicit declaration of a variable whose scope is the scoping
8 unit that contains the statement.

9 The name of a variable that appears as an *index-name* in a FORALL statement or FORALL or DO
10 CONCURRENT construct has a scope of the statement or construct. It is a scalar variable that has the
11 type and type parameters that it would have if it were the name of a variable in the scoping unit that
12 includes the FORALL, and this type shall be integer type; it has no other attributes. The appearance
13 of a name as an *index-name* in a FORALL statement or FORALL or DO CONCURRENT construct is
14 not an implicit declaration of a variable whose scope is the scoping unit that contains the statement or
15 construct.

16 The name of a variable that appears as a dummy argument in a statement function statement has a scope of the statement
17 in which it appears. It is a scalar that has the type and type parameters that it would have if it were the name of a variable
18 in the scoping unit that includes the statement function; it has no other attributes.

19 Except for a common block name or a scalar variable name, a global identifier or a local identifier of class
20 (1) (16.3) in the scoping unit of a FORALL statement, FORALL construct, or DO CONCURRENT
21 construct shall not be the same as any of its *index-names*. An *index-name* of a contained FORALL
22 statement, FORALL construct, or DO CONCURRENT construct shall not be the same as an *index-*
23 *name* of any of its containing FORALL or DO CONCURRENT constructs.

24 The associate name of a SELECT TYPE construct has a separate scope for each block of the construct.
25 Within each block, it has the declared type, dynamic type, type parameters, rank, and bounds specified
26 in 8.1.9.2.

27 The associate names of an ASSOCIATE construct have the scope of the block. They have the declared
28 type, dynamic type, type parameters, rank, and bounds specified in 8.1.3.2.

29 The macro local variables of a macro definition have the scope of the macro definition.

30 **16.5 Association**

31 **16.5.1 Name association**

32 **16.5.1.1 Forms of name association**

33 There are five forms of **name association**: argument association, use association, host association,
34 linkage association, and construct association. Argument, use, and host association provide mechanisms
35 by which entities known in one scoping unit may be accessed in another scoping unit.

36 **16.5.1.2 Argument association**

37 The rules governing argument association are given in Clause 12. As explained in 12.5, execution of a
38 procedure reference establishes an association between an actual argument and its corresponding dummy
39 argument. Argument association may be sequence association (12.5.2.12).

40 The name of the dummy argument may be different from the name, if any, of its associated actual
41 argument. The dummy argument name is the name by which the associated actual argument is known,

1 and by which it may be accessed, in the referenced procedure.

NOTE 16.6

An actual argument may be a nameless data entity, such as an expression that is not simply a variable or constant.

2 Upon termination of execution of a procedure reference, all argument associations established by that
3 reference are terminated. A dummy argument of that procedure may be associated with an entirely
4 different actual argument in a subsequent invocation of the procedure.

5 16.5.1.3 Use association

6 **Use association** is the association of names in different scoping units specified by a USE statement. The
7 rules governing use association are given in 11.2.2. They allow for renaming of entities being accessed.
8 Use association allows access in one scoping unit to entities defined in another scoping unit; it remains
9 in effect throughout the execution of the program.

10 16.5.1.4 Host association

11 An internal subprogram, a module subprogram, a submodule subprogram, a module procedure interface
12 body, or a derived-type definition has access to entities from its host via **host association**. An interface
13 body that is not a module procedure interface body has access via host association to the named entities
14 from its host that are made accessible by IMPORT statements in the interface body. The accessed
15 entities are identified by the same identifier and have the same attributes as in the host, except that an
16 accessed entity may have the VOLATILE or ASYNCHRONOUS attribute even if the host entity does
17 not. The accessed entities are named data objects, derived types, abstract interfaces, module procedure
18 interfaces, procedures, generic identifiers, macros, and namelist groups.

19 If an entity that is accessed by use association has the same nongeneric name as a host entity, the host
20 entity is inaccessible by that name. Within the scoping unit, a name that is declared to be an external
21 procedure name by an *external-stmt*, *procedure-declaration-stmt*, or *interface-body*, or that appears as a
22 *module-name* in a *use-stmt* is a global identifier; any entity of the host that has this as its nongeneric
23 name is inaccessible by that name. A name that appears in the scoping unit as

- 24 (1) a *function-name* in a *stmt-function-stmt* or in an *entity-decl* in a *type-declaration-stmt*,
- 25 (2) an *object-name* in an *entity-decl* in a *type-declaration-stmt*, in a *pointer-stmt*, in a *save-stmt*,
26 in an *allocatable-stmt*, or in a *target-stmt*,
- 27 (3) a *type-param-name* in a *derived-type-stmt*,
- 28 (4) a *named-constant* in a *named-constant-def* in a *parameter-stmt*,
- 29 (5) an *array-name* in a *dimension-stmt*,
- 30 (6) a *variable-name* in a *common-block-object* in a *common-stmt*,
- 31 (7) a *proc-pointer-name* in a *common-block-object* in a *common-stmt*,
- 32 (8) the name of a variable that is wholly or partially initialized in a *data-stmt*,
- 33 (9) the name of an object that is wholly or partially equivalenced in an *equivalence-stmt*,
- 34 (10) a *dummy-arg-name* in a *function-stmt*, in a *subroutine-stmt*, in an *entry-stmt*, or in a *stmt-*
35 *function-stmt*,
- 36 (11) a *result-name* in a *function-stmt* or in an *entry-stmt*,
- 37 (12) the name of an entity declared by an interface body,
- 38 (13) an *intrinsic-procedure-name* in an *intrinsic-stmt*,
- 39 (14) a *namelist-group-name* in a *namelist-stmt*,
- 40 (15) the name of a macro in a *define-macro-stmt*,
- 41 (16) a *generic-name* in a *generic-spec* in an *interface-stmt*, or

- 1 (17) the name of a named construct
- 2 is a local identifier in the scoping unit and any entity of the host that has this as its nongeneric name is
 3 inaccessible by that name by host association. If a scoping unit is the host of a derived-type definition
 4 or a subprogram that does not define a separate module procedure, the name of the derived type or of
 5 any procedure defined by the subprogram is a local identifier in the scoping unit; any entity of the host
 6 that has this as its nongeneric name is inaccessible by that name. Local identifiers of a subprogram are
 7 not accessible to its host.

NOTE 16.7

A name that appears in an `ASYNCHRONOUS` or `VOLATILE` statement is not necessarily the name of a local variable. In an internal or module procedure, if a variable that is accessible via host association is specified in an `ASYNCHRONOUS` or `VOLATILE` statement, that host variable is given the `ASYNCHRONOUS` or `VOLATILE` attribute in the local scope.

- 8 If a host entity is inaccessible only because a local variable with the same name is wholly or partially
 9 initialized in a `DATA` statement, the local variable shall not be referenced or defined prior to the `DATA`
 10 statement.
- 11 If a derived-type name of a host is inaccessible, data entities of that type or subobjects of such data
 12 entities still can be accessible.

NOTE 16.8

An interface body that is not a module procedure interface body accesses by host association only those entities made accessible by `IMPORT` statements.

- 13 If an external or dummy procedure with an implicit interface is accessed via host association, then it
 14 shall have the `EXTERNAL` attribute in the host scoping unit; if it is invoked as a function in the inner
 15 scoping unit, its type and type parameters shall be established in the host scoping unit. The type and
 16 type parameters of a function with the `EXTERNAL` attribute are established in a scoping unit if that
 17 scoping unit explicitly declares them, invokes the function, accesses the function from a module, or
 18 accesses the function from its host where its type and type parameters are established.
- 19 If an intrinsic procedure is accessed via host association, then it shall be established to be intrinsic in the
 20 host scoping unit. An intrinsic procedure is established to be intrinsic in a scoping unit if that scoping
 21 unit explicitly gives it the `INTRINSIC` attribute, invokes it as an intrinsic procedure, accesses it from a
 22 module, or accesses it from its host where it is established to be intrinsic.

NOTE 16.9

A host subprogram and an internal subprogram may contain the same and differing use-associated entities, as illustrated in the following example.

```

MODULE B; REAL BX, Q; INTEGER IX, JX; END MODULE B
MODULE C; REAL CX; END MODULE C
MODULE D; REAL DX, DY, DZ; END MODULE D
MODULE E; REAL EX, EY, EZ; END MODULE E
MODULE F; REAL FX; END MODULE F
MODULE G; USE F; REAL GX; END MODULE G
PROGRAM A
USE B; USE C; USE D
...

```

NOTE 16.9 (cont.)

CONTAINS

SUBROUTINE INNER_PROC (Q)

USE C ! Not needed

USE B, ONLY: BX ! Entities accessible are BX, IX, and JX

! if no other IX or JX

! is accessible to INNER_PROC

! Q is local to INNER_PROC,

! because Q is a dummy argument

USE D, X => DX ! Entities accessible are DX, DY, and DZ

! X is local name for DX in INNER_PROC

! X and DX denote same entity if no other

! entity DX is local to INNER_PROC

USE E, ONLY: EX ! EX is accessible in INNER_PROC, not in program A

! EY and EZ are not accessible in INNER_PROC

! or in program A

USE G ! FX and GX are accessible in INNER_PROC

...

END SUBROUTINE INNER_PROC

END PROGRAM A

Because program A contains the statement

USE B

all of the entities in module B, except for Q, are accessible in INNER_PROC, even though INNER-PROC contains the statement

USE B, ONLY: BX

The USE statement with the ONLY option means that this particular statement brings in only the entity named, not that this is the only variable from the module accessible in this scoping unit.

NOTE 16.10

For more examples of host association, see subclause C.12.1.

1 16.5.1.5 Linkage association

2 Linkage association occurs between a module variable that has the BIND attribute and the C variable
 3 with which it interoperates, or between a Fortran common block and the C variable with which it
 4 interoperates (15.4). Such association remains in effect throughout the execution of the program.

5 16.5.1.6 Construct association

6 Execution of a SELECT TYPE statement establishes an association between the selector and the asso-
 7 ciate name of the construct. Execution of an ASSOCIATE statement establishes an association between
 8 each selector and the corresponding associate name of the construct.

- 1 If the selector is allocatable, it shall be allocated; the associate name is associated with the data object
2 and does not have the ALLOCATABLE attribute.
- 3 If the selector has the POINTER attribute, it shall be associated; the associate name is associated with
4 the target of the pointer and does not have the POINTER attribute.
- 5 If the selector is a variable other than an array section having a vector subscript, the association is
6 with the data object specified by the selector; otherwise, the association is with the value of the selector
7 expression, which is evaluated prior to execution of the block.
- 8 Each associate name remains associated with the corresponding selector throughout the execution of the
9 executed block. Within the block, each selector is known by and may be accessed by the corresponding
10 associate name. Upon termination of the construct, the association is terminated.

NOTE 16.11

The association between the associate name and a data object is established prior to execution of the block and is not affected by subsequent changes to variables that were used in subscripts or substring ranges in the *selector*.

11 16.5.2 Pointer association**12 16.5.2.1 General**

13 Pointer association between a pointer and a target allows the target to be referenced by a reference to the
14 pointer. At different times during the execution of a program, a pointer may be undefined, associated
15 with different targets, or be disassociated. If a pointer is associated with a target, the definition status
16 of the pointer is either defined or undefined, depending on the definition status of the target. If the
17 pointer has deferred type parameters or shape, their values are assumed from the target. If the pointer
18 is polymorphic, its dynamic type is the dynamic type of the target.

19 16.5.2.2 Pointer association status

20 A pointer may have a **pointer association status** of associated, disassociated, or undefined. Its
21 association status may change during execution of a program. Unless a pointer is initialized (explicitly
22 or by default), it has an initial association status of undefined. A pointer may be initialized to have an
23 association status of disassociated or associated.

NOTE 16.12

A pointer from a module program unit may be accessible in a subprogram via use association. Such pointers have a lifetime that is greater than targets that are declared in the subprogram, unless such targets are saved. Therefore, if such a pointer is associated with a local target, there is the possibility that when a procedure defined by the subprogram completes execution, the target will cease to exist, leaving the pointer “dangling”. This standard considers such pointers to have an undefined association status. They are neither associated nor disassociated. They shall not be used again in the program until their status has been reestablished. There is no requirement on a processor to be able to detect when a pointer target ceases to exist.

24 16.5.2.2.1 Events that cause pointers to become associated

25 A pointer becomes associated when any of the following events occur.

- 26 • The pointer is allocated (6.3.1) as the result of the successful execution of an ALLOCATE statement
27 referencing the pointer.

- 1 • The pointer is pointer-assigned to a target (7.4.2) that is associated or is specified with the TAR-
- 2 GET attribute and, if allocatable, is allocated.
- 3 • The pointer is a dummy argument and its corresponding actual argument is not a pointer.
- 4 • The pointer is an ultimate component of an object of a type for which default initialization is
- 5 specified for the component, and the corresponding initializer is an initialization target, and
 - 6 – a procedure is invoked with this object as an actual argument corresponding to a nonpointer
 - 7 nonallocatable dummy argument with INTENT (OUT),
 - 8 – a procedure with this object as an unsaved nonpointer nonallocatable local object that is not
 - 9 accessed by use or host association is invoked, or
 - 10 – this object is allocated.

J3 internal note

Unresolved Technical Issue 076

Defining a bits entity associated with several numeric entities now causes all the numeric entities to become defined.

This has swapped the “bits doesn’t define anything” problem with the “bits define too much” problem.

I will draw people’s attention to one of the *primary* goals of the Fortran standard, which is to promote portability. This runs directly counter to that goal.

Since bits used in this manner are inherently nonportable, we need to limit them, not throw out the portability baby.

Furthermore, “invalid values”, i.e. values that don’t act consistently, or indeed crash the program, are not limited to LOGICALs.

We really *ought not* to go further than C here. C explicitly acknowledges that there may be “trap representations”; trap in this context means *don’t fall into it!* We ought not to attempt to require implementations not to have trap representations!!!!

Furthermore, note that defining a double precision entity does not define an associated single precision complex entity (in *all* previous Fortrans). And defining a real doesn’t define an associated integer. Allowing bits definition to do that (define two incompatible entities – note well that the real and integer, or the single complex and double real, are *not* bits compatible!), is bad.

It seems that the best way of resolving this tension between the goal of portability, the reality of trap representations, and the desire for unsafe low-level bit twiddling, is to say that whether the associated entity becomes defined (or does not become undefined) is *processor dependent*.

It’s not my preferred solution, but it seems to be the one which provides most traction to the feature with least effort in standard-writing.

11 16.5.2.2.2 Events that cause pointers to become disassociated

12 A pointer becomes disassociated when

- 13 (1) the pointer is nullified (6.3.2),
- 14 (2) the pointer is deallocated (6.3.3),
- 15 (3) the pointer is pointer-assigned (7.4.2) to a disassociated pointer, or
- 16 (4) the pointer is an ultimate component of an object of a type for which default initialization
- 17 is specified for the component, the corresponding initializer is a reference to the intrinsic
- 18 function NULL, and
 - 19 • a procedure is invoked with this object as an actual argument corresponding to a
 - 20 nonpointer nonallocatable dummy argument with INTENT (OUT),
 - 21 • a procedure with this object as an unsaved nonpointer nonallocatable local object
 - 22 that is not accessed by use or host association is invoked, or

1 •this object is allocated.

2 **16.5.2.2.3 Events that cause the association status of pointers to become undefined**

3 The association status of a pointer becomes undefined when

- 4 (1) the pointer is pointer-assigned to a target that has an undefined association status,
- 5 (2) the target of the pointer is deallocated other than through the pointer,
- 6 (3) the allocation transfer procedure (13.7.124) is executed, the pointer is associated with the
7 argument FROM, and an object without the target attribute associated with the argument
8 TO,
- 9 (4) execution of a RETURN or END statement causes the pointer's target to become undefined
10 (item (3) of 16.6.6),
- 11 (5) termination of a BLOCK construct causes the pointer's target to become undefined (item
12 (I16:BLOCK construct terminates) of 16.6.6),
- 13 (6) execution of the host instance of a procedure pointer is completed by execution of a RE-
14 TURN or END statement,
- 15 (7) a procedure is terminated by execution of a RETURN or END statement and the pointer
16 is declared or accessed in the subprogram that defines the procedure unless the pointer
17 (a) has the SAVE attribute,
18 (b) is in blank common,
19 (c) is in a named common block that appears in at least one other scoping unit that is
20 in execution,
21 (d) is in the scoping unit of a module if the module is also referenced by another scoping
22 unit that is in execution,
23 (e) is in the scoping unit of a submodule if any scoping unit in that submodule or any of
24 its descendants is in execution,
25 (f) is accessed by host association, or
26 (g) is the return value of a function declared to have the POINTER attribute,
- 27 (8) a BLOCK construct is terminated and the pointer is an unsaved local entity that is explicitly
28 declared in the BLOCK,
- 29 (9) the pointer is an ultimate component of an object, default initialization is not specified
30 for the component, and a procedure is invoked with this object as an actual argument
31 corresponding to a dummy argument with INTENT(OUT), or
- 32 (10) a procedure is invoked with the pointer as an actual argument corresponding to a pointer
33 dummy argument with INTENT(OUT).

34 **16.5.2.2.4 Other events that change the association status of pointers**

35 When a pointer becomes associated with another pointer by argument association, construct association,
36 or host association, the effects on its association status are specified in 16.5.5.

37 While two pointers are name associated, storage associated, or inheritance associated, if the association
38 status of one pointer changes, the association status of the other changes accordingly.

39 **16.5.2.3 Pointer definition status**

40 The definition status of a pointer is that of its target. If a pointer is associated with a definable target,
41 the definition status of the pointer may be defined or undefined according to the rules for a variable
42 (16.6).

1 16.5.2.4 Relationship between association status and definition status

2 If the association status of a pointer is disassociated or undefined, the pointer shall not be referenced
 3 or deallocated. Whatever its association status, a pointer always may be nullified, allocated, or pointer
 4 assigned. A nullified pointer is disassociated. When a pointer is allocated, it becomes associated but
 5 undefined. When a pointer is pointer assigned, its association and definition status become those of the
 6 specified *data-target* or *proc-target*.

7 16.5.3 Storage association

8 16.5.3.1 General

9 Storage sequences are used to describe relationships that exist among variables, common blocks, and
 10 result variables. **Storage association** is the association of two or more data objects that occurs when
 11 two or more storage sequences share or are aligned with one or more storage units.

12 16.5.3.2 Storage sequence

13 A **storage sequence** is a sequence of storage units. The **size of a storage sequence** is the number
 14 of storage units in the storage sequence. A **storage unit** is a character storage unit, a numeric storage
 15 unit, a file storage unit(9.2.4), or an unspecified storage unit. The sizes of the numeric storage unit, the
 16 character storage unit and the file storage unit are the value of constants in the ISO_FORTRAN_ENV
 17 intrinsic module (13.8.3).

18 In a storage association context

- 19 (1) a nonpointer scalar object of type default integer, default real, default logical, or default
 20 bits occupies a single **numeric storage unit**,
- 21 (2) a nonpointer scalar object of type double precision real or default complex occupies two
 22 contiguous numeric storage units,
- 23 (3) A nonpointer scalar object of type bits with a kind type parameter that is an integer multiple,
 24 N, of the size of a numeric storage unit occupies N contiguous numeric storage units. The
 25 ordering of these consecutive numeric storage units is processor dependent. A nonpointer
 26 scalar object of type bits with a kind type parameter that is not an integer multiple of the
 27 size of a numeric storage unit occupies an unspecified storage unit that is different for each
 28 such kind value.

J3 internal note

Unresolved Technical Issue 075

The “ordering” sentence is trying to say something subtle about endianness, but fails. The added note (some distance below) explains the subtlety, but the text here is still meaningless nonsense. Delete it.

- 29 (4) a nonpointer scalar object of type default character and character length *len* occupies *len*
 30 contiguous **character storage units**,
- 31 (5) a nonpointer scalar object of type character with the C character kind (15.2.2) and character
 32 length *len* occupies *len* contiguous **unspecified storage units**,
- 33 (6) a nonpointer scalar object of sequence type with no type parameters occupies a sequence of
 34 storage sequences corresponding to the sequence of its ultimate components,
- 35 (7) a nonpointer scalar object of any type not specified in items (1)-(6) occupies a single un-
 36 specified storage unit that is different for each case and each set of type parameter values,
 37 and that is different from the unspecified storage units of item (5),
- 38 (8) a nonpointer array occupies a sequence of contiguous storage sequences, one for each array

- 1 element, in array element order (6.2.2.2), and
 2 (9) a pointer occupies a single unspecified storage unit that is different from that of any non-
 3 pointer object and is different for each combination of type, type parameters, and rank. A
 4 pointer that has the CONTIGUOUS attribute occupies a storage unit that is different from
 5 that of a pointer that does not have the CONTIGUOUS attribute.
- 6 A sequence of storage sequences forms a storage sequence. The order of the storage units in such a
 7 composite storage sequence is that of the individual storage units in each of the constituent storage
 8 sequences taken in succession, ignoring any zero-sized constituent sequences.
- 9 Each common block has a storage sequence (5.7.2.2).
- 10 Two nonpointer nonallocatable scalar objects for which the intrinsic function BIT_SIZE returns the same
 11 value occupy the same amount of storage.

J3 internal note**Unresolved Technical Issue 077**

The purpose of the above paragraph here seems to be to promote rampant nonportability. We have always been at liberty to allow this in the past, and we have never done so, because it runs directly counter to our portability goal.

It appears to be attempting to extend storage association; I do not believe for one second that WG5 would have countenanced such a thing.

In any case, it contradicts item (7) above.

Furthermore, it apparently disallows padding in COMMON blocks, and maybe in some structures and arrays. Vendors do different amounts of padding (thus taking up different amounts of storage) for different types and objects right now, and it is highly desirable to allow that to continue.

NOTE 16.13

For a BITS value with a kind equal to an integer multiple of the size of a numeric storage unit, the order of its storage units may depend on the endian convention for the processor's memory. For example, on a system where the size of a numeric storage unit is 32 bits,

```

BITS(32) :: X(2)
BITS(64) :: Y, ZLE, XBE
...
X = TRANSFER(Y, X)
ZBE = X(1)//X(2)
ZLE = X(2)//X(1)

```

!

! On some processors Y==ZLE is true and on other processors Y==ZBE is true.

NOTE 16.14

A nonpointer nonallocatable scalar BITS object with a KIND value that is not an integer multiple of the size of a numeric storage unit in bits might be stored in a memory region larger than the minimum required to represent the value. For example, if BITS_KIND(x) has the value 13, the storage size for x might be 16 bits. Each element of a BITS array occupies the same size memory region as a scalar BITS object of the same kind.

12 **16.5.3.3 Association of storage sequences**

- 1 Two nonzero-sized storage sequences s_1 and s_2 are **storage associated** if the i th storage unit of s_1 is
 2 the same as the j th storage unit of s_2 . This causes the $(i + k)$ th storage unit of s_1 to be the same as
 3 the $(j + k)$ th storage unit of s_2 , for each integer k such that $1 \leq i + k \leq \text{size of } s_1$ and $1 \leq j + k \leq$
 4 $\text{size of } s_2$.
- 5 Storage association also is defined between two zero-sized storage sequences, and between a zero-sized
 6 storage sequence and a storage unit. A zero-sized storage sequence in a sequence of storage sequences is
 7 storage associated with its successor, if any. If the successor is another zero-sized storage sequence, the
 8 two sequences are storage associated. If the successor is a nonzero-sized storage sequence, the zero-sized
 9 sequence is storage associated with the first storage unit of the successor. Two storage units that are
 10 each storage associated with the same zero-sized storage sequence are the same storage unit.

NOTE 16.15

Zero-sized objects may occur in a storage association context as the result of changing a parameter. For example, a program might contain the following declarations:

```

INTEGER, PARAMETER :: PROBSIZE = 10
INTEGER, PARAMETER :: ARRAYSIZE = PROBSIZE * 100
REAL, DIMENSION (ARRAYSIZE) :: X
INTEGER, DIMENSION (ARRAYSIZE) :: IX
...
COMMON / EXAMPLE / A, B, C, X, Y, Z
EQUIVALENCE (X, IX)
...

```

If the first statement is subsequently changed to assign zero to PROBSIZE, the program still will conform to the standard.

11 16.5.3.4 Association of scalar data objects

- 12 Two scalar data objects are storage associated if their storage sequences are storage associated. Two
 13 scalar entities are **totally associated** if they have the same storage sequence. Two scalar entities are
 14 **partially associated** if they are associated without being totally associated.
- 15 The definition status and value of a data object affects the definition status and value of any storage
 16 associated entity. An EQUIVALENCE statement, a COMMON statement, or an ENTRY statement
 17 may cause storage association of storage sequences.
- 18 An EQUIVALENCE statement causes storage association of data objects only within one scoping unit,
 19 unless one of the equivalenced entities is also in a common block (5.7.1.2 and 5.7.2.2).
- 20 COMMON statements cause data objects in one scoping unit to become storage associated with data
 21 objects in another scoping unit.
- 22 A common block is permitted to contain a sequence of differing storage units. All scoping units that
 23 access named common blocks with the same name shall specify an identical sequence of storage units.
 24 Blank common blocks may be declared with differing sizes in different scoping units. For any two blank
 25 common blocks, the initial sequence of storage units of the longer blank common block shall be identical
 26 to the sequence of storage units of the shorter common block. If two blank common blocks are the same
 27 length, they shall have the same sequence of storage units.
- 28 An ENTRY statement in a function subprogram causes storage association of the result variables.

Partial association shall exist only between

- 1 (1) an object of default character or character sequence type and an object of default character
 2 or character sequence type, or
 3 (2) an object of default complex, double precision real, or numeric sequence type and an object
 4 of default integer, default real, default logical, double precision real, default complex, or
 5 numeric sequence type.
- 6 For noncharacter entities, partial association may occur only through the use of COMMON, EQUIV-
 7 ALENCE, or ENTRY statements. For character entities, partial association may occur only through
 8 argument association or the use of COMMON or EQUIVALENCE statements.

NOTE 16.16

In the example:

```
REAL A (4), B
COMPLEX C (2)
DOUBLE PRECISION D
EQUIVALENCE (C (2), A (2), B), (A, D)
```

the third storage unit of C, the second storage unit of A, the storage unit of B, and the second storage unit of D are specified as the same. The storage sequences may be illustrated as:

```
Storage unit      1      2      3      4      5
                  ----C(1)----|---C(2)----
                   A(1)  A(2)  A(3)  A(4)
                   --B--
                   -----D-----
```

A (2) and B are totally associated. The following are partially associated: A (1) and C (1), A (2) and C (2), A (3) and C (2), B and C (2), A (1) and D, A (2) and D, B and D, C (1) and D, and C (2) and D. Although C (1) and C (2) are each storage associated with D, C (1) and C (2) are not storage associated with each other.

- 9 Partial association of character entities occurs when some, but not all, of the storage units of the entities
 10 are the same.

NOTE 16.17

In the example:

```
CHARACTER A*4, B*4, C*3
EQUIVALENCE (A (2:3), B, C)
```

A, B, and C are partially associated.

- 11 A storage unit shall not be explicitly initialized more than once in a program. Explicit initialization
 12 overrides default initialization, and default initialization for an object of derived type overrides default
 13 initialization for a component of the object (4.5.2). Default initialization may be specified for a storage
 14 unit that is storage associated provided the objects supplying the default initialization are of the same
 15 type and type parameters, and supply the same value for the storage unit.

16 16.5.4 Inheritance association

1 Inheritance association occurs between components of the parent component and components inherited
2 by type extension into an extended type (4.5.7.2). This association is persistent; it is not affected by the
3 accessibility of the inherited components.

4 **16.5.5 Establishing associations**

5 When an association is established between two entities by argument association, host association,
6 or construct association, certain characteristics of the **associating entity** become those of the **pre-**
7 **existing** entity.

8 For argument association, if the dummy argument is not a pointer and the corresponding actual argument
9 is a pointer, the pre-existing entity is the target of that pointer. Otherwise the pre-existing entity is the
10 corresponding actual argument. In either case, the associating entity is the dummy argument.

11 For host association, the associating entity is the entity in the host scoping unit and the pre-existing
12 entity is the entity in the contained scoping unit. If an internal procedure is invoked via a dummy
13 procedure or procedure pointer, the pre-existing entity that participates in the association is the one
14 from the host instance. Otherwise, if the host scoping unit is a recursive procedure, the pre-existing entity
15 that participates in the association is the one from the innermost subprogram instance that invoked,
16 directly or indirectly, the contained procedure.

17 For construct association, the associating entity is identified by the associate name and the pre-existing
18 entity is the selector.

19 When an association is established by argument association, host association, or construct association,
20 the following applies.

- 21 • If the associating entity has the POINTER attribute, its pointer association status becomes the
22 same as that of the pre-existing entity. If the pre-existing entity has a pointer association status of
23 associated, the associating entity becomes pointer associated with the same target and, if they are
24 arrays, the bounds of the associating entity become the same as those of the pre-existing entity.
- 25 • If the associating entity has the ALLOCATABLE attribute, its allocation status becomes the same
26 as that of the pre-existing entity. If the pre-existing entity is allocated, the bounds (if it is an array),
27 values of deferred type parameters, definition status, and value (if it is defined) become the same
28 as those of the pre-existing entity. If the associating entity is polymorphic and the pre-existing
29 entity is allocated, the dynamic type of the associating entity becomes the same as that of the
30 pre-existing entity.
- 31 • If the associating entity is neither a pointer nor allocatable, its definition status, value (if it is
32 defined), and dynamic type (if it is polymorphic) become the same as those of the pre-existing
33 entity. If the entities are arrays and the association is not argument association, the bounds of the
34 associating entity become the same as those of the pre-existing entity.

35 **16.6 Definition and undefinition of variables**

36 **16.6.1 Definition of objects and subobjects**

37 A variable may be defined or may be undefined and its definition status may change during execution of
38 a program. An action that causes a variable to become undefined does not imply that the variable was
39 previously defined. An action that causes a variable to become defined does not imply that the variable
40 was previously undefined.

1 Arrays, including sections, and variables of derived, character, or complex type are objects that consist of
 2 zero or more subobjects. Associations may be established between variables and subobjects and between
 3 subobjects of different variables. These subobjects may become defined or undefined.

4 An array is defined if and only if all of its elements are defined.

5 A derived-type scalar object is defined if and only if all of its nonpointer components are defined.

6 A complex or character scalar object is defined if and only if all of its subobjects are defined.

7 If an object is undefined, at least one (but not necessarily all) of its subobjects are undefined.

8 **16.6.2 Variables that are always defined**

9 Zero-sized arrays and zero-length strings are always defined.

10 **16.6.3 Variables that are initially defined**

11 The following variables are initially defined:

- 12 (1) variables specified to have initial values by DATA statements;
- 13 (2) variables specified to have initial values by type declaration statements;
- 14 (3) nonpointer default-initialized subcomponents of variables that do not have the ALLOCAT-
 15 ABLE or POINTER attribute, and are saved, local variables of a main program, or declared
 16 in a MODULE or BLOCK DATA scoping unit;
- 17 (4) pointers specified to be initially associated with a variable that is initially defined;
- 18 (5) variables that are always defined;
- 19 (6) variables with the BIND attribute that are initialized by means other than Fortran.

NOTE 16.18

Fortran code:

```
module mod
  integer, bind(c,name="blivet") :: foo
end module mod
```

C code:

```
int blivet = 123;
```

In the above example, the Fortran variable foo is initially defined to have the value 123 by means other than Fortran.

20 **16.6.4 Variables that are initially undefined**

21 All other variables are initially undefined.

22 **16.6.5 Events that cause variables to become defined**

23 Variables become defined by the following events.

- 24 (1) Execution of an intrinsic assignment statement other than a masked array assignment or
 25 FORALL assignment statement causes the variable that precedes the equals to become

- 1 defined. Execution of a defined assignment statement may cause all or part of the variable
2 that precedes the equals to become defined.
- 3 (2) Execution of a masked array assignment or FORALL assignment statement may cause some
4 or all of the array elements in the assignment statement to become defined (7.4.3).
- 5 (3) As execution of an input statement proceeds, each variable that is assigned a value from
6 the input file becomes defined at the time that data is transferred to it. (See (4) in 16.6.6.)
7 Execution of a WRITE statement whose unit specifier identifies an internal file causes each
8 record that is written to become defined.
- 9 (4) Execution of a DO statement causes the DO variable, if any, to become defined.
- 10 (5) Beginning of execution of the action specified by an *io-implied-do* in a synchronous in-
11 put/output statement causes the *do-variable* to become defined.
- 12 (6) A reference to a procedure causes the entire dummy argument data object to become defined
13 if the dummy argument does not have INTENT(OUT) and the entire corresponding actual
14 argument is defined.
15 A reference to a procedure causes a subobject of a dummy argument to become defined if
16 the dummy argument does not have INTENT(OUT) and the corresponding subobject of
17 the corresponding actual argument is defined.
- 18 (7) Execution of an input/output statement containing an IOSTAT= specifier causes the spec-
19 ified integer variable to become defined.
- 20 (8) Execution of a synchronous READ statement containing a SIZE= specifier causes the spec-
21 ified integer variable to become defined.
- 22 (9) Execution of a wait operation corresponding to an asynchronous input statement containing
23 a SIZE= specifier causes the specified integer variable to become defined.
- 24 (10) Execution of an INQUIRE statement causes any variable that is assigned a value during the
25 execution of the statement to become defined if no error condition exists.
- 26 (11) If an error, end-of-file, or end-of-record condition occurs during execution of an input/output
27 statement that has an IOMSG= specifier, the *iormsg-variable* becomes defined.
- 28 (12) When a character storage unit becomes defined, all associated character storage units be-
29 come defined.
30 When a numeric storage unit becomes defined, all associated numeric storage units of the
31 same type or that are bits compatible become defined. When an entity of double precision
32 real type becomes defined, all totally associated entities of double precision real type become
33 defined.
34 When an unspecified storage unit becomes defined, all associated unspecified storage units
35 become defined.
- 36 (13) When a default complex entity becomes defined, all partially associated default real entities
37 become defined.
- 38 (14) When both parts of a default complex entity become defined as a result of partially associ-
39 ated default real or default complex entities becoming defined, the default complex entity
40 becomes defined.
- 41 (15) When all components of a structure of a numeric sequence type or character sequence type
42 become defined as a result of partially associated objects becoming defined, the structure
43 becomes defined.
- 44 (16) Execution of a statement with a STAT= specifier causes the variable specified by the STAT=
45 specifier to become defined.
- 46 (17) If an error condition occurs during execution of a statement that has an ERRMSG= specifier,
47 the variable specified by the ERRMSG= specifier becomes defined.
- 48 (18) Allocation of a zero-sized array causes the array to become defined.
- 49 (19) Allocation of an object that has a nonpointer default-initialized subcomponent causes that
50 subcomponent to become defined.

- 1 (20) Invocation of a procedure causes any automatic object of zero size in that procedure to
2 become defined.
- 3 (21) Execution of a pointer assignment statement that associates a pointer with a target that is
4 defined causes the pointer to become defined.
- 5 (22) Invocation of a procedure that contains an unsaved nonpointer nonallocatable local variable
6 causes all nonpointer default-initialized subcomponents of the object to become defined.
- 7 (23) Invocation of a procedure that has a nonpointer nonallocatable INTENT (OUT) dummy
8 argument causes all nonpointer default-initialized subcomponents of the dummy argument
9 to become defined.
- 10 (24) Invocation of a nonpointer function of a derived type causes all nonpointer default-initialized
11 subcomponents of the function result to become defined.
- 12 (25) In a FORALL or DO CONCURRENT construct, the *index-name* becomes defined when
13 the *index-name* value set is evaluated.
- 14 (26) An object with the VOLATILE attribute that is changed by a means not specified by the
15 program becomes defined (see 5.3.18).
- 16 (27) Execution of the BLOCK statement of a BLOCK construct that has an unsaved nonpointer
17 nonallocatable local variable causes all nonpointer default-initialized subcomponents of the
18 variable to become defined.
- 19 (28) Execution of an OPEN statement containing a NEWUNIT= specifier causes the specified
20 integer variable to become defined.

21 16.6.6 Events that cause variables to become undefined

22 Variables become undefined by the following events.

- 23 (1) When a variable of a given type becomes defined, all associated variables of different type
24 become undefined. However, when a variable of type default real is partially associated with
25 a variable of type default complex, the complex variable does not become undefined when
26 the real variable becomes defined and the real variable does not become undefined when
27 the complex variable becomes defined. When a variable of type default complex is partially
28 associated with another variable of type default complex, definition of one does not cause
29 the other to become undefined. When a dummy argument of type bits is associated with
30 an actual argument of a different type, definition of the dummy argument does not cause
31 the actual argument to become undefined.
- 32 (2) If the evaluation of a function may cause a variable to become defined and if a reference to
33 the function appears in an expression in which the value of the function is not needed to
34 determine the value of the expression, the variable becomes undefined when the expression
35 is evaluated.
- 36 (3) When execution of an instance of a subprogram completes,
37 (a) its unsaved local variables become undefined,
38 (b) unsaved variables in a named common block that appears in the subprogram become
39 undefined if they have been defined or redefined, unless another active scoping unit
40 is referencing the common block,
41 (c) unsaved nonfinalizable local variables of a module or submodule become undefined
42 unless another active scoping unit is referencing the module or submodule, and

NOTE 16.19

A module subprogram inherently references the module that is its host. Therefore, for processors that keep track of when modules are in use, a module is in use whenever any procedure in the module is active, even if no other active scoping units reference the module; this situation can arise if a module procedure is invoked via a procedure pointer, a type-bound procedure, or a companion processor.

- 1 (d) unsaved finalizable local variables of a module or submodule may be finalized if no
2 other active scoping unit is referencing the module or submodule – following which
3 they become undefined.
- 4 (4) When an error condition or end-of-file condition occurs during execution of an input state-
5 ment, all of the variables specified by the input list or namelist group of the statement
6 become undefined.
- 7 (5) When an error condition, end-of-file condition, or end-of-record condition occurs during
8 execution of an input/output statement and the statement contains any *io-implied-dos*, all
9 of the *do-variables* in the statement become undefined (9.10).
- 10 (6) Execution of a direct access input statement that specifies a record that has not been written
11 previously causes all of the variables specified by the input list of the statement to become
12 undefined.
- 13 (7) Execution of an INQUIRE statement may cause the NAME=, RECL=, and NEXTREC=
14 variables to become undefined (9.9).
- 15 (8) When a character storage unit becomes undefined, all associated character storage units
16 become undefined.
- 17 When a numeric storage unit becomes undefined, all associated numeric storage units be-
18 come undefined unless the undefinition is a result of defining an associated numeric storage
19 unit of different type (see (1) above).
- 20 When an entity of double precision real type becomes undefined, all totally associated
21 entities of double precision real type become undefined.
- 22 When an unspecified storage unit becomes undefined, all associated unspecified storage units
23 become undefined.
- 24 (9) When an allocatable entity is deallocated, it becomes undefined.
- 25 (10) When the allocation transfer procedure (13.5.17) causes the allocation status of an allocat-
26 able entity to become unallocated, the entity becomes undefined.
- 27 (11) Successful execution of an ALLOCATE statement for a nonzero-sized object that has a sub-
28 component for which default initialization has not been specified causes the subcomponent
29 to become undefined.
- 30 (12) Execution of an INQUIRE statement causes all inquiry specifier variables to become un-
31 defined if an error condition exists, except for any variable in an IOSTAT= or IOMSG=
32 specifier.
- 33 (13) When a procedure is invoked
- 34 (a) an optional dummy argument that is not associated with an actual argument becomes
35 undefined,
- 36 (b) a dummy argument with INTENT (OUT) becomes undefined except for any non-
37 pointer default-initialized subcomponents of the argument,
- 38 (c) an actual argument associated with a dummy argument with INTENT (OUT) be-
39 comes undefined except for any nonpointer default-initialized subcomponents of the
40 argument,
- 41 (d) a subobject of a dummy argument that does not have INTENT (OUT) becomes
42 undefined if the corresponding subobject of the actual argument is undefined, and
- 43 (e) the result variable of a function becomes undefined except for any of its nonpointer
44 default-initialized subcomponents.
- 45 (14) When the association status of a pointer becomes undefined or disassociated (16.5.2.2.2-
46 16.5.2.2.3), the pointer becomes undefined.
- 47 (15) When the execution of a FORALL or DO CONCURRENT construct completes, the *index-*
48 *names* become undefined.
- 49 (16) When a DO CONCURRENT construct terminates, a variable that is defined or becomes
50 undefined during more than one iteration of the construct becomes undefined.

- 1 (17) Execution of an asynchronous READ statement causes all of the variables specified by the
 2 input list or SIZE= specifier to become undefined. Execution of an asynchronous namelist
 3 READ statement causes any variable in the namelist group to become undefined if that
 4 variable will subsequently be defined during the execution of the READ statement or the
 5 corresponding WAIT operation.
- 6 (18) When execution of a RETURN or END statement causes a variable to become undefined,
 7 any variable of type C_PTR becomes undefined if its value is the C address of any part of
 8 the variable that becomes undefined.
- 9 (19) When a variable with the TARGET attribute is deallocated, any variable of type C_PTR
 10 becomes undefined if its value is the C address of any part of the variable that is deallocated.
- 11 (20) When a BLOCK construct terminates, its unsaved local variables become undefined.

NOTE 16.20

Execution of a defined assignment statement may leave all or part of the variable undefined.

12 **16.6.7 Variable definition context**

13 Some variables are prohibited from appearing in a syntactic context that would imply definition or
 14 undefinition of the variable (5.3.9, 5.3.14, 12.7). The following are the contexts in which the appearance
 15 of a variable implies such definition or undefinition of the variable:

- 16 • the *variable* of an *assignment-stmt*;
- 17 • a *pointer-object* in a *nullify-stmt*;
- 18 • a *data-pointer-object* or *proc-pointer-object* in a *pointer-assignment-stmt*;
- 19 • a *do-variable* in a *do-stmt* or *io-implied-do*;
- 20 • an *input-item* in a *read-stmt*;
- 21 • a *variable-name* in a *namelist-stmt* if the *namelist-group-name* appears in a NML= specifier in a
 22 *read-stmt*;
- 23 • an *internal-file-variable* in a *write-stmt*;
- 24 • an IOSTAT=, SIZE=, or IOMSG= specifier in an input/output statement;
- 25 • a specifier in an INQUIRE statement other than FILE=, ID=, and UNIT=;
- 26 • a NEWUNIT= specifier in an OPEN statement;
- 27 • a READY= specifier in a QUERY statement;
- 28 • a *stat-variable*, *allocate-object*, or *errmsg-variable*;
- 29 • an actual argument in a reference to a procedure with an explicit interface if the associated dummy
 30 argument has the INTENT(OUT) or INTENT(INOUT) attribute;
- 31 • a *variable* that is the *selector* in a SELECT TYPE or ASSOCIATE construct if the associate name
 32 of that construct appears in a variable definition context.

1 **16.6.8 Pointer association context**

2 Some pointers are prohibited from appearing in a syntactic context that would imply alteration of the
3 pointer association status (16.5.2.2, 5.3.9, 5.3.14). The following are the contexts in which the appearance
4 of a pointer implies such alteration of its pointer association status:

- 5 • a *pointer-object* in a *nullify-stmt*;
- 6 • a *data-pointer-object* or *proc-pointer-object* in a *pointer-assignment-stmt*;
- 7 • an *allocate-object* in an *allocate-stmt* or *deallocate-stmt*;
- 8 • an actual argument in a reference to a procedure if the associated dummy argument is a pointer
9 with the INTENT (OUT) or INTENT (INOUT) attribute.

10 If a reference to a function appears in a variable definition context the result of the function reference
11 shall be a pointer that is associated with a definable target. That target is the variable that becomes
12 defined or undefined.

Annex A

(Informative)

Glossary of technical terms

The following is a list of the principal technical terms used in the standard and their definitions. A reference in parentheses immediately after a term is to the clause or subclause where the term is defined or explained. The wording of a definition here is not necessarily the same as in the standard.

abstract type (4.5.7) : A type that has the ABSTRACT attribute. A nonpolymorphic object shall not be declared to be of abstract type. A polymorphic object shall not be constructed or allocated to have a dynamic abstract type.

actual argument (12, 12.5.2) : An expression, a variable, a procedure, or an alternate return specifier that is specified in a procedure reference.

allocatable variable (5.3.3) : A variable having the ALLOCATABLE attribute. It may be referenced or defined only when it is allocated. If it is an array, it has a shape only when it is allocated. It may be a named variable or a structure component.

argument (12) : An actual argument or a dummy argument.

argument association (16.5.1.2) : The relationship between an actual argument and a dummy argument during the execution of a procedure reference.

array (2.4.5) : A set of scalar data, all of the same type and type parameters, whose individual elements are arranged in a rectangular pattern. It may be a named array, an array section, a structure component, a function value, or an expression. Its rank is at least one. Note that in FORTRAN 77, arrays were always named and never constants.

array element (2.4.5, 6.2.2) : One of the scalar data that make up an array that is either named or is a structure component.

array pointer (5.3.7.4) : A pointer to an array.

array section (2.4.5, 6.2.2.3) : A subobject that is an array and is not a structure component.

assignment statement (7.4.1.1) : A statement that evaluates an expression and assigns its value to a variable.

associate name (8.1.3.1) : The name of the construct entity with which a selector of a SELECT TYPE or ASSOCIATE construct is associated within the construct.

association (16.5) : Name association, pointer association, storage association, or inheritance association.

assumed-shape array (5.3.7.3) : A nonpointer dummy array that takes its shape from the associated actual argument.

assumed-size array (5.3.7.5) : A dummy array whose size is assumed from the associated actual argument. Its last upper bound is specified by an asterisk.

attribute (5) : A property of an entity that determines its uses.

- 1 **automatic data object** (5.2) : A data object that is a local entity of a subprogram, that is not a dummy
2 argument, and that has a length type parameter or array bound that is specified by an expression that
3 is not an initialization expression.
- 4 **base type** (4.5.7) : An extensible type that is not an extension of another type.
- 5 **belong** (8.1.7.5.3, 8.1.7.5.4) : If an EXIT or a CYCLE statement contains a construct name, the
6 statement **belongs** to the construct using that name. Otherwise, it **belongs** to the innermost DO
7 construct in which it appears.
- 8 **binding label** (15.5.2, 15.4.2) : A value of type default character that uniquely identifies how a variable,
9 common block, subroutine, or function is known to a companion processor.
- 10 **bits compatible** (12.5.2.4) : An entity is bits compatible with another entity if and only if one is of
11 type bits, the other is of type bits, integer, real, complex, or logical, and scalar entities of these types
12 have the same size expressed in bits.
- 13 **block** (8.1) : A sequence of executable constructs embedded in another executable construct, bounded
14 by statements that are particular to the construct, and treated as an integral unit.
- 15 **block data program unit** (11.3) : A program unit that provides initial values for data objects in
16 named common blocks.
- 17 **bounds** (5.3.7.2) : For a named array, the limits within which the values of the subscripts of its array
18 elements shall lie.
- 19 **character length parameter** (2.4.1.1) : The type parameter that specifies the number of characters
20 for an entity of type character.
- 21 **character storage unit** (16.5.3.2) : The unit of storage for holding a scalar that is not a pointer and
22 is of type default character and character length one.
- 23 **character string** (4.4.5) : A sequence of characters numbered from left to right 1, 2, 3, ...
- 24 **characteristics** (12.3) :
- 25 Of a procedure, its classification as a function or subroutine, whether it is pure, whether it is elemental,
26 whether it has the BIND attribute, the value of its binding label, the characteristics of its dummy
27 arguments, and the characteristics of its function result if it is a function.
- 28 Of a dummy argument, whether it is a data object, is a procedure, is a procedure pointer, is an asterisk
29 (alternate return indicator), or has the OPTIONAL attribute.
- 30 Of a dummy data object, its type, type parameters, shape, the exact dependence of an array bound
31 or type parameter on other entities, intent, whether it is optional, whether it is a pointer or a target,
32 whether it is allocatable, whether it has the VALUE, ASYNCHRONOUS, or VOLATILE attributes,
33 whether it is polymorphic, and whether the shape, size, or a type parameter is assumed.
- 34 Of a dummy procedure or procedure pointer, whether the interface is explicit, the characteristics of the
35 procedure if the interface is explicit, and whether it is optional.
- 36 Of a function result, its type, type parameters, which type parameters are deferred, whether it is poly-
37 morphic, whether it is a pointer or allocatable, whether it is a procedure pointer, rank if it is a pointer
38 or allocatable, shape if it is not a pointer or allocatable, the exact dependence of an array bound or type
39 parameter on other entities, and whether the character length is assumed.
- 40 **class** (4.3.1.3) : A class named N is the set of types extended from the type named N.

- 1 **co-array** (2.4.6) A data entity that has nonzero co-rank. A non-dummy co-array has the same shape
2 on every image and its values may be referenced or defined by any image.
- 3 **co-bounds** (5.3.7) : For a co-array, the limits within which the values of the co-subscripts of its co-
4 indexed objects shall lie.
- 5 **co-indexed object** (2.4.6) : An object whose designator includes an *image-selector*.
- 6 **co-rank** (2.4.6) : The number of co-dimensions of a co-array.
- 7 **co-subscript** (6.2.3) : One of the list of scalar integer expressions in an *image-selector*.
- 8 **collating sequence** (4.4.5.4) : An ordering of all the different characters of a particular kind type
9 parameter.
- 10 **collective subroutine** (13.1) : An intrinsic subroutine that has an argument of type IMAGE_TEAM.
- 11 **common block** (5.7.2) : A block of physical storage that may be accessed by any of the scoping units
12 in a program.
- 13 **companion processor** (2.5.10): A mechanism by which global data and procedures may be referenced
14 or defined. It may be a mechanism that references and defines such entities by means other than Fortran.
15 The procedures can be described by a C function prototype.
- 16 **component** (4.5) : A constituent of a derived type.
- 17 **component order** (4.5.4.6) : The ordering of the components of a derived type that is used for intrinsic
18 formatted input/output and for structure constructors.
- 19 **conformable** (2.4.5) : Two arrays are **conformable** if they have the same shape. A scalar is con-
20 formable with any array.
- 21 **conformance** (1.5) : A program conforms to the standard if it uses only those forms and relationships
22 described therein and if the program has an interpretation according to the standard. A program unit
23 conforms to the standard if it can be included in a program in a manner that allows the program to be
24 standard-conforming. A processor conforms to the standard if it executes standard-conforming programs
25 in a manner that fulfills the interpretations prescribed in the standard and contains the capability of
26 detection and reporting as listed in 1.5.
- 27 **connect team** (9.4.5.17) : The set of images that are permitted to reference a particular external
28 input/output unit.
- 29 **connected** (9.4.3) :
- 30 For an external unit, the property of referring to an external file.
- 31 For an external file, the property of being referred to by an external unit.
- 32 **constant** (2.4.3.1.2) : A data object whose value shall not change during execution of a program. It
33 may be a named constant or a literal constant.
- 34 **construct** (7.4.3, 7.4.4, 8.1) : A sequence of statements starting with an ASSOCIATE, BLOCK, DO,
35 FORALL, IF, SELECT CASE, SELECT TYPE, or WHERE statement and ending with the correspond-
36 ing terminal statement.
- 37 **construct association** (16.5.1.6) : The association between the selector of an ASSOCIATE or SELECT
38 TYPE construct and the associate name.
- 39 **construct entity** (16) : An entity defined by a lexical token whose scope is a construct.

- 1 **control mask** (7.4.3) : In a WHERE statement or construct, an array of type logical whose value
2 determines which elements of an array, in a *where-assignment-stmt*, will be defined.
- 3 **data** : Plural of datum.
- 4 **data entity** (2.4.3) : A data object, the result of the evaluation of an expression, or the result of the
5 execution of a function reference (called the function result). A data entity has a type (either intrinsic
6 or derived) and has, or may have, a data value (the exception is an undefined variable). Every data
7 entity has a rank and is thus either a scalar or an array.
- 8 **data object** (2.4.3.1) : A data entity that is a constant, a variable, or a subobject of a constant.
- 9 **data type** (4) : See type.
- 10 **datum** : A single quantity that may have any of the set of values specified for its type.
- 11 **decimal symbol** (9.9.2.6, 10.6, 10.8.8) : The character that separates the whole and fractional parts in
12 the decimal representation of a real number in a file. By default the decimal symbol is a decimal point
13 (also known as a period). The current decimal symbol is determined by the current decimal edit mode.
- 14 **declared type** (4.3.1.3, 7.1.4) : The type that a data entity is declared to have. May differ from the
15 type during execution (the dynamic type) for polymorphic data entities.
- 16 **default initialization** (4.5) : If initialization is specified in a type definition, an object of the type will
17 be automatically initialized. Nonpointer components may be initialized with values by default; pointer
18 components may be initially disassociated by default. Default initialization is not provided for objects
19 of intrinsic type.
- 20 **default-initialized** (4.5.4.5) : A subcomponent is said to be default-initialized if it will be initialized
21 by default initialization.
- 22 **deferred binding** (4.5.5) : A binding with the DEFERRED attribute. A deferred binding shall appear
23 only in an abstract type definition (4.5.7).
- 24 **deferred type parameter** (4.3) : A length type parameter whose value is not specified in the decla-
25 ration of an object, but instead is specified when the object is allocated or pointer-assigned.
- 26 **definable** (2.5.5) : A variable is **definable** if its value may be changed by the appearance of its
27 designator on the left of an assignment statement. An allocatable variable that has not been allocated
28 is an example of a data object that is not definable. An example of a subobject that is not definable is
29 C (I) when C is an array that is a constant and I is an integer variable.
- 30 **defined** (2.5.5) : For a data object, the property of having or being given a valid value.
- 31 **defined assignment statement** (7.4.1.4, 12.4.3.3.2) : An assignment statement that is not an intrinsic
32 assignment statement; it is defined by a subroutine and a generic interface that specifies ASSIGNMENT
33 (=).
- 34 **defined operation** (7.1.3, 12.4.3.3.1) : An operation that is not an intrinsic operation and is defined
35 by a function that is associated with a generic identifier.
- 36 **deleted feature** (1.8) : A feature in a previous Fortran standard that is considered to have been
37 redundant and largely unused. See B.1 for a list of features that are in a previous Fortran standard, but
38 are not in this standard. A feature designated as an obsolescent feature in the standard may become a
39 deleted feature in the next revision.
- 40 **derived type** (2.4.1.2, 4.5) : A type whose data have components, each of which is either of intrinsic
41 type or of another derived type.

- 1 **designator** (2.5.1) : A name, followed by zero or more subobject selectors.
- 2 **disassociated** (2.4.7) : A disassociated pointer is not associated with a target. A pointer is disassociated
3 following execution of a NULLIFY statement, following pointer assignment with a disassociated pointer,
4 by default initialization, or by explicit initialization. A data pointer may also be disassociated by
5 execution of a DEALLOCATE statement.
- 6 **dummy argument** (12, 12.6.2.1, 12.6.2.2, 12.6.2.5, 12.6.4) : An entity by which an associated actual
7 argument is accessed during execution of a procedure.
- 8 **dummy array** : A dummy argument that is an array.
- 9 **dummy data object** (12.3.2.2, 12.5.2.5-12.5.2.8) : A dummy argument that is a data object.
- 10 **dummy procedure** (12.2.2.3) : A dummy argument that is specified or referenced as a procedure.
- 11 **dynamic type** (4.3.1.3, 7.1.4) : The type of a data entity during execution of a program. The dynamic
12 type of a data entity that is not polymorphic is the same as its declared type.
- 13 **effective item** (9.5.3) : A scalar object resulting from expanding an input/output list according to the
14 rules in 9.5.3.
- 15 **elemental** (2.4.5, 7.4.1.4, 12.8) : An adjective applied to an operation, procedure, or assignment state-
16 ment that is applied independently to elements of an array or corresponding elements of a set of con-
17 formable arrays and scalars.
- 18 **entity** : The term used for any of the following: an abstract interface, common block, construct, data
19 entity, external unit, generic interface, image, macro, namelist group, operator, procedure, program unit,
20 statement function, statement label, or type.
- 21 **executable construct** (2.1) : An action statement (R214) or an ASSOCIATE, CASE, DO, FORALL,
22 IF, SELECT TYPE, or WHERE construct.
- 23 **executable statement** (2.3.3) : An instruction to perform or control one or more computational
24 actions.
- 25 **explicit initialization** (5.2) : Explicit initialization may be specified for objects of intrinsic or derived
26 type in type declaration statements or DATA statements. An object of a derived type that specifies
27 default initialization shall not appear in a DATA statement.
- 28 **explicit interface** (12.4.2) : If a procedure has an explicit interface at the point of a reference to it, the
29 processor is able to verify that the characteristics of the reference and declaration are related as required
30 by this standard. A procedure has an explicit interface if it is an internal procedure, a module procedure,
31 an intrinsic procedure, an external procedure that has an interface body, a procedure reference in its
32 own scoping unit, or a dummy procedure that has an interface body.
- 33 **explicit-shape array** (5.3.7.2) : A named array that is declared with explicit bounds.
- 34 **expression** (2.4.3.2, 7.1) : A sequence of operands, operators, and parentheses (R722). It may be a
35 variable, a constant, a function reference, or may represent a computation.
- 36 **extended type** (4.5.7) : An extensible type that is an extension of another type. A type that is declared
37 with the EXTENDS attribute.
- 38 **extensible type** (4.5.7) : A type from which new types may be derived using the EXTENDS attribute.
39 A nonsequence type that does not have the BIND attribute.
- 40 **extension type** (4.5.7) : A base type is an extension type of itself only. An extended type is an

- 1 extension type of itself and of all types for which its parent type is an extension.
- 2 **extent** (2.4.5) : The size of one dimension of an array.
- 3 **external file** (9.2) : A sequence of records that exists in a medium external to the program.
- 4 **external linkage** : The characteristic describing that a C entity is global to the program; defined in
5 subclause 6.2.2 of the C International Standard.
- 6 **external procedure** (2.2.4.1) : A procedure that is defined by an external subprogram or by a means
7 other than Fortran.
- 8 **external subprogram** (2.2) : A subprogram that is not in a main program, module, or another
9 subprogram. Note that a module is not called a subprogram. Note that in FORTRAN 77, a block data
10 program unit is called a subprogram.
- 11 **external unit** (9.4) : A mechanism that is used to refer to an external file. It is identified by a
12 nonnegative integer.
- 13 **file** (9) : An internal file or an external file.
- 14 **file storage unit** (9.2.4) : The unit of storage for an unformatted or stream file.
- 15 **final subroutine** (4.5.6) : A subroutine that is called automatically by the processor during finalization.
- 16 **finalizable** (4.5.6) : A type that has final subroutines, or that has a finalizable component. An object
17 of finalizable type.
- 18 **finalization** (4.5.6.2) : The process of calling user-defined final subroutines immediately before destroy-
19 ing an object.
- 20 **function** (2.2.4) : A procedure that is invoked in an expression and computes a value which is then
21 used in evaluating the expression.
- 22 **function result** (12.6.2.1) : The data object that returns the value of a function.
- 23 **function subprogram** (12.6.2.1) : A sequence of statements beginning with a FUNCTION statement
24 that is not in an interface block and ending with the corresponding END statement.
- 25 **generic identifier** (12.4.3.2) : A lexical token that appears in an INTERFACE statement and is
26 associated with all the procedures in the interface block or that appears in a GENERIC statement and
27 is associated with the specific type-bound procedures.
- 28 **generic interface** (4.5.2, 12.4.3.2) : An interface specified by a generic procedure binding or a generic
29 interface block.
- 30 **generic interface block** (12.4.3.2) : An interface block with a generic specification.
- 31 **global entity** (16.2) : An entity with an identifier whose scope is a program.
- 32 **host** (2.2.1) : Host scoping unit.
- 33 **host association** (16.5.1.4) : The process by which a contained scoping unit accesses entities of its
34 host.
- 35 **host scoping unit** (2.2.1) : A scoping unit that immediately surrounds another scoping unit.
- 36 **image** (2.3.2) : A Fortran program executes as if it were replicated a number of times, the number of
37 replications remaining fixed during execution of the program. Each copy is called an **image** and each

- 1 image executes asynchronously.
- 2 **image index** (2.3.2) : An integer value that identifies an image.
- 3 **implicit interface** (12.4.2) : For a procedure referenced in a scoping unit, the property of not having
4 an explicit interface. A statement function always has an implicit interface
- 5 **inherit** (4.5.7) : To acquire from a parent. Type parameters, components, or procedure bindings of an
6 extended type that are automatically acquired from its parent type without explicit declaration in the
7 extended type are said to be inherited.
- 8 **inheritance association** (4.5.7.2, 16.5.4) : The relationship between the inherited components and the
9 parent component in an extended type.
- 10 **inquiry function** (13.1) : A function that is either intrinsic or is defined in an intrinsic module and
11 whose result depends on properties of one or more of its arguments instead of their values.
- 12 **instance of a subprogram** (12.6.2.3) : The copy of a subprogram that is created when a procedure
13 defined by the subprogram is invoked.
- 14 **intent** (5.3.9) : An attribute of a dummy data object that indicates whether it is used to transfer data
15 into the procedure, out of the procedure, or both.
- 16 **interface block** (12.4.3.2) : A sequence of statements from an INTERFACE statement to the corre-
17 sponding END INTERFACE statement.
- 18 **interface body** (12.4.3.2) : A sequence of statements in an interface block from a FUNCTION or
19 SUBROUTINE statement to the corresponding END statement.
- 20 **internal file** (9.3) : A character variable that is used to transfer and convert data from internal storage
21 to internal storage.
- 22 **internal procedure** (2.2.4.3) : A procedure that is defined by an internal subprogram.
- 23 **internal subprogram** (2.2) : A subprogram in a main program or another subprogram.
- 24 **interoperable** (15.3) : The property of a Fortran entity that ensures that an equivalent entity may be
25 defined by means of C.
- 26 **intrinsic** (2.5.7) : An adjective that may be applied to types, operators, assignment statements, proce-
27 dures, and modules. Intrinsic types, operators, and assignment statements are defined in this standard
28 and may be used in any scoping unit without further definition or specification. Intrinsic procedures are
29 defined in this standard or provided by a processor, and may be used in a scoping unit without further
30 definition or specification. Intrinsic modules are defined in this standard or provided by a processor,
31 and may be accessed by use association; procedures and types defined in an intrinsic module are not
32 themselves intrinsic.
- 33 Intrinsic procedures and modules that are not defined in this standard are called nonstandard intrinsic
34 procedures and modules.
- 35 **invoke** (2.2.4) :
- 36 To call a subroutine by a CALL statement or by a defined assignment statement.
- 37 To call a function by a reference to it by name or operator during the evaluation of an expression.
- 38 To call a final subroutine by finalization.
- 39 **keyword** (2.5.2) : A word that is part of the syntax of a statement or a name that is used to identify

- 1 an item in a list.
- 2 **kind type parameter** (2.4.1.1, 4.4.2, 4.4.3, 4.4.4, 4.4.5, 4.4.6, 4.5.3) : A parameter whose values label
3 the available kinds of an intrinsic type, or a derived-type parameter that is declared to have the KIND
4 attribute.
- 5 **label** : See binding label or statement label.
- 6 **length of a character string** (4.4.5) : The number of characters in the character string.
- 7 **lexical token** (3.2) : A sequence of one or more characters with a specified interpretation.
- 8 **line** (3.3) : A sequence of 0 to 132 characters, which may contain Fortran statements, a comment, or
9 an INCLUDE line.
- 10 **linkage association** (16.5.1.5) : The association between interoperable Fortran entities and their C
11 counterparts.
- 12 **literal constant** (2.4.3.1.2, 4.4) : A constant without a name. Note that in FORTRAN 77, this was
13 called simply a constant.
- 14 **local entity** (16.3) : An entity identified by a lexical token whose scope is a scoping unit.
- 15 **local variable** (2.4.3.1.1) : A variable local to a particular scoping unit; not imported through use or
16 host association, not a dummy argument, and not a variable in common.
- 17 **main program** (2.3.6, 11.1) : A Fortran main program or a replacement defined by means other than
18 Fortran.
- 19 **many-one array section** (6.2.2.3.2) : An array section with a vector subscript having two or more
20 elements with the same value.
- 21 **module** (2.2.5, 11.2) : A program unit that contains or accesses definitions to be accessed by other
22 program units.
- 23 **module procedure** (2.2.4.2) : A procedure that is defined by a module subprogram.

J3 internal note

Unresolved Technical Issue 5003

The following glossary entry is just plain wrong. A Module Procedure Interface Block declares a Module Procedure Interface, but not a Separate Module Procedure's Interface unless the Separate Module Procedure is defined by an Separate Module Subprogram (sic).

This raises the question of whether the whole definition/declaration stuff in C12 is consistent and correct. It looks like it is probably inconsistent, since one cannot know when translating the module whether the MPIB should be treated as declaring the SMP's interface or not. When C12 has been fixed, the glossary entry can be reconsidered – and deleted if it is not useful.

- 24 **module procedure interface** (12.4.3.2) : The interface for a separate module procedure.
- 25 **module subprogram** (2.2) : A subprogram that is in a module but is not an internal subprogram.
- 26 **name** (3.2.1) : A lexical token consisting of a letter followed by up to 62 alphanumeric characters
27 (letters, digits, and underscores). Note that in FORTRAN 77, this was called a symbolic name.
- 28 **name association** (16.5.1) : Argument association, use association, host association, linkage associa-
29 tion, or construct association.
- 30 **named** : Having a name. That is, in a phrase such as “named variable,” the word “named” signifies that

- 1 the variable name is not qualified by a subscript list, substring specification, and so on. For example,
2 if X is an array variable, the reference “X” is a named variable while the reference “X(1)” is an object
3 designator.
- 4 **named constant** (2.4.3.1.2) : A constant that has a name. Note that in FORTRAN 77, this was called
5 a symbolic constant.
- 6 **NaN** (14.7) : A Not-a-Number value of IEEE arithmetic. It represents an undefined value or a value
7 created by an invalid operation.
- 8 **nonexecutable statement** (2.3.3) : A statement used to configure the program environment in which
9 computational actions take place.
- 10 **numeric storage unit** (16.5.3.2) : The unit of storage for holding a scalar that is not a pointer and is
11 of type default real, default integer, or default logical.
- 12 **numeric type** (4.4) : Integer, real or complex type.
- 13 **object** (2.4.3.1) : Data object.
- 14 **object designator** (2.5.1) : A designator for a data object.
- 15 **obsolescent feature** (1.8) : A feature that is considered to have been redundant but that is still in
16 frequent use.
- 17 **operand** (2.5.8) : An expression that precedes or succeeds an operator.
- 18 **operation** (7.1.2) : A computation involving one or two operands.
- 19 **operator** (2.5.8) : A lexical token that specifies an operation.
- 20 **override** (4.5.2, 4.5.7) : When explicit initialization or default initialization overrides default initializa-
21 tion, it is as if only the overriding initialization were specified. If a procedure is bound to an extensible
22 type, it overrides the one that would have been inherited from the parent type.
- 23 **parent component** (4.5.7.2) : The component of an entity of extended type that corresponds to its
24 inherited portion.
- 25 **parent type** (4.5.7) : The extensible type from which an extended type is derived.
- 26 **passed-object dummy argument** (4.5.4.4) : The dummy argument of a type-bound procedure or
27 procedure pointer component that becomes associated with the object through which the procedure was
28 invoked.
- 29 **pointer** (2.4.7) : An entity that has the POINTER attribute.
- 30 **pointer assignment** (7.4.2) : The pointer association of a pointer with a target by the execution of a
31 pointer assignment statement or the execution of an assignment statement for a data object of derived
32 type having the pointer as a subobject.
- 33 **pointer assignment statement** (7.4.2) : A statement of the form “pointer-object => target”.
- 34 **pointer associated** (6.3, 7.4.2) : The relationship between a pointer and a target following a pointer
35 assignment or a valid execution of an ALLOCATE statement.
- 36 **pointer association** (16.5.2) : The process by which a pointer becomes pointer associated with a target.
- 37 **polymorphic** (4.3.1.3) : Able to be of differing types during program execution. An object declared
38 with the CLASS keyword is polymorphic.

- 1 **preconnected** (9.4.4) : A property describing a unit that is connected to an external file at the beginning
2 of execution of a program. Such a unit may be specified in input/output statements without an OPEN
3 statement being executed for that unit.
- 4 **procedure** (2.2.4, 12.2) : A computation that may be invoked during program execution. It may be a
5 function or a subroutine. It may be an intrinsic procedure, an external procedure, a module procedure,
6 an internal procedure, a dummy procedure, or a statement function. A subprogram may define more than
7 one procedure if it contains ENTRY statements.
- 8 **procedure designator** (2.5.1) : A designator for a procedure.
- 9 **procedure interface** (12.4) : The characteristics of a procedure, the name of the procedure, the name
10 of each dummy argument, and the generic identifiers (if any) by which it may be referenced.
- 11 **processor** (1.2) : The combination of a computing system and the mechanism by which programs are
12 transformed for use on that computing system.
- 13 **processor dependent** (1.5) : The designation given to a facility that is not completely specified by
14 this standard. Such a facility shall be provided by a processor, with methods or semantics determined
15 by the processor.
- 16 **program** (2.2.2) : A set of program units that includes exactly one main program.
- 17 **program unit** (2.2.1) : The fundamental component of a program. A sequence of statements, comments,
18 and INCLUDE lines. It may be a main program, a module, an external subprogram, or a block data
19 program unit.
- 20 **prototype** : The C analog of a function interface body; defined in 6.7.5.3 of the C International
21 Standard.
- 22 **pure procedure** (12.7) : A procedure that is a pure intrinsic procedure (13.1), is defined by a pure
23 subprogram, or is a statement function that references only pure functions.
- 24 **rank** (2.4.4, 2.4.5) : The number of dimensions of an array. Zero for a scalar.
- 25 **record** (9.1) : A sequence of values or characters that is treated as a whole within a file.
- 26 **reference** (2.5.6) : The appearance of an object designator in a context requiring the value at that
27 point during execution, the appearance of a procedure designator, its operator symbol, or a defined
28 assignment statement in a context requiring execution of the procedure at that point, or the appearance
29 of a module name in a USE statement. Neither the act of defining a variable nor the appearance of the
30 name of a procedure as an actual argument is regarded as a reference.
- 31 **result variable** (2.2.4, 12.6.2.1) : The variable that returns the value of a function.
- 32 **rounding mode** (14.3, 10.7.2.3.7) : The method used to choose the result of an operation that cannot
33 be represented exactly. In IEEE arithmetic, there are four modes; nearest, towards zero, up (towards
34 ∞), and down (towards $-\infty$). In addition, for input/output the two additional modes COMPATIBLE
35 and PROCESSOR_DEFINED are provided.
- 36 **scalar** (2.4.4) :
37 Not being an array.
- 38 **scope** (16) : That part of a program within which a lexical token has a single interpretation. It may be
39 a program, a scoping unit, a construct, a single statement, or a part of a statement.
- 40 **scoping unit** (2.2.1) : One of the following:

- 1 (1) A program unit or subprogram, excluding any scoping units in it,
2 (2) A derived-type definition, or
3 (3) An interface body, excluding any scoping units in it.
- 4 **section subscript** (6.2.2) : A subscript, vector subscript, or subscript triplet in an array section selector.
- 5 **selector** (6.1.1, 6.1.2, 6.1.4, 8.1.5, 8.1.3) : A syntactic mechanism for designating
- 6 (1) a subobject,
7 (2) the set of values for which a CASE block is executed,
8 (3) the object whose type determines which branch of a SELECT TYPE construct is executed,
9 or
10 (4) the object that is associated with the *associate-name* in an ASSOCIATE construct.
- 11 **shape** (2.4.5) : The rank and extents of an array. The shape may be represented by the rank-one array
12 whose elements are the extents in each dimension.
- 13 **size** (2.4.5) : The total number of elements of an array.
- 14 **specification expression** (7.1.6) : An expression with limitations that make it suitable for use in
15 specifications such as length type parameters or array bounds.
- 16 **specification function** (7.1.6) : A nonintrinsic function that may be used in a specification expression.
- 17 **standard-conforming program** (1.5) : A program that uses only those forms and relationships de-
18 scribed in this standard, and that has an interpretation according to this standard.
- 19 **statement** (3.3) : A sequence of lexical tokens. It usually consists of a single line, but a statement may
20 be continued from one line to another and the semicolon symbol may be used to separate statements
21 within a line.
- 22 **statement entity** (16) : An entity identified by a lexical token whose scope is a single statement or
23 part of a statement.
- 24 **statement function** (12.6.4) : A procedure specified by a single statement that is similar in form to an assignment
25 statement.
- 26 **statement label** (3.2.4) : A lexical token consisting of up to five digits that precedes a statement and
27 may be used to refer to the statement.
- 28 **storage association** (16.5.3) : The relationship between two storage sequences if a storage unit of one
29 is the same as a storage unit of the other.
- 30 **storage sequence** (16.5.3.2) : A sequence of contiguous storage units.
- 31 **storage unit** (16.5.3.2) : A character storage unit, a numeric storage unit, a file storage unit, or an
32 unspecified storage unit.
- 33 **stride** (6.2.2.3.1) : The increment specified in a subscript triplet.
- 34 **struct** : The C analog of a sequence derived type; defined in 6.2.5 of the C International Standard.
- 35 **structure** (2.4.1.2) : A scalar data object of derived type.
- 36 **structure component** (6.1.2) : A part of an object of derived type. It may be referenced by an object
37 designator.
- 38 **structure constructor** (4.5.10) : A syntactic mechanism for constructing a value of derived type.

- 1 **subcomponent** (6.1.2) : A subcomponent of an object of derived type is a component of that object
2 or of a subobject of that object.
- 3 **submodule** (2.2.6, 11.2.3) : A program unit that extends a module or another submodule.
- 4 **subobject** (2.4.3.1) : A portion of a data object that may be referenced or defined independently of
5 other portions. It may be an array element, an array section, a structure component, a substring, or the
6 real or imaginary part of a complex object.
- 7 **subprogram** (2.2.1) : A function subprogram or a subroutine subprogram. Note that in FORTRAN 77,
8 a block data program unit was called a subprogram.
- 9 **subroutine** (2.2.4) : A procedure that is invoked by a CALL statement or by a defined assignment
10 statement.
- 11 **subroutine subprogram** (12.6.2.2) : A sequence of statements beginning with a SUBROUTINE state-
12 ment that is not in an interface block and ending with the corresponding END statement.
- 13 **subscript** (6.2.2) : One of the list of scalar integer expressions in an array element selector. Note that
14 in FORTRAN 77, the whole list was called the subscript.
- 15 **subscript triplet** (6.2.2) : An item in the list of an array section selector that contains a colon and
16 specifies a regular sequence of integer values.
- 17 **substring** (6.1.1) : A contiguous portion of a scalar character string. Note that an array section can
18 include a substring selector; the result is called an array section and not a substring.
- 19 **target** (2.4.7, 5.3.16, 6.3.1.2) : A data entity that has the TARGET attribute, or an entity that is
20 associated with a pointer.
- 21 **team** (2.3.2) : A set of images formed by invoking the intrinsic collective subroutine FORM_TEAM
22 (13.7.69).
- 23 **team synchronization** (8.5.3) : Synchronization of the images in a team.
- 24 **transformational function** (13.1) : A function that is either intrinsic or is defined in an intrinsic
25 module and that is neither an elemental function nor an inquiry function.
- 26 **type** (2.4.1) : A named category of data that is characterized by a set of values, together with a way to
27 denote these values and a collection of operations that interpret and manipulate the values. The set of
28 data values depends on the values of the type parameters.
- 29 **type-bound procedure** (4.5.5) : A procedure binding in a type definition. The procedure may be
30 referenced by the *binding-name* via any object of that dynamic type, as a defined operator, by defined
31 assignment, or as part of the finalization process.
- 32 **type compatible** (4.3.1.3) : All entities are type compatible with other entities of the same type.
33 Unlimited polymorphic entities are type compatible with all entities; other polymorphic entities are type
34 compatible with entities whose dynamic type is an extension type of the polymorphic entity's declared
35 type.
- 36 **type declaration statement** (5.2) : An INTEGER, REAL, DOUBLE PRECISION, COMPLEX,
37 CHARACTER, LOGICAL, TYPE (*type-name*), or CLASS (*type-name*) statement.
- 38 **type parameter** (2.4.1.1) : A parameter of a data type. KIND and LEN are the type parameters of
39 intrinsic types. The type parameters of a derived type are defined in the derived-type definition.
- 40 **type parameter order** (4.5.3.2) : The ordering of the type parameters of a derived type that is used

- 1 for derived-type specifiers.
- 2 **ultimate component** (4.5) : For a structure, a component that is of intrinsic type, has the ALLOCAT-
3 ABLE attribute, or has the POINTER attribute, or an ultimate component of a derived-type component
4 that does not have the POINTER attribute or the ALLOCATABLE attribute.
- 5 **undefined** (2.5.5) : For a data object, the property of not having a determinate value.
- 6 **unsigned** : A qualifier of a C numeric type indicating that it is comprised only of nonnegative values;
7 defined in 6.2.5 of the C International Standard. There is nothing analogous in Fortran.
- 8 **unspecified storage unit** (16.5.3.2) : A unit of storage for holding a pointer or a scalar that is not a
9 pointer and is of type other than default integer, default character, default real, double precision real,
10 default logical, or default complex.
- 11 **use association** (16.5.1.3) : The association of names in different scoping units specified by a USE
12 statement.
- 13 **variable** (2.4.3.1.1) : A data object whose value can be defined and redefined during the execution of
14 a program. It may be a named data object, an array element, an array section, a structure component,
15 or a substring. Note that in FORTRAN 77, a variable was always scalar and named.
- 16 **vector subscript** (6.2.2.3.2) : A section subscript that is an integer expression of rank one.
- 17 **void** : A C type comprising an empty set of values; defined in 6.2.5 of the C International Standard.
18 There is nothing analogous in Fortran.
- 19 **whole array** (6.2.1) : A named array.

Annex B

(Informative)

Decremental features

B.1 Deleted features

The deleted features are those features of Fortran 90 that were redundant and considered largely unused.

The following Fortran 90 features are not required by Fortran 95, Fortran 2003, or this part of ISO/IEC 1539.

- (1) Real and double precision DO variables.

The ability present in FORTRAN 77, and for consistency also in Fortran 90, for a DO variable to be of type real or double precision in addition to type integer, has been deleted. A similar result can be achieved by using a DO construct with no loop control and the appropriate exit test.

- (2) Branching to an END IF statement from outside its block.

In FORTRAN 77, and for consistency also in Fortran 90, it was possible to branch to an END IF statement from outside the IF construct; this has been deleted. A similar result can be achieved by branching to a CONTINUE statement that is immediately after the END IF statement.

- (3) PAUSE statement.

The PAUSE statement, present in FORTRAN 66, FORTRAN 77 and for consistency also in Fortran 90, has been deleted. A similar result can be achieved by writing a message to the appropriate unit, followed by reading from the appropriate unit.

- (4) ASSIGN and assigned GO TO statements and assigned format specifiers.

The ASSIGN statement and the related assigned GO TO statement, present in FORTRAN 66, FORTRAN 77 and for consistency also in Fortran 90, have been deleted. Further, the ability to use an assigned integer as a format, present in FORTRAN 77 and Fortran 90, has been deleted. A similar result can be achieved by using other control constructs instead of the assigned GOTO statement and by using a default character variable to hold a format specification instead of using an assigned integer.

- (5) H edit descriptor.

In FORTRAN 77, and for consistency also in Fortran 90, there was an alternative form of character string edit descriptor, which had been the only such form in FORTRAN 66; this has been deleted. A similar result can be achieved by using a character string edit descriptor.

The following is a list of the previous editions of the Fortran International Standard, along with their informal names.

- ISO R 1539-1972, FORTRAN 66;
- ISO 1539-1980, FORTRAN 77;
- ISO/IEC 1539:1991, Fortran 90;
- ISO/IEC 1539-1:1997, Fortran 95.

See ISO/IEC 1539:1991 for detailed rules of how these deleted features work.

1 **B.2 Obsolescent features**

2 The obsolescent features are those features of Fortran 90 that were redundant and for which better
3 methods were available in Fortran 90. Subclause 1.8.3 describes the nature of the obsolescent features.
4 The obsolescent features in this standard are the following.

- 5 (1) Arithmetic IF — use the IF statement or IF construct (8.1.8).
- 6 (2) Shared DO termination and termination on a statement other than END DO or CON-
7 TINUE — use an END DO or a CONTINUE statement for each DO statement.
- 8 (3) Alternate return — see B.2.1.
- 9 (4) Computed GO TO statement — see B.2.2.
- 10 (5) Statement functions — see B.2.3.
- 11 (6) DATA statements amongst executable statements — see B.2.4.
- 12 (7) Assumed length character functions — see B.2.5.
- 13 (8) Fixed form source — see B.2.6.
- 14 (9) CHARACTER* form of CHARACTER declaration — see B.2.7.

15 **B.2.1 Alternate return**

16 An alternate return introduces labels into an argument list to allow the called procedure to direct the
17 execution of the caller upon return. The same effect can be achieved with a return code that is used
18 in a CASE construct on return. This avoids an irregularity in the syntax and semantics of argument
19 association. For example,

```
20 CALL SUBR_NAME (X, Y, Z, *100, *200, *300)
```

21 may be replaced by

```
22 CALL SUBR_NAME (X, Y, Z, RETURN_CODE)  
23 SELECT CASE (RETURN_CODE)  
24 CASE (1)  
25 ...  
26 CASE (2)  
27 ...  
28 CASE (3)  
29 ...  
30 CASE DEFAULT  
31 ...  
32 END SELECT
```

33 **B.2.2 Computed GO TO statement**

34 The computed GO TO has been superseded by the CASE construct, which is a generalized, easier to
35 use and more efficient means of expressing the same computation.

36 **B.2.3 Statement functions**

37 Statement functions are subject to a number of nonintuitive restrictions and are a potential source of
38 error because their syntax is easily confused with that of an assignment statement.

39 The internal function is a more generalized form of the statement function and completely supersedes
40 it.

1 **B.2.4 DATA statements among executables**

2 The statement ordering rules of FORTRAN 66, and hence of FORTRAN 77 and Fortran 90 for compatibility,
3 allowed DATA statements to appear anywhere in a program unit after the specification statements. The
4 ability to position DATA statements amongst executable statements is very rarely used, is unnecessary
5 and is a potential source of error.

6 **B.2.5 Assumed character length functions**

7 Assumed character length for functions is an irregularity in the language in that elsewhere in Fortran
8 the philosophy is that the attributes of a function result depend only on the actual arguments of the
9 invocation and on any data accessible by the function through host or use association. Some uses of this
10 facility can be replaced with an automatic character length function, where the length of the function
11 result is declared in a specification expression. Other uses can be replaced by the use of a subroutine
12 whose arguments correspond to the function result and the function arguments.

13 Note that dummy arguments of a function may be assumed character length.

14 **B.2.6 Fixed form source**

15 Fixed form source was designed when the principal machine-readable input medium for new programs
16 was punched cards. Now that new and amended programs are generally entered via keyboards with
17 screen displays, it is an unnecessary overhead, and is potentially error-prone, to have to locate positions
18 6, 7, or 72 on a line. Free form source was designed expressly for this more modern technology.

19 It is a simple matter for a software tool to convert from fixed to free form source.

20 **B.2.7 CHARACTER* form of CHARACTER declaration**

21 Fortran 90 had two different forms of specifying the length selector in CHARACTER declarations. The
22 older form (CHARACTER*char-length) is redundant.

Annex C

(Informative)

Extended notes

C.1 Clause 4 notes

C.1.1 Selection of the approximation methods (4.4.3)

One can select the real approximation method for an entire program through the use of a module and the parameterized real type. This is accomplished by defining a named integer constant to have a particular kind type parameter value and using that named constant in all real, complex, and derived-type declarations. For example, the specification statements

```

10 INTEGER, PARAMETER :: LONG_FLOAT = 8
11 REAL (LONG_FLOAT) X, Y
12 COMPLEX (LONG_FLOAT) Z

```

specify that the approximation method corresponding to a kind type parameter value of 8 is supplied for the data objects X, Y, and Z in the program unit. The kind type parameter value LONG_FLOAT can be made available to an entire program by placing the INTEGER specification statement in a module and accessing the named constant LONG_FLOAT with a USE statement. Note that by changing 8 to 4 once in the module, a different approximation method is selected.

To avoid the use of the processor-dependent values 4 or 8, replace 8 by KIND (0.0) or KIND (0.0D0). Another way to avoid these processor-dependent values is to select the kind value using the intrinsic inquiry function SELECTED_REAL_KIND. This function, given integer arguments P and R specifying minimum requirements for decimal precision and decimal exponent range, respectively, returns the kind type parameter value of the approximation method that has at least P decimal digits of precision and at least a range for positive numbers of 10^{-R} to 10^R . In the above specification statement, the 8 may be replaced by, for instance, SELECTED_REAL_KIND (10, 50), which requires an approximation method to be selected with at least 10 decimal digits of precision and a range from 10^{-50} to 10^{50} . There are no magnitude or ordering constraints placed on kind values, in order that implementers may have flexibility in assigning such values and may add new kinds without changing previously assigned kind values.

As kind values have no portable meaning, a good practice is to use them in programs only through named constants as described above (for example, SINGLE, IEEE_SINGLE, DOUBLE, and QUAD), rather than using the kind values directly.

C.1.2 Type extension and component accessibility (4.5.2.2, 4.5.4)

The default accessibility of an extended type may be specified in the type definition. The accessibility of its components may be specified individually.

```

34 module types
35     type base_type
36         private                !-- Sets default accessibility

```

```

1     integer :: i           !-- a private component
2     integer, private :: j !-- another private component
3     integer, public  :: k !-- a public component
4     end type base_type
5
6     type, extends(base_type) :: my_type
7     private                   !-- Sets default for components declared in my_type
8     integer :: l             !-- A private component.
9     integer, public :: m !-- A public component.
10    end type my_type
11
12    end module types
13
14    subroutine sub
15    use types
16    type (my_type) :: x
17
18    ....
19
20    call another_sub( &
21    x%base_type, & !-- ok because base_type is a public subobject of x
22    x%base_type%k, & !-- ok because x%base_type is ok and has k as a
23    !-- public component.
24    x%k, & !-- ok because it is shorthand for x%base_type%k
25    x%base_type%i, & !-- Invalid because i is private.
26    x%i) !-- Invalid because it is shorthand for x%base_type%i
27    end subroutine sub

```

28 C.1.3 Abstract types

29 The following defines an object that can be displayed in an X window:

```

30     TYPE, ABSTRACT :: DRAWABLE_OBJECT
31     REAL, DIMENSION(3) :: RGB_COLOR=(/1.0,1.0,1.0/) ! White
32     REAL, DIMENSION(2) :: POSITION=(/0.0,0.0/) ! Centroid
33     CONTAINS
34     PROCEDURE(RENDER_X), PASS(OBJECT), DEFERRED :: RENDER
35     END TYPE DRAWABLE_OBJECT
36
37     ABSTRACT INTERFACE
38     SUBROUTINE RENDER_X(OBJECT, WINDOW)
39     CLASS(DRAWABLE_OBJECT), INTENT(IN) :: OBJECT
40     CLASS(X_WINDOW), INTENT(INOUT) :: WINDOW
41     END SUBROUTINE RENDER_X
42     END INTERFACE

```

1 We can declare a nonabstract type by extending the abstract type:

```

2     TYPE, EXTENDS(DRAWABLE_OBJECT) :: DRAWABLE_TRIANGLE ! Not ABSTRACT
3         REAL, DIMENSION(2,3) :: VERTICES ! In relation to centroid
4     CONTAINS
5         PROCEDURE, PASS(OBJECT) :: RENDER=>RENDER_TRIANGLE_X
6     END TYPE DRAWABLE_TRIANGLE

```

7 The actual drawing procedure will draw a triangle in WINDOW with vertices at x coordinates
8 OBJECT%POSITION(1)+OBJECT%VERTICES(1,:) and y coordinates
9 OBJECT%POSITION(2)+OBJECT%VERTICES(2,:):

```

10    SUBROUTINE RENDER_TRIANGLE_X(OBJECT, WINDOW)
11        CLASS(DRAWABLE_TRIANGLE), INTENT(IN) :: OBJECT
12        CLASS(X_WINDOW), INTENT(INOUT) :: WINDOW
13        ...
14    END SUBROUTINE RENDER_TRIANGLE_X

```

15 C.1.4 Pointers (4.5.2)

16 Pointers are names that can change dynamically their association with a target object. In a sense, a
17 normal variable is a name with a fixed association with a particular object. A normal variable name
18 refers to the same storage space throughout the lifetime of the variable. A pointer name may refer
19 to different storage space, or even no storage space, at different times. A variable may be considered
20 to be a descriptor for space to hold values of the appropriate type, type parameters, and array rank
21 such that the values stored in the descriptor are fixed when the variable is created. A pointer also may
22 be considered to be a descriptor, but one whose values may be changed dynamically so as to describe
23 different pieces of storage. When a pointer is declared, space to hold the descriptor is created, but the
24 space for the target object is not created.

25 A derived type may have one or more components that are defined to be pointers. It may have a
26 component that is a pointer to an object of the same derived type. This “recursive” data definition
27 allows dynamic data structures such as linked lists, trees, and graphs to be constructed. For example:

```

28 TYPE NODE          ! Define a ''recursive'' type
29     INTEGER :: VALUE = 0
30     TYPE (NODE), POINTER :: NEXT_NODE => NULL ( )
31 END TYPE NODE
32
33 TYPE (NODE), TARGET :: HEAD          ! Automatically initialized
34 TYPE (NODE), POINTER :: CURRENT, TEMP ! Declare pointers
35 INTEGER :: IOEM, K
36
37 CURRENT => HEAD                      ! CURRENT points to head of list
38
39 DO
40     READ (*, *, IOSTAT = IOEM) K ! Read next value, if any

```

```

1   IF (IOEM /= 0) EXIT
2   ALLOCATE (TEMP)           ! Create new cell each iteration
3   TEMP % VALUE = K         ! Assign value to cell
4   CURRENT % NEXT_NODE => TEMP ! Attach new cell to list
5   CURRENT => TEMP          ! CURRENT points to new end of list
6   END DO

```

7 A list is now constructed and the last linked cell contains a disassociated pointer. A loop can be used
8 to “walk through” the list.

```

9   CURRENT => HEAD
10  DO
11   IF (.NOT. ASSOCIATED (CURRENT % NEXT_NODE)) EXIT
12   CURRENT => CURRENT % NEXT_NODE
13   WRITE (*, *) CURRENT % VALUE
14  END DO

```

15 C.1.5 Structure constructors and generic names

16 A generic name may be the same as a type name. This can be used to emulate user-defined structure
17 constructors for that type, even if the type has private components. For example:

```

18  MODULE mytype_module
19   TYPE mytype
20   PRIVATE
21   COMPLEX value
22   LOGICAL exact
23  END TYPE
24  INTERFACE mytype
25   MODULE PROCEDURE int_to_mytype
26  END INTERFACE
27   ! Operator definitions etc.
28   ...
29  CONTAINS
30   TYPE(mytype) FUNCTION int_to_mytype(i)
31   INTEGER,INTENT(IN) :: i
32   int_to_mytype%value = i
33   int_to_mytype%exact = .TRUE.
34  END FUNCTION
35   ! Procedures to support operators etc.
36   ...
37  END
38
39  PROGRAM example
40   USE mytype_module

```



```

1  TYPE(mytype) x
2  x = mytype(17)
3  END

```

4 The type name may still be used as a generic name if the type has type parameters. For example:

```

5  MODULE m
6  TYPE t(kind)
7  INTEGER, KIND :: kind
8  COMPLEX(kind) value
9  END TYPE
10 INTEGER,PARAMETER :: single = KIND(0.0), double = KIND(0d0)
11 INTERFACE t
12 MODULE PROCEDURE real_to_t1, dble_to_t2, int_to_t1, int_to_t2
13 END INTERFACE
14 ...
15 CONTAINS
16 TYPE(t(single)) FUNCTION real_to_t1(x)
17 REAL(single) x
18 real_to_t1%value = x
19 END FUNCTION
20 TYPE(t(double)) FUNCTION dble_to_t2(x)
21 REAL(double) x
22 dble_to_t2%value = x
23 END FUNCTION
24 TYPE(t(single)) FUNCTION int_to_t1(x,mold)
25 INTEGER x
26 TYPE(t(single)) mold
27 int_to_t1%value = x
28 END FUNCTION
29 TYPE(t(double)) FUNCTION int_to_t2(x,mold)
30 INTEGER x
31 TYPE(t(double)) mold
32 int_to_t2%value = x
33 END FUNCTION
34 ...
35 END
36
37 PROGRAM example
38 USE m
39 TYPE(t(single)) x
40 TYPE(t(double)) y
41 x = t(1.5) ! References real_to_t1
42 x = t(17,mold=x) ! References int_to_t1
y = t(1.5d0) ! References dble_to_t2

```

```

1
2   y = t(42,mold=y)           ! References int_to_t2
3   y = t(kind(0d0)) ((0,1)) ! Uses the structure constructor for type t
4 END

```

5 C.1.6 Generic type-bound procedures

6 Example of a derived type with generic type-bound procedures:

7 The only difference between this example and the same thing rewritten to use generic interface blocks
8 is that with type-bound procedures,

```
9     USE(rational_numbers),ONLY :: rational
```

10 does not block the type-bound procedures; the user still gets access to the defined assignment and
11 extended operations.

```

12 MODULE rational_numbers
13   IMPLICIT NONE
14   PRIVATE
15   TYPE,PUBLIC :: rational
16     PRIVATE
17     INTEGER n,d
18   CONTAINS
19     ! ordinary type-bound procedure
20     PROCEDURE :: real => rat_to_real
21     ! specific type-bound procedures for generic support
22     PROCEDURE,PRIVATE :: rat_asgn_i
23     PROCEDURE,PRIVATE :: rat_plus_rat
24     PROCEDURE,PRIVATE :: rat_plus_i
25     PROCEDURE,PRIVATE,PASS(b) :: i_plus_rat
26     ! generic type-bound procedures
27     GENERIC :: ASSIGNMENT(=) => rat_asgn_i
28     GENERIC :: OPERATOR(+) => rat_plus_rat, rat_plus_i, i_plus_rat
29   END TYPE
30 CONTAINS
31   ELEMENTAL REAL FUNCTION rat_to_real(this) RESULT(r)
32     CLASS(rational),INTENT(IN) :: this
33     r = REAL(this%n)/this%d
34   END FUNCTION
35   ELEMENTAL SUBROUTINE rat_asgn_i(a,b)
36     CLASS(rational),INTENT(OUT) :: a
37     INTEGER,INTENT(IN) :: b
38     a%n = b
39     a%d = 1
40   END SUBROUTINE

```

```

1  ELEMENTAL TYPE(rational) FUNCTION rat_plus_i(a,b) RESULT(r)
2      CLASS(rational),INTENT(IN) :: a
3      INTEGER,INTENT(IN) :: b
4      r%n = a%n + b*a%d
5      r%d = a%d
6  END FUNCTION
7  ELEMENTAL TYPE(rational) FUNCTION i_plus_rat(a,b) RESULT(r)
8      INTEGER,INTENT(IN) :: a
9      CLASS(rational),INTENT(IN) :: b
10     r%n = b%n + a*b%d
11     r%d = b%d
12 END FUNCTION
13 ELEMENTAL TYPE(rational) FUNCTION rat_plus_rat(a,b) RESULT(r)
14     CLASS(rational),INTENT(IN) :: a,b
15     r%n = a%n*b%d + b%n*a%d
16     r%d = a%d*b%d
17 END FUNCTION
18 END

```

19 C.1.7 Final subroutines (4.5.6, 4.5.6.2, 4.5.6.3, 4.5.6.4)

20 Example of a parameterized derived type with final subroutines:

```

21 MODULE m
22     TYPE t(k)
23         INTEGER, KIND :: k
24         REAL(k),POINTER :: vector(:) => NULL()
25     CONTAINS
26         FINAL :: finalize_t1s, finalize_t1v, finalize_t2e
27     END TYPE
28 CONTAINS
29     SUBROUTINE finalize_t1s(x)
30         TYPE(t(KIND(0.0))) x
31         IF (ASSOCIATED(x%vector)) DEALLOCATE(x%vector)
32     END SUBROUTINE
33     SUBROUTINE finalize_t1v(x)
34         TYPE(t(KIND(0.0))) x(:)
35         DO i=LBOUND(x,1),UBOUND(x,1)
36             IF (ASSOCIATED(x(i)%vector)) DEALLOCATE(x(i)%vector)
37         END DO
38     END SUBROUTINE
39     ELEMENTAL SUBROUTINE finalize_t2e(x)
40         TYPE(t(KIND(0.0d0))),INTENT(INOUT) :: x
41         IF (ASSOCIATED(x%vector)) DEALLOCATE(x%vector)
42     END SUBROUTINE

```

```

1  END MODULE
2
3  SUBROUTINE example(n)
4    USE m
5    TYPE(t(KIND(0.0))) a,b(10),c(n,2)
6    TYPE(t(KIND(0.0d0))) d(n,n)
7    ...
8    ! Returning from this subroutine will effectively do
9    !   CALL finalize_t1s(a)
10   !   CALL finalize_t1v(b)
11   !   CALL finalize_t2e(d)
12   ! No final subroutine will be called for variable C because the user
13   ! omitted to define a suitable specific procedure for it.
14  END SUBROUTINE

```

15 Example of extended types with final subroutines:

```

16  MODULE m
17    TYPE t1
18      REAL a,b
19    END TYPE
20    TYPE,EXTENDS(t1) :: t2
21      REAL,POINTER :: c(:),d(:)
22    CONTAINS
23      FINAL :: t2f
24    END TYPE
25    TYPE,EXTENDS(t2) :: t3
26      REAL,POINTER :: e
27    CONTAINS
28      FINAL :: t3f
29    END TYPE
30    ...
31  CONTAINS
32    SUBROUTINE t2f(x) ! Finalizer for TYPE(t2)'s extra components
33      TYPE(t2) :: x
34      IF (ASSOCIATED(x%c)) DEALLOCATE(x%c)
35      IF (ASSOCIATED(x%d)) DEALLOCATE(x%d)
36    END SUBROUTINE
37    SUBROUTINE t3f(y) ! Finalizer for TYPE(t3)'s extra components
38      TYPE(t3) :: y
39      IF (ASSOCIATED(y%e)) DEALLOCATE(y%e)
40    END SUBROUTINE
41  END MODULE
42
43  SUBROUTINE example

```

```

1
2  USE m
3  TYPE(t1) x1
4  TYPE(t2) x2
5  TYPE(t3) x3
6  ...
7  ! Returning from this subroutine will effectively do
8  !   ! Nothing to x1; it is not finalizable
9  !   CALL t2f(x2)
10 !   CALL t3f(x3)
11 !   CALL t2f(x3%t2)
12 END SUBROUTINE

```

13 C.2 Clause 5 notes

14 C.2.1 The POINTER attribute (5.3.13)

15 The POINTER attribute shall be specified to declare a pointer. The type, type parameters, and rank,
16 which may be specified in the same statement or with one or more attribute specification statements,
17 determine the characteristics of the target objects that may be associated with the pointers declared
18 in the statement. An obvious model for interpreting declarations of pointers is that such declarations
19 create for each name a descriptor. Such a descriptor includes all the data necessary to describe fully
20 and locate in memory an object and all subobjects of the type, type parameters, and rank specified.
21 The descriptor is created empty; it does not contain values describing how to access an actual memory
22 space. These descriptor values will be filled in when the pointer is associated with actual target space.

23 The following example illustrates the use of pointers in an iterative algorithm:

```

24 PROGRAM DYNAM_ITER
25   REAL, DIMENSION (:, :), POINTER :: A, B, SWAP ! Declare pointers
26   ...
27   READ (*, *) N, M
28   ALLOCATE (A (N, M), B (N, M)) ! Allocate target arrays
29   ! Read values into A
30   ...
31   ITER: DO
32     ...
33     ! Apply transformation of values in A to produce values in B
34     ...
35     IF (CONVERGED) EXIT ITER
36     ! Swap A and B
37     SWAP => A; A => B; B => SWAP
38   END DO ITER
39   ...
40 END PROGRAM DYNAM_ITER

```

C.2.2 The TARGET attribute (5.3.16)

1

2 The TARGET attribute shall be specified for any nonpointer object that may, during the execution of the
 3 program, become associated with a pointer. This attribute is defined primarily for optimization purposes.
 4 It allows the processor to assume that any nonpointer object not explicitly declared as a target may
 5 be referred to only by way of its original declared name. It also means that implicitly-declared objects
 6 shall not be used as pointer targets. This will allow a processor to perform optimizations that otherwise
 7 would not be possible in the presence of certain pointers.

8 The following example illustrates the use of the TARGET attribute in an iterative algorithm:

```

9 PROGRAM ITER
10   REAL, DIMENSION (1000, 1000), TARGET :: A, B
11   REAL, DIMENSION (:, :), POINTER      :: IN, OUT, SWAP
12   ...
13   ! Read values into A
14   ...
15   IN => A           ! Associate IN with target A
16   OUT => B          ! Associate OUT with target B
17   ...
18   ITER:DO
19     ...
20     ! Apply transformation of IN values to produce OUT
21     ...
22     IF (CONVERGED) EXIT ITER
23     ! Swap IN and OUT
24     SWAP => IN; IN => OUT; OUT => SWAP
25   END DO ITER
26   ...
27 END PROGRAM ITER

```

28 C.2.3 The VOLATILE attribute (5.3.18)

29 The following example shows the use of a variable with the VOLATILE attribute to communicate with
 30 an asynchronous process, in this case the operating system. The program detects a user keystroke on
 31 the terminal and reacts at a convenient point in its processing.

32 The VOLATILE attribute is necessary to prevent an optimizing compiler from storing the communication
 33 variable in a register or from doing flow analysis and deciding that the EXIT statement can never be
 34 executed.

```

35 SUBROUTINE TERMINATE_ITERATIONS
36
37   LOGICAL, VOLATILE :: USER_HIT_ANY_KEY
38
39   ! Have the OS start to look for a user keystroke and set the variable
40   ! "USER_HIT_ANY_KEY" to TRUE as soon as it detects a keystroke.
41   ! This pseudo call is operating system dependent.

```

```

1
2 CALL OS_BEGIN_DETECT_USER_KEYSTROKE( USER_HIT_ANY_KEY )
3
4 USER_HIT_ANY_KEY = .FALSE.      ! This will ignore any recent keystrokes
5
6 PRINT *, " Hit any key to terminate iterations!"
7
8 DO I = 1,100
9     ...                          ! Compute a value for R
10    PRINT *, I, R
11    IF (USER_HIT_ANY_KEY)        EXIT
12 ENDDO
13
14 ! Have the OS stop looking for user keystrokes
15 CALL OS_STOP_DETECT_USER_KEYSTROKE
16
17 END SUBROUTINE TERMINATE_ITERATIONS

```

18 C.3 Clause 6 notes

19 C.3.1 Structure components (6.1.2)

20 Components of a structure are referenced by writing the components of successive levels of the structure
 21 hierarchy until the desired component is described. For example,

```

22 TYPE ID_NUMBERS
23     INTEGER SSN
24     INTEGER EMPLOYEE_NUMBER
25 END TYPE ID_NUMBERS
26
27 TYPE PERSON_ID
28     CHARACTER (LEN=30) LAST_NAME
29     CHARACTER (LEN=1) MIDDLE_INITIAL
30     CHARACTER (LEN=30) FIRST_NAME
31     TYPE (ID_NUMBERS) NUMBER
32 END TYPE PERSON_ID
33
34 TYPE PERSON
35     INTEGER AGE
36     TYPE (PERSON_ID) ID
37 END TYPE PERSON
38
39 TYPE (PERSON) GEORGE, MARY
40
41 PRINT *, GEORGE % AGE          ! Print the AGE component

```

```

1
2 PRINT *, MARY % ID % LAST_NAME      ! Print LAST_NAME of MARY
3 PRINT *, MARY % ID % NUMBER % SSN ! Print SSN of MARY
4 PRINT *, GEORGE % ID % NUMBER      ! Print SSN and EMPLOYEE_NUMBER of GEORGE

```

5 A structure component may be a data object of intrinsic type as in the case of GEORGE % AGE or it
6 may be of derived type as in the case of GEORGE % ID % NUMBER. The resultant component may
7 be a scalar or an array of intrinsic or derived type.

```

8 TYPE LARGE
9     INTEGER ELT (10)
10    INTEGER VAL
11 END TYPE LARGE
12
13 TYPE (LARGE) A (5)      ! 5 element array, each of whose elements
14                        ! includes a 10 element array ELT and
15                        ! a scalar VAL.
16 PRINT *, A (1)         ! Prints 10 element array ELT and scalar VAL.
17 PRINT *, A (1) % ELT (3) ! Prints scalar element 3
18                        ! of array element 1 of A.
19 PRINT *, A (2:4) % VAL ! Prints scalar VAL for array elements
20                        ! 2 to 4 of A.

```

21 Components of an object of extensible type that are inherited from the parent type may be accessed as
22 a whole by using the parent component name, or individually, either with or without qualifying them
23 by the parent component name.

24 For example:

```

25 TYPE POINT              ! A base type
26     REAL :: X, Y
27 END TYPE POINT
28 TYPE, EXTENDS(POINT) :: COLOR_POINT ! An extension of TYPE(POINT)
29     ! Components X and Y, and component name POINT, inherited from parent
30     INTEGER :: COLOR
31 END TYPE COLOR_POINT
32
33 TYPE(POINT) :: PV = POINT(1.0, 2.0)
34 TYPE(COLOR_POINT) :: CPV = COLOR_POINT(POINT=PV, COLOR=3)
35
36 PRINT *, CPV%POINT      ! Prints 1.0 and 2.0
37 PRINT *, CPV%POINT%X, CPV%POINT%Y ! And this does, too
38 PRINT *, CPV%X, CPV%Y   ! And this does, too

```

39 C.3.2 Allocation with dynamic type (6.3.1)

1 The following example illustrates the use of allocation with the value and dynamic type of the allocated
 2 object given by another object. The example copies a list of objects of any type. It copies the list
 3 starting at IN_LIST. After copying, each element of the list starting at LIST_COPY has a polymorphic
 4 component, ITEM, for which both the value and type are taken from the ITEM component of the
 5 corresponding element of the list starting at IN_LIST.

```

6 TYPE :: LIST ! A list of anything
7   TYPE(LIST), POINTER :: NEXT => NULL()
8   CLASS(*), ALLOCATABLE :: ITEM
9 END TYPE LIST
10 ...
11 TYPE(LIST), POINTER :: IN_LIST, LIST_COPY => NULL()
12 TYPE(LIST), POINTER :: IN_WALK, NEW_TAIL
13 ! Copy IN_LIST to LIST_COPY
14 IF (ASSOCIATED(IN_LIST)) THEN
15   IN_WALK => IN_LIST
16   ALLOCATE(LIST_COPY)
17   NEW_TAIL => LIST_COPY
18   DO
19     ALLOCATE(NEW_TAIL%ITEM, SOURCE=IN_WALK%ITEM)
20     IN_WALK => IN_WALK%NEXT
21     IF (.NOT. ASSOCIATED(IN_WALK)) EXIT
22     ALLOCATE(NEW_TAIL%NEXT)
23     NEW_TAIL => NEW_TAIL%NEXT
24   END DO
25 END IF

```

26 C.3.3 Pointer allocation and association

27 The effect of ALLOCATE, DEALLOCATE, NULLIFY, and pointer assignment is that they are inter-
 28 preted as changing the values in the descriptor that is the pointer. An ALLOCATE is assumed to create
 29 space for a suitable object and to “assign” to the pointer the values necessary to describe that space.
 30 A NULLIFY breaks the association of the pointer with the space. A DEALLOCATE breaks the asso-
 31 ciation and releases the space. Depending on the implementation, it could be seen as setting a flag in
 32 the pointer that indicates whether the values in the descriptor are valid, or it could clear the descriptor
 33 values to some (say zero) value indicative of the pointer not pointing to anything. A pointer assignment
 34 copies the values necessary to describe the space occupied by the target into the descriptor that is the
 35 pointer. Descriptors are copied, values of objects are not.

36 If PA and PB are both pointers and PB is associated with a target, then

```
37 PA => PB
```

38 results in PA being associated with the same target as PB. If PB was disassociated, then PA becomes
 39 disassociated.

40 The standard is specified so that such associations are direct and independent. A subsequent statement

```
41 PB => D
```

42 or

1 ALLOCATE (PB)
2 has no effect on the association of PA with its target. A statement
3 DEALLOCATE (PB)
4 deallocates the space that is associated with both PA and PB. PB becomes disassociated, but there is
5 no requirement that the processor make it explicitly recognizable that PA no longer has a target. This
6 leaves PA as a “dangling pointer” to space that has been released. The program shall not use PA again
7 until it becomes associated via pointer assignment or an ALLOCATE statement.
8 DEALLOCATE may only be used to release space that was created by a previous ALLOCATE. Thus
9 the following is invalid:

```
10 REAL, TARGET :: T
11 REAL, POINTER :: P
12     ...
13 P = > T
14 DEALLOCATE (P) ! Not allowed: P's target was not allocated
```

15 The basic principle is that ALLOCATE, NULLIFY, and pointer assignment primarily affect the pointer
16 rather than the target. ALLOCATE creates a new target but, other than breaking its connection with
17 the specified pointer, it has no effect on the old target. Neither NULLIFY nor pointer assignment has
18 any effect on targets. A piece of memory that was allocated and associated with a pointer will become
19 inaccessible to a program if the pointer is nullified or associated with a different target and no other
20 pointer was associated with this piece of memory. Such pieces of memory may be reused by the processor
21 if this is expedient. However, whether such inaccessible memory is in fact reused is entirely processor
22 dependent.

23 C.4 Clause 7 notes

24 C.4.1 Character assignment

25 The FORTRAN 77 restriction that none of the character positions being defined in the character assign-
26 ment statement may be referenced in the expression has been removed (7.4.1.3).

27 C.4.2 Evaluation of function references

28 If more than one function reference appears in a statement, they may be executed in any order (subject to
29 a function result being evaluated after the evaluation of its arguments) and their values shall not depend
30 on the order of execution. This lack of dependence on order of evaluation permits parallel execution of
31 the function references (7.1.8.2).

32 C.4.3 Pointers in expressions

33 A pointer is basically considered to be like any other variable when it is used as a primary in an expression.
34 If a pointer is used as an operand to an operator that expects a value, the pointer will automatically
35 deliver the value stored in the space described by the pointer, that is, the value of the target object
36 associated with the pointer.

1 C.4.4 Pointers on the left side of an assignment

2 A pointer that appears on the left of an intrinsic assignment statement also is dereferenced and is taken
3 to be referring to the space that is its current target. Therefore, the assignment statement specifies the
4 normal copying of the value of the right-hand expression into this target space. All the normal rules of
5 intrinsic assignment hold; the type and type parameters of the expression and the pointer target shall
6 agree and the shapes shall be conformable.

7 For intrinsic assignment of derived types, nonpointer components are assigned and pointer components
8 are pointer assigned. Dereferencing is applied only to entire scalar objects, not selectively to pointer
9 subobjects.

10 For example, suppose a type such as

```
11 TYPE CELL
12     INTEGER :: VAL
13     TYPE (CELL), POINTER :: NEXT_CELL
14 END TYPE CELL
```

15 is defined and objects such as HEAD and CURRENT are declared using

```
16 TYPE (CELL), TARGET :: HEAD
17 TYPE (CELL), POINTER :: CURRENT
```

18 If a linked list has been created and attached to HEAD and the pointer CURRENT has been allocated
19 space, statements such as

```
20 CURRENT = HEAD
21 CURRENT = CURRENT % NEXT_CELL
```

22 cause the contents of the cells referenced on the right to be copied to the cell referred to by CURRENT.
23 In particular, the right-hand side of the second statement causes the pointer component in the cell,
24 CURRENT, to be selected. This pointer is dereferenced because it is in an expression context to produce
25 the target's integer value and a pointer to a cell that is in the target's NEXT_CELL component. The
26 left-hand side causes the pointer CURRENT to be dereferenced to produce its present target, namely
27 space to hold a cell (an integer and a cell pointer). The integer value on the right is copied to the integer
28 space on the left and the pointer component is pointer assigned (the descriptor on the right is copied
29 into the space for a descriptor on the left). When a statement such as

```
30 CURRENT => CURRENT % NEXT_CELL
```

31 is executed, the descriptor value in CURRENT % NEXT_CELL is copied to the descriptor named
32 CURRENT. In this case, CURRENT is made to point at a different target.

33 In the intrinsic assignment statement, the space associated with the current pointer does not change but
34 the values stored in that space do. In the pointer assignment, the current pointer is made to associate
35 with different space. Using the intrinsic assignment causes a linked list of cells to be moved up through
36 the current "window"; the pointer assignment causes the current pointer to be moved down through the
37 list of cells.

38 C.4.5 An example of a FORALL construct containing a WHERE construct

```
INTEGER :: A(5,5)
```

```

1
2 ...
3 FORALL (I = 1:5)
4   WHERE (A(I,:) == 0)
5     A(:,I) = I
6   ELSEWHERE (A(I,:) > 2)
7     A(I,:) = 6
8   END WHERE
9 END FORALL

```

10 If prior to execution of the FORALL, A has the value

```

11 A =      1  0  0  0  0
12         2  1  1  1  0
13         1  2  2  0  2
14         2  1  0  2  3
15         1  0  0  0  0

```

16 After execution of the assignment statements following the WHERE statement A has the value A'. The
 17 mask created from row one is used to mask the assignments to column one; the mask from row two is
 18 used to mask assignments to column two; etc.

```

19 A' =     1  0  0  0  0
20         1  1  1  1  5
21         1  2  2  4  5
22         1  1  3  2  5
23         1  2  0  0  5

```

24 The masks created for assignments following the ELSEWHERE statement are

```

25 .NOT. (A(I,:) == 0) .AND. (A'(I,:) > 2)

```

26 Thus the only elements affected by the assignments following the ELSEWHERE statement are A(3, 5)
 27 and A(4, 5). After execution of the FORALL construct, A has the value

```

28 A =      1  0  0  0  0
29         1  1  1  1  5
30         1  2  2  4  6
31         1  1  3  2  6
32         1  2  0  0  5

```

33 C.4.6 Examples of FORALL statements

34 Example 1:

```

35 FORALL (J=1:M, K=1:N) X(K, J) = Y(J, K)
   FORALL (K=1:N) X(K, 1:M) = Y(1:M, K)

```

1

2 These statements both copy columns 1 through N of array Y into rows 1 through N of array X. They
3 are equivalent to

4 `X(1:N, 1:M) = TRANSPOSE (Y(1:M, 1:N))`

5 Example 2:

6 The following FORALL statement computes five partial sums of subarrays of J.

7 `J = (/ 1, 2, 3, 4, 5 /)`

8 `FORALL (K = 1:5) J(K) = SUM (J(1:K))`

9 SUM is allowed in a FORALL because intrinsic functions are pure (12.7). After execution of the FORALL
10 statement, `J = (/ 1, 3, 6, 10, 15 /)`.

11 Example 3:

12 `FORALL (I = 2:N-1) X(I) = (X(I-1) + 2*X(I) + X(I+1)) / 4`

13 has the same effect as

14 `X(2:N-1) = (X(1:N-2) + 2*X(2:N-1) + X(3:N)) / 4`

15 C.5 Clause 8 notes

16 C.5.1 Loop control

17 Fortran provides several forms of loop control:

- 18 (1) With an iteration count and a DO variable. This is the classic Fortran DO loop.
- 19 (2) Test a logical condition before each execution of the loop (DO WHILE).
- 20 (3) DO “forever”.

21 C.5.2 The CASE construct

22 At most one case block is selected for execution within a CASE construct, and there is no fall-through
23 from one block into another block within a CASE construct. Thus there is no requirement for the user
24 to exit explicitly from a block.

25 C.5.3 Examples of DO constructs

26 The following are all valid examples of block DO constructs.

27 Example 1:

```

28     SUM = 0.0
29     READ (IUN) N
30     OUTER: DO L = 1, N           ! A DO with a construct name
31         READ (IUN) IQUAL, M, ARRAY (1:M)
32         IF (IQUAL < IQUAL_MIN) CYCLE OUTER    ! Skip inner loop
33         INNER: DO 40 I = 1, M       ! A DO with a label and a name

```

```

1          CALL CALCULATE (ARRAY (I), RESULT)
2          IF (RESULT < 0.0) CYCLE
3          SUM = SUM + RESULT
4          IF (SUM > SUM_MAX) EXIT OUTER
5      40    END DO INNER
6          END DO OUTER

```

7 The outer loop has an iteration count of MAX (N, 0), and will execute that number of times or until
8 SUM exceeds SUM_MAX, in which case the EXIT OUTER statement terminates both loops. The inner
9 loop is skipped by the first CYCLE statement if the quality flag, IQUAL, is too low. If CALCULATE
10 returns a negative RESULT, the second CYCLE statement prevents it from being summed. Both loops
11 have construct names and the inner loop also has a label. A construct name is required in the EXIT
12 statement in order to terminate both loops, but is optional in the CYCLE statements because each
13 belongs to its innermost loop.

14 Example 2:

```

15          N = 0
16          DO 50, I = 1, 10
17              J = I
18              DO K = 1, 5
19                  L = K
20                  N = N + 1 ! This statement executes 50 times
21              END DO ! Nonlabeled DO inside a labeled DO
22          50 CONTINUE

```

23 After execution of the above program fragment, I = 11, J = 10, K = 6, L = 5, and N = 50.

24 Example 3:

```

25          N = 0
26          DO I = 1, 10
27              J = I
28              DO 60, K = 5, 1 ! This inner loop is never executed
29                  L = K
30                  N = N + 1
31          60    CONTINUE ! Labeled DO inside a nonlabeled DO
32          END DO

```

33 After execution of the above program fragment, I = 11, J = 10, K = 5, N = 0, and L is not defined by
34 these statements.

35 The following are all valid examples of nonblock DO constructs:

36 Example 4:

```

37          DO 70
38              READ (IUN, '(1X, G14.7)', IOSTAT = IOS) X

```

```

1      IF (IOS /= 0) EXIT
2      IF (X < 0.) GOTO 70
3      CALL SUBA (X)
4      CALL SUBB (X)
5      ...
6      CALL SUBY (X)
7      CYCLE
8      70  CALL SUBNEG (X) ! SUBNEG called only when X < 0.

```

9 This is not a block DO construct because it ends with a statement other than END DO or CONTINUE. The loop will
10 continue to execute until an end-of-file condition or input/output error occurs.

11 Example 5:

```

12      SUM = 0.0
13      READ (IUN) N
14      DO 80, L = 1, N
15          READ (IUN) IQUAL, M, ARRAY (1:M)
16          IF (IQUAL < IQUAL_MIN) M = 0 ! Skip inner loop
17          DO 80 I = 1, M
18              CALL CALCULATE (ARRAY (I), RESULT)
19              IF (RESULT < 0.) CYCLE
20              SUM = SUM + RESULT
21              IF (SUM > SUM_MAX) GOTO 81
22      80  CONTINUE ! This CONTINUE is shared by both loops
23      81  CONTINUE

```

24 This example is similar to Example 1 above, except that the two loops are not block DO constructs because they share
25 the CONTINUE statement with the label 80. The terminal construct of the outer DO is the entire inner DO construct.
26 The inner loop is skipped by forcing M to zero. If SUM grows too large, both loops are terminated by branching to the
27 CONTINUE statement labeled 81. The CYCLE statement in the inner loop is used to skip negative values of RESULT.

28 Example 6:

```

29      N = 0
30      DO 100 I = 1, 10
31          J = I
32          DO 100 K = 1, 5
33              L = K
34      100  N = N + 1 ! This statement executes 50 times

```

35 In this example, the two loops share an assignment statement. After execution of this program fragment, I = 11, J = 10,
36 K = 6, L = 5, and N = 50.

37 Example 7:

```

38      N = 0
39      DO 200 I = 1, 10
40          J = I
41          DO 200 K = 5, 1 ! This inner loop is never executed
42              L = K
43      200  N = N + 1

```

44 This example is very similar to the previous one, except that the inner loop is never executed. After execution of this
45 program fragment, I = 11, J = 10, K = 5, N = 0, and L is not defined by these statements.

1 C.5.4 Examples of invalid DO constructs

2 The following are all examples of invalid skeleton DO constructs:

3 Example 1:

```
4 DO I = 1, 10
5     ...
6 END DO LOOP    ! No matching construct name
```

7 Example 2:

```
8 LOOP: DO 1000 I = 1, 10    ! No matching construct name
9     ...
10 1000 CONTINUE
```

11 Example 3:

```
12 LOOP1: DO
13     ...
14 END DO LOOP2    ! Construct names don't match
```

15 Example 4:

```
16 DO I = 1, 10    ! Label required or ...
17     ...
18 1010 CONTINUE ! ... END DO required
```

19 Example 5:

```
20 DO 1020 I = 1, 10
21     ...
22 1021 END DO      ! Labels don't match
```

23 Example 6:

```
24 FIRST: DO I = 1, 10
25     SECOND: DO J = 1, 5
26     ...
27     END DO FIRST    ! Improperly nested DOs
28 END DO SECOND
```


1 C.6 Clause 9 notes

2 C.6.1 External files (9.2)

3 This standard accommodates, but does not require, file cataloging. To do this, several concepts are
4 introduced.

5 C.6.1.1 File connection (9.4)

6 Before any input/output may be performed on a file, it shall be connected to a unit. The unit then serves
7 as a designator for that file as long as it is connected. To be connected does not imply that “buffers”
8 have or have not been allocated, that “file-control tables” have or have not been filled, or that any other
9 method of implementation has been used. Connection means that (barring some other fault) a READ
10 or WRITE statement may be executed on the unit, hence on the file. Without a connection, a READ
11 or WRITE statement shall not be executed.

12 C.6.1.2 File existence (9.2.1)

13 Totally independent of the connection state is the property of existence, this being a file property. The
14 processor “knows” of a set of files that exist at a given time for a given program. This set would include
15 tapes ready to read, files in a catalog, a keyboard, a printer, etc. The set may exclude files inaccessible
16 to the program because of security, because they are already in use by another program, etc. This
17 standard does not specify which files exist, hence wide latitude is available to a processor to implement
18 security, locks, privilege techniques, etc. Existence is a convenient concept to designate all of the files
19 that a program can potentially process.

20 All four combinations of connection and existence may occur:

Connect	Exist	Examples
Yes	Yes	A card reader loaded and ready to be read
Yes	No	A printer before the first line is written
No	Yes	A file named 'JOAN' in the catalog
No	No	A file on a reel of tape, not known to the processor

21 Means are provided to create, delete, connect, and disconnect files.

22 C.6.1.3 File names (9.4.5.8)

23 A file may have a name. The form of a file name is not specified. If a system does not have some form of
24 cataloging or tape labeling for at least some of its files, all file names will disappear at the termination
25 of execution. This is a valid implementation. Nowhere does this standard require names to survive for
26 any period of time longer than the execution time span of a program. Therefore, this standard does not
27 impose cataloging as a prerequisite. The naming feature is intended to allow use of a cataloging system
28 where one exists.

29 C.6.1.4 File access (9.2.2)

30 This standard does not address problems of security, protection, locking, and many other concepts that
31 may be part of the concept of “right of access”. Such concepts are considered to be in the province of
32 an operating system.

33 The OPEN and INQUIRE statements can be extended naturally to consider these things.

1 Possible access methods for a file are: sequential and direct. The processor may implement two different
2 types of files, each with its own access method. It might also implement one type of file with two different
3 access methods.

4 Direct access to files is of a simple and commonly available type, that is, fixed-length records. The key
5 is a positive integer.

6 **C.6.2 Nonadvancing input/output (9.2.3.1)**

7 Data transfer statements affect the positioning of an external file. In FORTRAN 77, if no error or end-of-
8 file condition exists, the file is positioned after the record just read or written and that record becomes
9 the preceding record. This standard contains the record positioning ADVANCE= specifier in a data
10 transfer statement that provides the capability of maintaining a position within the current record from
11 one formatted data transfer statement to the next data transfer statement. The value NO provides this
12 capability. The value YES positions the file after the record just read or written. The default is YES.

13 The tab edit descriptor and the slash are still appropriate for use with this type of record access but the
14 tab will not reposition before the left tab limit.

15 A BACKSPACE of a file that is positioned within a record causes the specified unit to be positioned
16 before the current record.

17 If the last data transfer statement was WRITE and the file is positioned within a record, the file will be
18 positioned implicitly after the current record before an ENDFILE record is written to the file, that is, a
19 REWIND, BACKSPACE, or ENDFILE statement following a nonadvancing WRITE statement causes
20 the file to be positioned at the end of the current output record before the endfile record is written to
21 the file.

22 This standard provides a SIZE= specifier to be used with nonadvancing data transfer statements. The
23 variable in the SIZE= specifier will contain the count of the number of characters that make up the
24 sequence of values read by the data edit descriptors in this input statement.

25 The count is especially helpful if there is only one list item in the input list because it will contain the
26 number of characters that were present for the item.

27 The EOR= specifier is provided to indicate when an end-of-record condition has been encountered
28 during a nonadvancing data transfer statement. The end-of-record condition is not an error condition.
29 If this specifier is present, the current input list item that required more characters than the record
30 contained will be padded with blanks if PAD= 'YES' is in effect. This means that the input list item
31 was successfully completed. The file will then be positioned after the current record. The IOSTAT=
32 specifier, if present, will be defined with the value of the named constant IOSTAT_EOR from the ISO-
33 FORTRAN_ENV module and the data transfer statement will be terminated. Program execution will
34 continue with the statement specified in the EOR= specifier. The EOR= specifier gives the capability
35 of taking control of execution when the end-of-record has been found. The *do-variables* in *io-implied-*
36 *dos* retain their last defined value and any remaining items in the *input-item-list* retain their definition
37 status when an end-of-record condition occurs. The SIZE= specifier, if present, will contain the number
38 of characters read with the data edit descriptors during this READ statement.

39 For nonadvancing input, the processor is not required to read partial records. The processor may read
40 the entire record into an internal buffer and make successive portions of the record available to successive
41 input statements.

42 In an implementation of nonadvancing input/output in which a nonadvancing write to a terminal device
43 causes immediate display of the output, such a write can be used as a mechanism to output a prompt.

44 In this case, the statement

1 WRITE (*, FMT='(A)', ADVANCE='NO') 'CONTINUE?(Y/N): '

2 would result in the prompt

3 CONTINUE?(Y/N):

4 being displayed with no subsequent line feed.

5 The response, which might be read by a statement of the form

6 READ (*, FMT='(A)') ANSWER

7 can then be entered on the same line as the prompt as in

8 CONTINUE?(Y/N): Y

9 The standard does not require that an implementation of nonadvancing input/output operate in this
 10 manner. For example, an implementation of nonadvancing output in which the display of the output is
 11 deferred until the current record is complete is also standard-conforming. Such an implementation will
 12 not, however, allow a prompting mechanism of this kind to operate.

13 C.6.3 Asynchronous input/output

14 Rather than limit support for asynchronous input/output to what has been traditionally provided by
 15 facilities such as BUFFERIN/BUFFEROUT, this standard builds upon existing Fortran syntax. This
 16 permits alternative approaches for implementing asynchronous input/output, and simplifies the task of
 17 adapting existing standard-conforming programs to use asynchronous input/output.

18 Not all processors will actually perform input/output asynchronously, nor will every processor that does
 19 be able to handle data transfer statements with complicated input/output item lists in an asynchronous
 20 manner. Such processors can still be standard-conforming. Hopefully, the documentation for each
 21 Fortran processor will describe when, if ever, input/output will be performed asynchronously.

22 This standard allows for at least two different conceptual models for asynchronous input/output.

23 Model 1: the processor will perform asynchronous input/output when the item list is simple (perhaps
 24 one contiguous named array) and the input/output is unformatted. The implementation cost is reduced,
 25 and this is the scenario most likely to be beneficial on traditional "big-iron" machines.

26 Model 2: The processor is free to do any of the following:

- 27 (1) on output, create a buffer inside the input/output library, completely formatted, and then
 28 start an asynchronous write of the buffer, and immediately return to the next statement in
 29 the program. The processor is free to wait for previously issued WRITES, or not, or
 30 (2) pass the input/output list addresses to another processor/process, which will process the
 31 list items independently of the processor that executes the user's code. The addresses of the
 32 list items must be computed before the asynchronous READ/WRITE statement completes.
 33 There is still an ordering requirement on list item processing to handle things like READ
 34 (...) N,(a(i),i=1,N).

35 The standard allows a user to issue a large number of asynchronous input/output requests, without
 36 waiting for any of them to complete, and then wait for any or all of them. It may be impossible, and
 37 undesirable to keep track of each of these input/output requests individually.

38 It is not necessary for all requests to be tracked by the runtime library. If an ID= specifier does not
 39 appear in on a READ or WRITE statement, the runtime is free to forget about this particular request
 40 once it has successfully completed. If it gets an ERR or END condition, the processor is free to report
 41 this during any input/output operation to that unit. When an ID= specifier is present, the processor's

1 runtime input/output library is required to keep track of any END or ERR conditions for that particular
2 input/output request. However, if the input/output request succeeds without any exceptional conditions
3 occurring, then the runtime can forget that ID= value if it wishes. Typically, a runtime might only keep
4 track of the last request made, or perhaps a very few. Then, when a user WAITs for a particular request,
5 either the library knows about it (and does the right thing with respect to error handling, etc.), or will
6 assume it is one of those requests that successfully completed and was forgotten about (and will just
7 return without signaling any end or error conditions). It is incumbent on the user to pass valid ID=
8 values. There is no requirement on the processor to detect invalid ID= values. There is of course,
9 a processor dependent limit on how many outstanding input/output requests that generate an end or
10 error condition can be handled before the processor runs out of memory to keep track of such conditions.
11 The restrictions on the SIZE= variables are designed to allow the processor to update such variables at
12 any time (after the request has been processed, but before the WAIT operation), and then forget about
13 them. That's why there is no SIZE= specifier allowed in the various WAIT operations. Only exceptional
14 conditions (errors or ends of files) are expected to be tracked by individual request by the runtime, and
15 then only if an ID= specifier was present. The END= and EOR= specifiers have not been added to all
16 statements that can be WAIT operations. Instead, the IOSTAT variable will have to be queried after a
17 WAIT operation to handle this situation. This choice was made because we expect the WAIT statement
18 to be the usual method of waiting for input/output to complete (and WAIT does support the END=
19 and EOR= specifiers). This particular choice is philosophical, and was not based on significant technical
20 difficulties.

21 Note that the requirement to set the IOSTAT variable correctly requires an implementation to remember
22 which input/output requests got an EOR condition, so that a subsequent wait operation will return the
23 correct IOSTAT value. This means there is a processor defined limit on the number of outstanding
24 nonadvancing input/output requests that got an EOR condition (constrained by available memory to
25 keep track of this information, similar to END/ERR conditions).

26 **C.6.4 OPEN statement (9.4.5)**

27 A file may become connected to a unit either by preconnection or by execution of an OPEN statement.
28 Preconnection is performed prior to the beginning of execution of a program by means external to For-
29 tran. For example, it may be done by job control action or by processor-established defaults. Execution
30 of an OPEN statement is not required to access preconnected files (9.4.4).

31 The OPEN statement provides a means to access existing files that are not preconnected. An OPEN
32 statement may be used in either of two ways: with a file name (open-by-name) and without a file name
33 (open-by-unit). A unit is given in either case. Open-by-name connects the specified file to the specified
34 unit. Open-by-unit connects a processor-dependent default file to the specified unit. (The default file
35 might or might not have a name.)

36 Therefore, there are three ways a file may become connected and hence processed: preconnection, open-
37 by-name, and open-by-unit. Once a file is connected, there is no means in standard Fortran to determine
38 how it became connected.

39 An OPEN statement may also be used to create a new file. In fact, any of the foregoing three connection
40 methods may be performed on a file that does not exist. When a unit is preconnected, writing the first
41 record creates the file. With the other two methods, execution of the OPEN statement creates the file.

42 When an OPEN statement is executed, the unit specified in the OPEN might or might not already be
43 connected to a file. If it is already connected to a file (either through preconnection or by a prior OPEN),
44 then omitting the FILE= specifier in the OPEN statement implies that the file is to remain connected
45 to the unit. Such an OPEN statement may be used to change the values of the blank interpretation
46 mode, decimal edit mode, pad mode, input/output rounding mode, delimiter mode, and sign mode.

47 If the value of the ACTION= specifier is WRITE, then READ statements shall not refer to this connec-

1 tion. ACTION = 'WRITE' does not restrict positioning by a BACKSPACE statement or positioning
 2 specified by the POSITION= specifier with the value APPEND. However, a BACKSPACE statement
 3 or an OPEN statement containing POSITION = 'APPEND' may fail if the processor requires reading
 4 of the file to achieve the positioning.

5 The following examples illustrate these rules. In the first example, unit 10 is preconnected to a SCRATCH
 6 file; the OPEN statement changes the value of PAD= to YES.

```

7 CHARACTER (LEN = 20) CH1
8 WRITE (10, '(A)') 'THIS IS RECORD 1'
9 OPEN (UNIT = 10, STATUS = 'OLD', PAD = 'YES')
10 REWIND 10
11 READ (10, '(A20)') CH1 ! CH1 now has the value
12 ! 'THIS IS RECORD 1  '
```

13 In the next example, unit 12 is first connected to a file named FRED, with a status of OLD. The second
 14 OPEN statement then opens unit 12 again, retaining the connection to the file FRED, but changing the
 15 value of the DELIM= specifier to QUOTE.

```

16 CHARACTER (LEN = 25) CH2, CH3
17 OPEN (12, FILE = 'FRED', STATUS = 'OLD', DELIM = 'NONE')
18 CH2 = '''THIS STRING HAS QUOTES.'''
19 ! Quotes in string CH2
20 WRITE (12, *) CH2 ! Written with no delimiters
21 OPEN (12, DELIM = 'QUOTE') ! Now quote is the delimiter
22 REWIND 12
23 READ (12, *) CH3 ! CH3 now has the value
24 ! 'THIS STRING HAS QUOTES.  '
```

25 The next example is invalid because it attempts to change the value of the STATUS= specifier.

```

26 OPEN (10, FILE = 'FRED', STATUS = 'OLD')
27 WRITE (10, *) A, B, C
28 OPEN (10, STATUS = 'SCRATCH') ! Attempts to make FRED
29 ! a SCRATCH file
```

30 The previous example could be made valid by closing the unit first, as in the next example.

```

31 OPEN (10, FILE = 'FRED', STATUS = 'OLD')
32 WRITE (10, *) A, B, C
33 CLOSE (10)
34 OPEN (10, STATUS = 'SCRATCH') ! Opens a different
35 ! SCRATCH file
```

36 C.6.5 Connection properties (9.4.3)

1 When a unit becomes connected to a file, either by execution of an OPEN statement or by preconnection,
2 the following connection properties, among others, may be established.

3 (1) An access method, which is sequential, direct, or stream, is established for the connection
4 (9.4.5.1).

5 (2) A form, which is formatted or unformatted, is established for a connection to a file that
6 exists or is created by the connection. For a connection that results from execution of an
7 OPEN statement, a default form (which depends on the access method, as described in
8 9.2.2) is established if no form is specified. For a preconnected file that exists, a form is
9 established by preconnection. For a preconnected file that does not exist, a form may be
10 established, or the establishment of a form may be delayed until the file is created (for
11 example, by execution of a formatted or unformatted WRITE statement) (9.4.5.9).

12 (3) A record length may be established. If the access method is direct, the connection establishes
13 a record length that specifies the length of each record of the file. An existing file with records
14 that are not all of equal length shall not be connected for direct access.

15 If the access method is sequential, records of varying lengths are permitted. In this case,
16 the record length established specifies the maximum length of a record in the file (9.4.5.13).

17 A processor has wide latitude in adapting these concepts and actions to its own cataloging and job
18 control conventions. Some processors may require job control action to specify the set of files that
19 exist or that will be created by a program. Some processors may require no job control action prior to
20 execution. This standard enables processors to perform dynamic open, close, or file creation operations,
21 but it does not require such capabilities of the processor.

22 The meaning of “open” in contexts other than Fortran may include such things as mounting a tape,
23 console messages, spooling, label checking, security checking, etc. These actions may occur upon job
24 control action external to Fortran, upon execution of an OPEN statement, or upon execution of the first
25 read or write of the file. The OPEN statement describes properties of the connection to the file and
26 might or might not cause physical activities to take place. It is a place for an implementation to define
27 properties of a file beyond those required in standard Fortran.

28 **C.6.6 CLOSE statement (9.4.6)**

29 Similarly, the actions of dismounting a tape, protection, etc. of a “close” may be implicit at the end of
30 a run. The CLOSE statement might or might not cause such actions to occur. This is another place to
31 extend file properties beyond those of standard Fortran. Note, however, that the execution of a CLOSE
32 statement on a unit followed by an OPEN statement on the same unit to the same file or to a different
33 file is a permissible sequence of events. The processor shall not deny this sequence solely because the
34 implementation chooses to do the physical act of closing the file at the termination of execution of the
35 program.

36 **C.7 Clause 10 notes**

37 **C.7.1 Number of records (10.4, 10.5, 10.8.2)**

38 The number of records read by an explicitly formatted advancing input statement can be determined
39 from the following rule: a record is read at the beginning of the format scan (even if the input list is
40 empty), at each slash edit descriptor encountered in the format, and when a format rescan occurs at the
41 end of the format.

42 The number of records written by an explicitly formatted advancing output statement can be determined
43 from the following rule: a record is written when a slash edit descriptor is encountered in the format,
44 when a format rescan occurs at the end of the format, and at completion of execution of the output

1 statement (even if the output list is empty). Thus, the occurrence of n successive slashes between two
2 other edit descriptors causes $n - 1$ blank lines if the records are printed. The occurrence of n slashes at
3 the beginning or end of a complete format specification causes n blank lines if the records are printed.
4 However, a complete format specification containing n slashes ($n > 0$) and no other edit descriptors
5 causes $n + 1$ blank lines if the records are printed. For example, the statements

```
6 PRINT 3  
7 3 FORMAT (/)
```

8 will write two records that cause two blank lines if the records are printed.

9 **C.7.2 List-directed input (10.10.3)**

10 The following examples illustrate list-directed input. A blank character is represented by b.

11 Example 1:

12 Program:

```
13 J = 3  
14 READ *, I  
15 READ *, J
```

16 Sequential input file:

```
17 record 1: b1b,4bbbb  
18 record 2: ,2bbbbbbbb
```

19 Result: I = 1, J = 3.

20 Explanation: The second READ statement reads the second record. The initial comma in the record
21 designates a null value; therefore, J is not redefined.

22 Example 2:

23 Program:

```
24 CHARACTER A *8, B *1  
25 READ *, A, B
```

26 Sequential input file:

```
27 record 1: 'bbbbbbbb'  
28 record 2: 'QXY'b'Z'
```

29 Result: A = 'bbbbbbbb', B = 'Q'

30 Explanation: In the first record, the rightmost apostrophe is interpreted as delimiting the constant (it
31 cannot be the first of a pair of embedded apostrophes representing a single apostrophe because this
32 would involve the prohibited “splitting” of the pair by the end of a record); therefore, A is assigned
33 the character constant 'bbbbbbbb'. The end of a record acts as a blank, which in this case is a value
34 separator because it occurs between two constants.

1 C.8 Clause 11 notes

2 C.8.1 Main program and block data program unit (11.1, 11.3)

3 The name of the main program or of a block data program unit has no explicit use within the Fortran
4 language. It is available for documentation and for possible use by a processor.

5 A processor may implement an unnamed main program or unnamed block data program unit by assigning
6 it a default name. However, this name shall not conflict with any other global name in a standard-
7 conforming program. This might be done by making the default name one that is not permitted in a
8 standard-conforming program (for example, by including a character not normally allowed in names)
9 or by providing some external mechanism such that for any given program the default name can be
10 changed to one that is otherwise unused.

11 C.8.2 Dependent compilation (11.2)

12 This standard, like its predecessors, is intended to permit the implementation of conforming processors
13 in which a program can be broken into multiple units, each of which can be separately translated in
14 preparation for execution. Such processors are commonly described as supporting separate compilation.
15 There is an important difference between the way separate compilation can be implemented under this
16 standard and the way it could be implemented under the FORTRAN 77 International Standard. Under
17 the FORTRAN 77 standard, any information required to translate a program unit was specified in that
18 program unit. Each translation was thus totally independent of all others. Under this standard, a
19 program unit can use information that was specified in a separate module and thus may be dependent
20 on that module. The implementation of this dependency in a processor may be that the translation of a
21 program unit may depend on the results of translating one or more modules. Processors implementing
22 the dependency this way are commonly described as supporting dependent compilation.

23 The dependencies involved here are new only in the sense that the Fortran processor is now aware of
24 them. The same information dependencies existed under the FORTRAN 77 International Standard, but
25 it was the programmer's responsibility to transport the information necessary to resolve them by making
26 redundant specifications of the information in multiple program units. The availability of separate but
27 dependent compilation offers several potential advantages over the redundant textual specification of
28 information.

- 29 (1) Specifying information at a single place in the program ensures that different program units
30 using that information will be translated consistently. Redundant specification leaves the
31 possibility that different information will erroneously be specified. Even if an INCLUDE line
32 is used to ensure that the text of the specifications is identical in all involved program units,
33 the presence of other specifications (for example, an IMPLICIT statement) may change the
34 interpretation of that text.
- 35 (2) During the revision of a program, it is possible for a processor to assist in determining
36 whether different program units have been translated using different (incompatible) versions
37 of a module, although there is no requirement that a processor provide such assistance.
38 Inconsistencies in redundant textual specification of information, on the other hand, tend
39 to be much more difficult to detect.
- 40 (3) Putting information in a module provides a way of packaging it. Without modules, redun-
41 dant specifications frequently shall be interleaved with other specifications in a program
42 unit, making convenient packaging of such information difficult.
- 43 (4) Because a processor may be implemented such that the specifications in a module are
44 translated once and then repeatedly referenced, there is the potential for greater efficiency
45 than when the processor shall translate redundant specifications of information in multiple
46 program units.

1 The exact meaning of the requirement that the public portions of a module be available at the time of
2 reference is processor dependent. For example, a processor could consider a module to be available only
3 after it has been compiled and require that if the module has been compiled separately, the result of
4 that compilation shall be identified to the compiler when compiling program units that use it.

5 **C.8.2.1 USE statement and dependent compilation (11.2.2)**

6 Another benefit of the USE statement is its enhanced facilities for name management. If one needs to
7 use only selected entities in a module, one can do so without having to worry about the names of all
8 the other entities in that module. If one needs to use two different modules that happen to contain
9 entities with the same name, there are several ways to deal with the conflict. If none of the entities with
10 the same name are to be used, they can simply be ignored. If the name happens to refer to the same
11 entity in both modules (for example, if both modules obtained it from a third module), then there is no
12 confusion about what the name denotes and the name can be freely used. If the entities are different
13 and one or both is to be used, the local renaming facility in the USE statement makes it possible to give
14 those entities different names in the program unit containing the USE statements.

15 A benefit of using the ONLY option consistently, as compared to USE without it, is that the module
16 from which each accessed entity is accessed is explicitly specified in each program unit. This means that
17 one need not search other program units to find where each one is defined. This reduces maintenance
18 costs.

19 A typical implementation of dependent but separate compilation may involve storing the result of trans-
20 lating a module in a file (or file element) whose name is derived from the name of the module. Note,
21 however, that the name of a module is limited only by the Fortran rules and not by the names allowed
22 in the file system. Thus the processor may have to provide a mapping between Fortran names and file
23 system names.

24 The result of translating a module could reasonably either contain only the information textually specified
25 in the module (with “pointers” to information originally textually specified in other modules) or contain
26 all information specified in the module (including copies of information originally specified in other
27 modules). Although the former approach would appear to save on storage space, the latter approach
28 can greatly simplify the logic necessary to process a USE statement and can avoid the necessity of
29 imposing a limit on the logical “nesting” of modules via the USE statement.

30 Variables declared in a module retain their definition status on much the same basis as variables in
31 a common block. That is, saved variables retain their definition status throughout the execution of a
32 program, while variables that are not saved retain their definition status only during the execution of
33 scoping units that reference the module. In some cases, it may be appropriate to put a USE statement
34 such as

```
35 USE MY_MODULE, ONLY:
```

36 in a scoping unit in order to assure that other procedures that it references can communicate through
37 the module. In such a case, the scoping unit would not access any entities from the module, but the
38 variables not saved in the module would retain their definition status throughout the execution of the
39 scoping unit.

40 There is an increased potential for undetected errors in a scoping unit that uses both implicit typing
41 and the USE statement. For example, in the program fragment

```
42 SUBROUTINE SUB  
43     USE MY_MODULE  
     IMPLICIT INTEGER (I-N), REAL (A-H, O-Z)
```

```
1   X = F (B)
2   A = G (X) + H (X + 1)
3 END SUBROUTINE SUB
```

4 X could be either an implicitly typed real variable or a variable obtained from the module MY_MODULE
5 and might change from one to the other because of changes in MY_MODULE unrelated to the action
6 performed by SUB. Logic errors resulting from this kind of situation can be extremely difficult to locate.
7 Thus, the use of these features together is discouraged.

8 **C.8.2.2 Accessibility attributes**

9 The PUBLIC and PRIVATE attributes, which can be declared only in modules, divide the entities in a
10 module into those that are actually relevant to a scoping unit referencing the module and those that are
11 not. This information may be used to improve the performance of a Fortran processor. For example,
12 it may be possible to discard much of the information about the private entities once a module has
13 been translated, thus saving on both storage and the time to search it. Similarly, it may be possible
14 to recognize that two versions of a module differ only in the private entities they contain and avoid
15 retranslating program units that use that module when switching from one version of the module to the
16 other.

17 **C.8.3 Examples of the use of modules**

18 **C.8.3.1 Identical common blocks**

19 A common block and all its associated specification statements may be placed in a module named, for
20 example, MY_COMMON and accessed by a USE statement of the form

```
21 USE MY_COMMON
```

22 that accesses the whole module without any renaming. This ensures that all instances of the common
23 block are identical. Module MY_COMMON could contain more than one common block.

24 **C.8.3.2 Global data**

25 A module may contain only data objects, for example:

```
26 MODULE DATA_MODULE
27   SAVE
28   REAL A (10), B, C (20,20)
29   INTEGER :: I=0
30   INTEGER, PARAMETER :: J=10
31   COMPLEX D (J,J)
32 END MODULE DATA_MODULE
```

33 Data objects made global in this manner may have any combination of data types.

34 Access to some of these may be made by a USE statement with the ONLY option, such as:

```
35 USE DATA_MODULE, ONLY: A, B, D
```

36 and access to all of them may be made by the following USE statement:

1 USE DATA_MODULE

2 Access to all of them with some renaming to avoid name conflicts may be made by:

3 USE DATA_MODULE, AMODULE => A, DMODULE => D

4 **C.8.3.3 Derived types**

5 A derived type may be defined in a module and accessed in a number of program units. For example:

```
6 MODULE SPARSE
7     TYPE NONZERO
8         REAL A
9         INTEGER I, J
10    END TYPE NONZERO
11 END MODULE SPARSE
```

12 defines a type consisting of a real component and two integer components for holding the numerical
13 value of a nonzero matrix element and its row and column indices.

14 **C.8.3.4 Global allocatable arrays**

15 Many programs need large global allocatable arrays whose sizes are not known before program execution.

16 A simple form for such a program is:

```
17 PROGRAM GLOBAL_WORK
18     CALL CONFIGURE_ARRAYS      ! Perform the appropriate allocations
19     CALL COMPUTE              ! Use the arrays in computations
20 END PROGRAM GLOBAL_WORK
21 MODULE WORK_ARRAYS          ! An example set of work arrays
22     INTEGER N
23     REAL, ALLOCATABLE, SAVE :: A (:), B (:, :), C (:, :, :)
24 END MODULE WORK_ARRAYS
25 SUBROUTINE CONFIGURE_ARRAYS ! Process to set up work arrays
26     USE WORK_ARRAYS
27     READ (*, *) N
28     ALLOCATE (A (N), B (N, N), C (N, N, 2 * N))
29 END SUBROUTINE CONFIGURE_ARRAYS
30 SUBROUTINE COMPUTE
31     USE WORK_ARRAYS
32     ... ! Computations involving arrays A, B, and C
33 END SUBROUTINE COMPUTE
```

34 Typically, many subprograms need access to the work arrays, and all such subprograms would contain
35 the statement

36 USE WORK_ARRAYS

1 C.8.3.5 Procedure libraries

2 Interface bodies for external procedures in a library may be gathered into a module. This permits the
3 use of argument keywords and optional arguments, and allows static checking of the references. Different
4 versions may be constructed for different applications, using argument keywords in common use in each
5 application.

6 An example is the following library module:

```
7 MODULE LIBRARY_LLS
8   INTERFACE
9     SUBROUTINE LLS (X, A, F, FLAG)
10      REAL X (:, :)
11      ! The SIZE in the next statement is an intrinsic function
12      REAL, DIMENSION (SIZE (X, 2)) :: A, F
13      INTEGER FLAG
14    END SUBROUTINE LLS
15    ...
16  END INTERFACE
17  ...
18 END MODULE LIBRARY_LLS
```

19 This module allows the subroutine LLS to be invoked:

```
20 USE LIBRARY_LLS
21   ...
22 CALL LLS (X = ABC, A = D, F = XX, FLAG = IFLAG)
23   ...
```

24 C.8.3.6 Operator extensions

25 In order to extend an intrinsic operator symbol to have an additional meaning, an interface block
26 specifying that operator symbol in the OPERATOR option of the INTERFACE statement may be
27 placed in a module.

28 For example, // may be extended to perform concatenation of two derived-type objects serving as varying
29 length character strings and + may be extended to specify matrix addition for type MATRIX or interval
30 arithmetic addition for type INTERVAL.

31 A module might contain several such interface blocks. An operator may be defined by an external
32 function (either in Fortran or some other language) and its procedure interface placed in the module.

33 C.8.3.7 Data abstraction

34 In addition to providing a portable means of avoiding the redundant specification of information in
35 multiple program units, a module provides a convenient means of “packaging” related entities, such as
36 the definitions of the representation and operations of an abstract data type. The following example
37 of a module defines a data abstraction for a SET type where the elements of each set are of type
38 integer. The standard set operations of UNION, INTERSECTION, and DIFFERENCE are provided.
39 The CARDINALITY function returns the cardinality of (number of elements in) its set argument.

1 Two functions returning logical values are included, ELEMENT and SUBSET. ELEMENT defines the
 2 operator .IN. and SUBSET extends the operator <=. ELEMENT determines if a given scalar integer
 3 value is an element of a given set, and SUBSET determines if a given set is a subset of another given
 4 set. (Two sets may be checked for equality by comparing cardinality and checking that one is a subset
 5 of the other, or checking to see if each is a subset of the other.)

6 The transfer function SETF converts a vector of integer values to the corresponding set, with duplicate
 7 values removed. Thus, a vector of constant values can be used as set constants. An inverse transfer
 8 function VECTOR returns the elements of a set as a vector of values in ascending order. In this SET
 9 implementation, set data objects have a maximum cardinality of 200.

```

10 MODULE INTEGER_SETS
11 ! This module is intended to illustrate use of the module facility
12 ! to define a new type, along with suitable operators.
13
14 INTEGER, PARAMETER :: MAX_SET_CARD = 200
15
16 TYPE SET                                ! Define SET type
17     PRIVATE
18     INTEGER CARD
19     INTEGER ELEMENT (MAX_SET_CARD)
20 END TYPE SET
21
22 INTERFACE OPERATOR (.IN.)
23     MODULE PROCEDURE ELEMENT
24 END INTERFACE OPERATOR (.IN.)
25
26 INTERFACE OPERATOR (<=)
27     MODULE PROCEDURE SUBSET
28 END INTERFACE OPERATOR (<=)
29
30 INTERFACE OPERATOR (+)
31     MODULE PROCEDURE UNION
32 END INTERFACE OPERATOR (+)
33
34 INTERFACE OPERATOR (-)
35     MODULE PROCEDURE DIFFERENCE
36 END INTERFACE OPERATOR (-)
37
38 INTERFACE OPERATOR (*)
39     MODULE PROCEDURE INTERSECTION
40 END INTERFACE OPERATOR (*)
41
42 CONTAINS
43
44 INTEGER FUNCTION CARDINALITY (A)        ! Returns cardinality of set A
      TYPE (SET), INTENT (IN) :: A

```

```

1
2   CARDINALITY = A % CARD
3 END FUNCTION CARDINALITY
4
5 LOGICAL FUNCTION ELEMENT (X, A)           ! Determines if
6   INTEGER, INTENT(IN) :: X               ! element X is in set A
7   TYPE (SET), INTENT(IN) :: A
8   ELEMENT = ANY (A % ELEMENT (1 : A % CARD) == X)
9 END FUNCTION ELEMENT
10
11 FUNCTION UNION (A, B)                    ! Union of sets A and B
12   TYPE (SET) UNION
13   TYPE (SET), INTENT(IN) :: A, B
14   INTEGER J
15   UNION = A
16   DO J = 1, B % CARD
17     IF (.NOT. (B % ELEMENT (J) .IN. A)) THEN
18       IF (UNION % CARD < MAX_SET_CARD) THEN
19         UNION % CARD = UNION % CARD + 1
20         UNION % ELEMENT (UNION % CARD) = &
21           B % ELEMENT (J)
22       ELSE
23         ! Maximum set size exceeded . . .
24       END IF
25     END IF
26   END DO
27 END FUNCTION UNION
28
29 FUNCTION DIFFERENCE (A, B)               ! Difference of sets A and B
30   TYPE (SET) DIFFERENCE
31   TYPE (SET), INTENT(IN) :: A, B
32   INTEGER J, X
33   DIFFERENCE % CARD = 0                   ! The empty set
34   DO J = 1, A % CARD
35     X = A % ELEMENT (J)
36     IF (.NOT. (X .IN. B)) DIFFERENCE = DIFFERENCE + SET (1, X)
37   END DO
38 END FUNCTION DIFFERENCE
39
40 FUNCTION INTERSECTION (A, B)             ! Intersection of sets A and B
41   TYPE (SET) INTERSECTION
42   TYPE (SET), INTENT(IN) :: A, B
43   INTERSECTION = A - (A - B)
44 END FUNCTION INTERSECTION
45

```

```

1 LOGICAL FUNCTION SUBSET (A, B)           ! Determines if set A is
2   TYPE (SET), INTENT(IN) :: A, B       ! a subset of set B
3   INTEGER I
4   SUBSET = A % CARD <= B % CARD
5   IF (.NOT. SUBSET) RETURN              ! For efficiency
6   DO I = 1, A % CARD
7     SUBSET = SUBSET .AND. (A % ELEMENT (I) .IN. B)
8   END DO
9 END FUNCTION SUBSET
10
11 TYPE (SET) FUNCTION SETF (V)           ! Transfer function between a vector
12   INTEGER V (:)                        ! of elements and a set of elements
13   INTEGER J                             ! removing duplicate elements
14   SETF % CARD = 0
15   DO J = 1, SIZE (V)
16     IF (.NOT. (V (J) .IN. SETF)) THEN
17       IF (SETF % CARD < MAX_SET_CARD) THEN
18         SETF % CARD = SETF % CARD + 1
19         SETF % ELEMENT (SETF % CARD) = V (J)
20       ELSE
21         ! Maximum set size exceeded . . .
22       END IF
23     END IF
24   END DO
25 END FUNCTION SETF
26
27 FUNCTION VECTOR (A)                    ! Transfer the values of set A
28   TYPE (SET), INTENT (IN) :: A ! into a vector in ascending order
29   INTEGER, POINTER :: VECTOR (:)
30   INTEGER I, J, K
31   ALLOCATE (VECTOR (A % CARD))
32   VECTOR = A % ELEMENT (1 : A % CARD)
33   DO I = 1, A % CARD - 1                ! Use a better sort if
34     DO J = I + 1, A % CARD              ! A % CARD is large
35       IF (VECTOR (I) > VECTOR (J)) THEN
36         K = VECTOR (J); VECTOR (J) = VECTOR (I); VECTOR (I) = K
37       END IF
38     END DO
39   END DO
40 END FUNCTION VECTOR
41
42 END MODULE INTEGER_SETS

43 Examples of using INTEGER_SETS (A, B, and C are variables of type SET; X
44 is an integer variable):

```

```

1 ! Check to see if A has more than 10 elements
2 IF (CARDINALITY (A) > 10) ...
3
4 ! Check for X an element of A but not of B
5 IF (X .IN. (A - B)) ...
6
7 ! C is the union of A and the result of B intersected
8 ! with the integers 1 to 100
9 C = A + B * SETF ((/ (I, I = 1, 100) /))
10
11 ! Does A have any even numbers in the range 1:100?
12 IF (CARDINALITY (A * SETF ((/ (I, I = 2, 100, 2) /))) > 0) ...
13
14 PRINT *, VECTOR (B) ! Print out the elements of set B, in ascending order

```

15 C.8.3.8 Public entities renamed

16 At times it may be necessary to rename entities that are accessed with USE statements. Care should be
 17 taken if the referenced modules also contain USE statements.

18 The following example illustrates renaming features of the USE statement.

```

19 MODULE J; REAL JX, JY, JZ; END MODULE J
20 MODULE K
21   USE J, ONLY : KX => JX, KY => JY
22   ! KX and KY are local names to module K
23   REAL KZ      ! KZ is local name to module K
24   REAL JZ      ! JZ is local name to module K
25 END MODULE K
26 PROGRAM RENAME
27   USE J; USE K
28   ! Module J's entity JX is accessible under names JX and KX
29   ! Module J's entity JY is accessible under names JY and KY
30   ! Module K's entity KZ is accessible under name KZ
31   ! Module J's entity JZ and K's entity JZ are different entities
32   ! and shall not be referenced
33   ...
34 END PROGRAM RENAME

```

35 C.8.4 Modules with submodules

36 Each submodule specifies that it is the child of exactly one parent module or submodule. Therefore, a
 37 module and all of its descendant submodules stand in a tree-like relationship one to another.

38 If a module procedure interface body that is specified in a module has public accessibility, and its
 39 corresponding separate module procedure is defined in a descendant of that module, the procedure can
 40 be accessed by use association. No other entity in a submodule can be accessed by use association. Each

1 program unit that references a module by use association depends on it, and each submodule depends
 2 on its ancestor module. Therefore, if one changes a separate module procedure body in a submodule but
 3 does not change its corresponding module procedure interface, a tool for automatic program translation
 4 would not need to reprocess program units that reference the module by use association. This is so
 5 even if the tool exploits the relative modification times of files as opposed to comparing the result of
 6 translating the module to the result of a previous translation.

7 By constructing taller trees, one can put entities at intermediate levels that are shared by submodules
 8 at lower levels; changing these entities cannot change the interpretation of anything that is accessible
 9 from the module by use association. Developers of modules that embody large complicated concepts
 10 can exploit this possibility to organize components of the concept into submodules, while preserving the
 11 privacy of entities that are shared by the submodules and that ought not to be exposed to users of the
 12 module. Putting these shared entities at an intermediate level also prevents cascades of reprocessing
 13 and testing if some of them are changed.

14 The following example illustrates a module, `color_points`, with a submodule, `color_points_a`, that in
 15 turn has a submodule, `color_points_b`. Public entities declared within `color_points` can be accessed
 16 by use association. The submodules `color_points_a` and `color_points_b` can be changed without
 17 causing retranslation of program units that reference the module `color_points`.

18 The module `color_points` does not have a *module-subprogram-part*, but a *module-subprogram-part* is
 19 not prohibited. The module could be published as definitive specification of the interface, without
 20 revealing trade secrets contained within `color_points_a` or `color_points_b`. Of course, a similar
 21 module without the `module` prefix in the interface bodies would serve equally well as documentation –
 22 but the procedures would be external procedures. It would make little difference to the consumer, but
 23 the developer would forfeit all of the advantages of modules.

```

24  module color_points
25
26      type color_point
27          private
28          real :: x, y
29          integer :: color
30      end type color_point
31
32      interface                ! Interfaces for procedures with separate
33                              ! bodies in the submodule color_points_a
34      module subroutine color_point_del ( p ) ! Destroy a color_point object
35          type(color_point), allocatable :: p
36      end subroutine color_point_del
37      ! Distance between two color_point objects
38      real module function color_point_dist ( a, b )
39          type(color_point), intent(in) :: a, b
40      end function color_point_dist
41      module subroutine color_point_draw ( p ) ! Draw a color_point object
42          type(color_point), intent(in) :: p
43      end subroutine color_point_draw
44      module subroutine color_point_new ( p ) ! Create a color_point object
45          type(color_point), allocatable :: p
46      end subroutine color_point_new
  
```

```

1
2     end interface
3
4     end module color_points

```

5 The only entities within `color_points_a` that can be accessed by use association are separate module
6 procedures for which corresponding module procedure interface bodies are provided in `color_points`.
7 If the procedures are changed but their interfaces are not, the interface from program units that access
8 them by use association is unchanged. If the module and submodule are in separate files, utilities that
9 examine the time of modification of a file would notice that changes in the module could affect the
10 translation of its submodules or of program units that reference the module by use association, but
11 that changes in submodules could not affect the translation of the parent module or program units that
12 reference it by use association.

13 The variable `instance_count` is not accessible by use association of `color_points`, but is accessible
14 within `color_points_a`, and its submodules.

```

15     submodule ( color_points ) color_points_a ! Submodule of color_points
16
17     integer, save :: instance_count = 0
18
19     interface                ! Interface for a procedure with a separate
20                             ! body in submodule color_points_b
21     module subroutine inquire_palette ( pt, pal )
22         use palette_stuff    ! palette_stuff, especially submodules
23                             ! thereof, can reference color_points by use
24                             ! association without causing a circular
25                             ! dependence during translation because this
26                             ! use is not in the module. Furthermore,
27                             ! changes in the module palette_stuff do not
28                             ! affect the translation of color_points.
29         type(color_point), intent(in) :: pt
30         type(palette), intent(out) :: pal
31     end subroutine inquire_palette
32
33     end interface
34
35     contains ! Invisible bodies for public module procedure interfaces
36         ! declared in the module
37
38     module subroutine color_point_del ( p )
39         type(color_point), allocatable :: p
40         instance_count = instance_count - 1
41         deallocate ( p )
42     end subroutine color_point_del
43     real module function color_point_dist ( a, b ) result ( dist )
44         type(color_point), intent(in) :: a, b

```

```

1
2     dist = sqrt( (b%x - a%x)**2 + (b%y - a%y)**2 )
3 end function color_point_dist
4 module subroutine color_point_new ( p )
5     type(color_point), allocatable :: p
6     instance_count = instance_count + 1
7     allocate ( p )
8 end subroutine color_point_new
9
10 end submodule color_points_a

```

11 The subroutine `inquire_palette` is accessible within `color_points_a` because its interface is declared
12 therein. It is not, however, accessible by use association, because its interface is not declared in the
13 module, `color_points`. Since the interface is not declared in the module, changes in the interface
14 cannot affect the translation of program units that reference the module by use association.

```

15 submodule ( color_points:color_points_a ) color_points_b ! Subsidiary**2 submodule
16
17 contains
18     ! Invisible body for interface declared in the ancestor module
19     module subroutine color_point_draw ( p )
20         use palette_stuff, only: palette
21         type(color_point), intent(in) :: p
22         type(palette) :: MyPalette
23         ...; call inquire_palette ( p, MyPalette ); ...
24     end subroutine color_point_draw
25
26     ! Invisible body for interface declared in the parent submodule
27     module procedure inquire_palette
28         ... implementation of inquire_palette
29     end procedure inquire_palette
30
31     subroutine private_stuff ! not accessible from color_points_a
32         ...
33     end subroutine private_stuff
34
35 end submodule color_points_b
36
37 module palette_stuff
38     type :: palette ; ... ; end type palette
39 contains
40     subroutine test_palette ( p )
41         ! Draw a color wheel using procedures from the color_points module
42         type(palette), intent(in) :: p
43         use color_points ! This does not cause a circular dependency because
44             ! the "use palette_stuff" that is logically within

```

```

1
2           ! color_points is in the color_points_a submodule.
3       ...
4       end subroutine test_palette
5   end module palette_stuff

```

6 There is a use palette_stuff in color_points_a, and a use color_points in palette_stuff. The
7 use palette_stuff would cause a circular reference if it appeared in color_points. In this case, it
8 does not cause a circular dependence because it is in a submodule. Submodules cannot be referenced
9 by use association, and therefore what would be a circular appearance of use palette_stuff is not
10 accessed.


```


11  program main
12      use color_points
13      ! "instance_count" and "inquire_palette" are not accessible here
14      ! because they are not declared in the "color_points" module.
15      ! "color_points_a" and "color_points_b" cannot be referenced by
16      ! use association.
17      interface draw                ! just to demonstrate it's possible
18          module procedure color_point_draw
19      end interface
20      type(color_point) :: C_1, C_2
21      real :: RC
22      ...
23      call color_point_new (c_1)      ! body in color_points_a, interface in color_points
24      ...
25      call draw (c_1)                ! body in color_points_b, specific interface
26                                     ! in color_points, generic interface here.
27      ...
28      rc = color_point_dist (c_1, c_2) ! body in color_points_a, interface in color_points
29      ...
30      call color_point_del (c_1)     ! body in color_points_a, interface in color_points
31      ...
32  end program main

```

33 A multilevel submodule system can be used to package and organize a large and interconnected concept
34 without exposing entities of one subsystem to other subsystems.

35 Consider a Plasma module from a Tokamak simulator. A plasma simulation requires attention at least to
36 fluid flow, thermodynamics, and electromagnetism. Fluid flow simulation requires simulation of subsonic,
37 supersonic, and hypersonic flow. This problem decomposition can be reflected in the submodule structure
38 of the Plasma module:





Entities can be shared among the `Subsonic`, `Supersonic`, and `Hypersonic` submodules by putting them within the `Flow` submodule. One then need not worry about accidental use of these entities by use association or by the `Thermal` or `Electromagnetics` modules, or the development of a dependency of correct operation of those subsystems upon the representation of entities of the `Flow` subsystem as a consequence of maintenance. Since these these entities are not accessible by use association, if any of them are changed, the new values cannot be accessed in program units that reference the `Plasma` module by use association; the answer to the question “where are these entities used” is therefore confined to the set of descendant submodules of the `Flow` submodule.

C.9 Clause 12 notes

C.9.1 Portability problems with external procedures (12.4.3.4)

There is a potential portability problem in a scoping unit that references an external procedure without explicitly declaring it to have the `EXTERNAL` attribute (5.3.8). On a different processor, the name of that procedure may be the name of a nonstandard intrinsic procedure and the processor would be permitted to interpret those procedure references as references to that intrinsic procedure. (On that processor, the program would also be viewed as not conforming to the standard because of the references to the nonstandard intrinsic procedure.) Declaration of the `EXTERNAL` attribute causes the references to be to the external procedure regardless of the availability of an intrinsic procedure with the same name. Note that declaration of the type of a procedure is not enough to make it external, even if the type is inconsistent with the type of the result of an intrinsic of the same name.

C.9.2 Procedures defined by means other than Fortran (12.6.3)

A processor is not required to provide any means other than Fortran for defining external procedures. Among the means that might be supported are the machine assembly language, other high level languages, the Fortran language extended with nonstandard features, and the Fortran language as supported by another Fortran processor (for example, a previously existing FORTRAN 77 processor).

Procedures defined by means other than Fortran are considered external procedures because their definitions are not in a Fortran program unit and because they are referenced using global names. The use of the term external should not be construed as any kind of restriction on the way in which these procedures may be defined. For example, if the means other than Fortran has its own facilities for internal and external procedures, it is permissible to use them. If the means other than Fortran can create an “internal” procedure with a global name, it is permissible for such an “internal” procedure to be considered by Fortran to be an external procedure. The means other than Fortran for defining external procedures, including any restrictions on the structure for organization of those procedures, are entirely processor dependent.

A Fortran processor may limit its support of procedures defined by means other than Fortran such that these procedures may affect entities in the Fortran environment only on the same basis as procedures written in Fortran. For example, it might prohibit the value of a local variable from being changed by a procedure reference unless that variable were one of the arguments to the procedure.

C.9.3 Procedure interfaces (12.4)

1 In FORTRAN 77, the interface to an external procedure was always deduced from the form of references
 2 to that procedure and any declarations of the procedure name in the referencing program unit. In this
 3 standard, features such as argument keywords and optional arguments make it impossible to deduce
 4 sufficient information about the dummy arguments from the nature of the actual arguments to be
 5 associated with them, and features such as array function results and pointer function results make
 6 necessary extensions to the declaration of a procedure that cannot be done in a way that would be
 7 analogous with the handling of such declarations in FORTRAN 77. Hence, mechanisms are provided
 8 through which all the information about a procedure's interface may be made available in a scoping
 9 unit that references it. A procedure whose interface shall be deduced as in FORTRAN 77 is described
 10 as having an implicit interface. A procedure whose interface is fully known is described as having an
 11 explicit interface.

12 A scoping unit is allowed to contain an interface body for a procedure that does not exist in the program,
 13 provided the procedure described is never referenced or used in any other way. The purpose of this rule is
 14 to allow implementations in which the use of a module providing interface bodies describing the interface
 15 of every routine in a library would not automatically cause each of those library routines to be a part of
 16 the program referencing the module. Instead, only those library procedures actually referenced would
 17 be a part of the program. (In implementation terms, the mere presence of an interface body would not
 18 generate an external reference in such an implementation.)

19 **C.9.4 Abstract interfaces (12.4) and procedure pointer components (4.5)**

20 This is an example of a library module providing lists of callbacks that the user may register and invoke.

```

21 MODULE callback_list_module
22   !
23   ! Type for users to extend with their own data, if they so desire
24   !
25   TYPE callback_data
26   END TYPE
27   !
28   ! Abstract interface for the callback procedures
29   !
30   ABSTRACT INTERFACE
31     SUBROUTINE callback_procedure(data)
32       IMPORT callback_data
33       CLASS(callback_data),OPTIONAL :: data
34     END SUBROUTINE
35   END INTERFACE
36   !
37   ! The callback list type.
38   !
39   TYPE callback_list
40     PRIVATE
41     CLASS(callback_record),POINTER :: first => NULL()
42   END TYPE
43   !
44   ! Internal: each callback registration creates one of these
45   !

```

```

1  TYPE,PRIVATE :: callback_record
2      PROCEDURE(callback_procedure),POINTER,NOPASS :: proc
3      CLASS(callback_record),POINTER :: next
4      CLASS(callback_data),POINTER :: data => NULL();
5  END TYPE
6  PRIVATE invoke,forward_invoke
7  CONTAINS
8  !
9  ! Register a callback procedure with optional data
10 !
11 SUBROUTINE register_callback(list, entry, data)
12     TYPE(callback_list),INTENT(INOUT) :: list
13     PROCEDURE(callback_procedure) :: entry
14     CLASS(callback_data),OPTIONAL :: data
15     TYPE(callback_record),POINTER :: new,last
16     ALLOCATE(new)
17     new%proc => entry
18     IF (PRESENT(data)) ALLOCATE(new%data,SOURCE=data)
19     new%next => list%first
20     list%first => new
21 END SUBROUTINE
22 !
23 ! Internal: Invoke a single callback and destroy its record
24 !
25 SUBROUTINE invoke(callback)
26     TYPE(callback_record),POINTER :: callback
27     IF (ASSOCIATED(callback%data) THEN
28         CALL callback%proc(list%first%data)
29         DEALLOCATE(callback%data)
30     ELSE
31         CALL callback%proc
32     END IF
33     DEALLOCATE(callback)
34 END SUBROUTINE
35 !
36 ! Call the procedures in reverse order of registration
37 !
38 SUBROUTINE invoke_callback_reverse(list)
39     TYPE(callback_list),INTENT(INOUT) :: list
40     TYPE(callback_record),POINTER :: next,current
41     current => list%first
42     NULLIFY(list%first)
43     DO WHILE (ASSOCIATED(current))
44         next => current%next
45         CALL invoke(current)

```

```
1      current => next
2      END DO
3      END SUBROUTINE
4      !
5      ! Internal: Forward mode invocation
6      !
7      RECURSIVE SUBROUTINE forward_invoke(callback)
8          IF (ASSOCIATED(callback%next)) CALL forward_invoke(callback%next)
9          CALL invoke(callback)
10     END SUBROUTINE
11     !
12     ! Call the procedures in forward order of registration
13     !
14     SUBROUTINE invoke_callback_forward(list)
15         TYPE(callback_list),INTENT(INOUT) :: list
16         IF (ASSOCIATED(list%first)) CALL forward_invoke(list%first)
17     END SUBROUTINE
18 END
```

19 C.9.5 Argument association and evaluation (12.5.2)

20 There is a significant difference between the argument association allowed in this standard and that
21 supported by FORTRAN 77 and FORTRAN 66. In FORTRAN 77 and 66, actual arguments were limited
22 to consecutive storage units. With the exception of assumed length character dummy arguments, the
23 structure imposed on that sequence of storage units was always determined in the invoked procedure and
24 not taken from the actual argument. Thus it was possible to implement FORTRAN 66 and FORTRAN 77
25 argument association by supplying only the location of the first storage unit (except for character argu-
26 ments, where the length would also have to be supplied). However, this standard allows arguments that
27 do not reside in consecutive storage locations (for example, an array section), and dummy arguments that
28 assume additional structural information from the actual argument (for example, assumed-shape dummy
29 arguments). Thus, the mechanism to implement the argument association allowed in this standard needs
30 to be more general.

31 Because there are practical advantages to a processor that can support references to and from procedures
32 defined by a FORTRAN 77 processor, requirements for explicit interfaces make it possible to determine
33 whether a simple (FORTRAN 66/FORTRAN 77) argument association implementation mechanism is suffi-
34 cient or whether the more general mechanism is necessary (12.4.2). Thus a processor can be implemented
35 whose procedures expect the simple mechanism to be used whenever the procedure's interface is one that
36 uses only FORTRAN 77 features and that expects the more general mechanism otherwise (for example, if
37 there are assumed-shape or optional arguments). At the point of reference, the appropriate mechanism
38 can be determined from the interface if it is explicit and can be assumed to be the simple mechanism
39 if it is not. Note that if the simple mechanism is determined to be what the procedure expects, it may
40 be necessary for the processor to allocate consecutive temporary storage for the actual argument, copy
41 the actual argument to the temporary storage, reference the procedure using the temporary storage in
42 place of the actual argument, copy the contents of temporary storage back to the actual argument, and
43 deallocate the temporary storage.

44 While this is the particular implementation method these rules were designed to support, it is not
45 the only one possible. For example, on some processors, it may be possible to implement the general
46 argument association in such a way that the information involved in FORTRAN 77 argument association
47 may be found in the same places and the "extra" information is placed so it does not disturb a procedure

1 expecting only FORTRAN 77 argument association. With such an implementation, argument association
2 could be translated without regard to whether the interface is explicit or implicit.

3 The provisions for expression evaluation give the processor considerable flexibility for obtaining expres-
4 sion values in the most efficient way possible. This includes not evaluating or only partially evaluating
5 an operand, for example, if the value of the expression can be determined otherwise (7.1.8.2). This
6 flexibility applies to function argument evaluation, including the order of argument evaluation, delay-
7 ing argument evaluation, and omitting argument evaluation. A processor may delay the evaluation of
8 an argument in a procedure reference until the execution of the procedure refers to the value of that
9 argument, provided delaying the evaluation of the argument does not otherwise affect the results of
10 the program. The processor may, with similar restrictions, entirely omit the evaluation of an argument
11 not referenced in the execution of the procedure. This gives processors latitude for optimization (for
12 example, for parallel processing).

13 C.9.6 Pointers and targets as arguments (12.5.2.5, 12.5.2.7, 12.5.2.8)

J3 internal note

Unresolved Technical Issue 082

The following paragraph takes no account of auto-targetting. Rewrite or delete.

14 If a dummy argument is declared to be a pointer, it may be matched only by an actual argument that
15 also is a pointer, and the characteristics of both arguments shall agree. A model for such an association is
16 that descriptor values of the actual pointer are copied to the dummy pointer. If the actual pointer has an
17 associated target, this target becomes accessible via the dummy pointer. If the dummy pointer becomes
18 associated with a different target during execution of the procedure, this target will be accessible via the
19 actual pointer after the procedure completes execution. If the dummy pointer becomes associated with
20 a local target that ceases to exist when the procedure completes, the actual pointer will be left dangling
21 in an undefined state. Such dangling pointers shall not be used.

22 When execution of a procedure completes, any pointer that remains defined and that is associated with
23 a dummy argument that has the TARGET attribute and is either a scalar or an assumed-shape array,
24 remains associated with the corresponding actual argument if the actual argument has the TARGET
25 attribute and is not an array section with a vector subscript.

```

26 REAL, POINTER      :: PBEST
27 REAL, TARGET       :: B (10000)
28 CALL BEST (PBEST, B)          ! Upon return PBEST is associated
29 ...                          ! with the 'best' element of B
30 CONTAINS
31   SUBROUTINE BEST (P, A)
32     REAL, POINTER, INTENT (OUT) :: P
33     REAL, TARGET, INTENT (IN)  :: A (:)
34     ...                        ! Find the 'best' element A(I)
35     P => A (I)
36   RETURN
37 END SUBROUTINE BEST
38 END

```

39 When procedure BEST completes, the pointer PBEST is associated with an element of B.

40 An actual argument without the TARGET attribute can become associated with a dummy argument

1 with the TARGET attribute. This permits pointers to become associated with the dummy argument
2 during execution of the procedure that contains the dummy argument. For example:

```

3  INTEGER LARGE(100,100)
4  CALL SUB (LARGE)
5  ...
6  CALL SUB ()
7  CONTAINS
8  SUBROUTINE SUB(ARG)
9      INTEGER, TARGET, OPTIONAL :: ARG(100,100)
10     INTEGER, POINTER, DIMENSION(:,:) :: PARG
11     IF (PRESENT(ARG)) THEN
12         PARG => ARG
13     ELSE
14         ALLOCATE (PARG(100,100))
15         PARG = 0
16     ENDIF
17     ... ! Code with lots of references to PARG
18     IF (.NOT. PRESENT(ARG)) DEALLOCATE(PARG)
19 END SUBROUTINE SUB
20 END

```

21 Within subroutine SUB the pointer PARG is either associated with the dummy argument ARG or it is
22 associated with an allocated target. The bulk of the code can reference PARG without further calls to
23 the PRESENT intrinsic.

24 C.9.7 Polymorphic Argument Association (12.5.2.10)

25 The following example illustrates polymorphic argument association rules using the derived types defined
26 in Note 4.59.

```

27  TYPE(POINT) :: T2
28  TYPE(COLOR_POINT) :: T3
29  CLASS(POINT) :: P2
30  CLASS(COLOR_POINT) :: P3
31  ! Dummy argument is polymorphic and actual argument is of fixed type
32  SUBROUTINE SUB2 ( X2 ); CLASS(POINT) :: X2; ...
33  SUBROUTINE SUB3 ( X3 ); CLASS(COLOR_POINT) :: X3; ...
34
35  CALL SUB2 ( T2 ) ! Valid -- The declared type of T2 is the same as the
36                   !           declared type of X2.
37  CALL SUB2 ( T3 ) ! Valid -- The declared type of T3 is extended from
38                   !           the declared type of X2.
39  CALL SUB3 ( T2 ) ! Invalid -- The declared type of T2 is neither the
40                   !           same as nor extended from the declared type
41                   !           type of X3.

```

```

1
2 CALL SUB3 ( T3 ) ! Valid -- The declared type of T3 is the same as the
3           !           declared type of X3.
4 ! Actual argument is polymorphic and dummy argument is of fixed type
5 SUBROUTINE TUB2 ( D2 ); TYPE(POINT) :: D2; ...
6 SUBROUTINE TUB3 ( D3 ); TYPE(COLOR_POINT) :: D3; ...
7
8 CALL TUB2 ( P2 ) ! Valid -- The declared type of P2 is the same as the
9           !           declared type of D2.
10 CALL TUB2 ( P3 ) ! Invalid -- The declared type of P3 differs from the
11           !           declared type of D2.
12 CALL TUB2 ( P3%POINT ) ! Valid alternative to the above
13 CALL TUB3 ( P2 ) ! Invalid -- The declared type of P2 differs from the
14           !           declared type of D3.
15 SELECT TYPE ( P2 ) ! Valid conditional alternative to the above
16 CLASS IS ( COLOR_POINT ) ! Works if the dynamic type of P2 is the same
17   CALL TUB3 ( P2 )       ! as the declared type of D3, or a type
18                           ! extended therefrom.
19 CLASS DEFAULT
20                           ! Cannot work if not.
21 END SELECT
22 CALL TUB3 ( P3 ) ! Valid -- The declared type of P3 is the same as the
23           !           declared type of D3.
24 ! Both the actual and dummy arguments are of polymorphic type.
25 CALL SUB2 ( P2 ) ! Valid -- The declared type of P2 is the same as the
26           !           declared type of X2.
27 CALL SUB2 ( P3 ) ! Valid -- The declared type of P3 is extended from
28           !           the declared type of X2.
29 CALL SUB3 ( P2 ) ! Invalid -- The declared type of P2 is neither the
30           !           same as nor extended from the declared
31           !           type of X3.
32 SELECT TYPE ( P2 ) ! Valid conditional alternative to the above
33 CLASS IS ( COLOR_POINT ) ! Works if the dynamic type of P2 is the
34   CALL SUB3 ( P2 )       ! same as the declared type of X3, or a
35                           ! type extended therefrom.
36 CLASS DEFAULT
37                           ! Cannot work if not.
38 END SELECT
39 CALL SUB3 ( P3 ) ! Valid -- The declared type of P3 is the same as the
40           !           declared type of X3.

```

41 C.10 Clause 13 notes

1 C.10.1 Module for THIS_IMAGE and IMAGE_INDEX

2 The intrinsics THIS_IMAGE (CO_ARRAY) and IMAGE_INDEX (CO_ARRAY, SUB) cannot be written
3 in Fortran since CO_ARRAY may be of any type and THIS_IMAGE (CO_ARRAY) needs to know the
4 index of the image on which the code is running.

5 As an example, here are simple versions that require the co-bounds to be specified as integer arrays and
6 require the image index for THIS_IMAGE (CO_ARRAY).

```

7 MODULE index
8 CONTAINS
9     INTEGER FUNCTION image_index(lbound, ubound, sub)
10        INTEGER, INTENT(IN) :: lbound(:), ubound(:), sub(:)
11        INTEGER              :: i, n
12        n = SIZE(sub)
13        image_index = sub(n) - lbound(n)
14        DO i = n-1, 1, -1
15            image_index = image_index*(ubound(i)-lbound(i)+1) + sub(i) - lbound(i)
16        END DO
17        image_index = image_index + 1
18    END FUNCTION image_index
19
20    INTEGER FUNCTION this_image(lbound, ubound, me) RESULT(sub)
21        INTEGER, INTENT(IN) :: lbound(:), ubound(:), me
22        INTEGER              :: sub(SIZE(lbound))
23        INTEGER              :: extent, i, m, ml, n
24        n = SIZE(sub)
25        m = me - 1
26        DO i = 1, n-1
27            extent = ubound(i) - lbound(i) + 1
28            ml = m
29            m = m/extent
30            sub(i) = ml - m*extent + lbound(i)
31        END DO
32        sub(n) = m + lbound(n)
33    END FUNCTION this_IMAGE
34 END MODULE index

```

35 C.10.2 Collective co-array subroutine variations

36 The subroutines of 13.5.15 return an array of the same shape as the given co-array after having applied
37 an operation on the images involved. Simple routines can be written to also apply the operation to the
38 elements of the co-array on an image. Various versions of a global sum can be programmed, for example:

```

39 MODULE global_sum_module
40     INTRINSIC, PRIVATE :: CO_SUM, SIZE, SUM
41 CONTAINS

```

```

1  REAL FUNCTION global_sum(array)
2      REAL,INTENT(IN) :: array(:,:)[*]
3      REAL,SAVE      :: temp[*]
4      temp = SUM(array)          ! Sum on the executing image
5      CALL CO_SUM(temp, global_sum)
6  END FUNCTION global_sum
7
8  REAL FUNCTION global_sum_mask(array, mask)
9      REAL,INTENT(IN)  :: array(:,:)[*]
10     LOGICAL,INTENT(IN) :: mask(:,:)
11     REAL,SAVE        :: temp[*]
12     temp = SUM(array, MASK=mask) ! Sum on the executing image
13     CALL CO_SUM(temp, global_sum_mask)
14 END FUNCTION global_sum_mask
15
16 FUNCTION global_sum_dim(array, dim)
17     REAL, INTENT(IN)  :: array(:,:)[*]
18     INTEGER, INTENT(IN) :: dim
19     REAL, ALLOCATABLE :: global_sum_dim(:)
20     REAL, ALLOCATABLE :: temp(:):[]
21     ALLOCATE (global_sum_dim(SIZE(array, 3-dim)))
22     ALLOCATE (      temp(SIZE(array, 3-dim))[*])
23     temp = SUM(array, dim)      ! Sum of the local part of the co-array.
24     CALL CO_SUM(temp, global_sum_dim)
25 END FUNCTION global_sum_dim
26 END MODULE global_sum_module

```

27 C.11 Clause 15 notes

28 C.11.1 Runtime environments

29 This standard allows programs to contain procedures defined by means other than Fortran. That raises
30 the issues of initialization of and interaction between the runtime environments involved.

31 Implementations are free to solve these issues as they see fit, provided that

- 32 (1) heap allocation/deallocation (e.g., (DE)ALLOCATE in a Fortran subprogram and mal-
33 loc/free in a C function) can be performed without interference,
- 34 (2) input/output to and from external files can be performed without interference, as long as
35 procedures defined by different means do not do input/output with the same external file,
- 36 (3) input/output preconnections exist as required by the respective standards, and
- 37 (4) initialized data is initialized according to the respective standards.

38 C.11.2 Examples of Interoperation between Fortran and C Functions

39 The following examples illustrate the interoperation of Fortran and C functions. Two examples are
40 shown: one of Fortran calling C, and one of C calling Fortran. In each of the examples, the correspon-
41 dences of Fortran actual arguments, Fortran dummy arguments, and C formal parameters are described.

1 **C.11.2.1 Example of Fortran calling C**

2 C Function Prototype:

```
3     int C_Library_Function(void* sendbuf, int sendcount,
4         int *recvcounts);
```

5 Fortran Modules:

```
6     MODULE FTN_C_1
7         USE, INTRINSIC :: ISO_C_BINDING
8     END MODULE FTN_C_1
9
10    MODULE FTN_C_2
11        INTERFACE
12
13            INTEGER (C_INT) FUNCTION C_LIBRARY_FUNCTION &
14                (SENDBUF, SENDCOUNT, RECVCOUNTS) &
15                BIND(C, NAME='C_Library_Function')
16                USE FTN_C_1
17                IMPLICIT NONE
18                TYPE (C_PTR), VALUE :: SENDBUF
19                INTEGER (C_INT), VALUE :: SENDCOUNT
20                TYPE (C_PTR), VALUE :: RECVCOUNTS
21            END FUNCTION C_LIBRARY_FUNCTION
22        END INTERFACE
23    END MODULE FTN_C_2
```

23 The module FTN_C.2 contains the declaration of the Fortran dummy arguments, which correspond to
 24 the C formal parameters. The intrinsic module ISO_C_BINDING is referenced in the module FTN_C.1.
 25 The NAME specifier is used in the BIND attribute in order to handle the case-sensitive name change
 26 between Fortran and C from 'C_LIBRARY_FUNCTION' to 'C_Library_Function'. See also Note 12.46.

27 The first C formal parameter is the pointer to void `sendbuf`, which corresponds to the Fortran dummy
 28 argument `SENDBUF`, which has the type `C_PTR` and the `VALUE` attribute.

29 The second C formal parameter is the int `sendcount`, which corresponds to the Fortran dummy argument
 30 `SENDCOUNT`, which has the type `INTEGER(C_INT)` and the `VALUE` attribute.

31 The third C formal parameter is the pointer to int `recvcounts`, which corresponds to the Fortran dummy
 32 argument `RECVCOUNTS`, which has the type `C_PTR` and the `VALUE` attribute.

33 Fortran Calling Sequence:

```
34     USE, INTRINSIC :: ISO_C_BINDING, ONLY: C_INT, C_FLOAT, C_LOC
35     USE FTN_C_2
36     ...
37     REAL (C_FLOAT), TARGET :: SEND(100)
```

```

1
2     INTEGER (C_INT)           :: SENDCOUNT
3     INTEGER (C_INT), ALLOCATABLE, TARGET :: RECVCOUNTS(100)
4     ...
5     ALLOCATE( RECVCOUNTS(100) )
6     ...
7     CALL C_LIBRARY_FUNCTION(C_LOC(SEND), SENDCOUNT, &
8     C_LOC(RECVCOUNTS))
9     ...

```

10 The preceding code contains the declaration of the Fortran actual arguments associated with the above-
11 listed Fortran dummy arguments.

12 The first Fortran actual argument is the address of the first element of the array SEND, which has the
13 type REAL(C_FLOAT) and the TARGET attribute. This address is returned by the intrinsic function
14 C_LOC. This actual argument is associated with the Fortran dummy argument SENDBUF, which has
15 the type C_PTR and the VALUE attribute.

16 The second Fortran actual argument is SENDCOUNT of type INTEGER(C_INT), which is associated
17 with the Fortran dummy argument SENDCOUNT, which has the type INTEGER(C_INT) and the
18 VALUE attribute.

19 The third Fortran actual argument is the address of the first element of the allocatable array RECV-
20 COUNTS, with has the type REAL(C_FLOAT) and the TARGET attribute. This address is returned
21 by the intrinsic function C_LOC. This actual argument is associated with the Fortran dummy argument
22 RECVCOUNTS, which has the type C_PTR and the VALUE attribute.

23 C.11.2.2 Example of C calling Fortran

24 Fortran Code:

```

25 SUBROUTINE SIMULATION(ALPHA, BETA, GAMMA, DELTA, ARRAYS) BIND(C)
26     USE, INTRINSIC :: ISO_C_BINDING
27     IMPLICIT NONE
28     INTEGER (C_LONG), VALUE           :: ALPHA
29     REAL (C_DOUBLE), INTENT(INOUT)    :: BETA
30     INTEGER (C_LONG), INTENT(OUT)     :: GAMMA
31     REAL (C_DOUBLE), DIMENSION(*), INTENT(IN) :: DELTA
32     TYPE, BIND(C) :: PASS
33         INTEGER (C_INT) :: LENC, LENF
34         TYPE (C_PTR)   :: C, F
35     END TYPE PASS
36     TYPE (PASS), INTENT(INOUT) :: ARRAYS
37     REAL (C_FLOAT), ALLOCATABLE, TARGET, SAVE :: ETA(:)
38     REAL (C_FLOAT), POINTER :: C_ARRAY(:)
39     ...
40     ! Associate C_ARRAY with an array allocated in C
41     CALL C_F_POINTER (ARRAYS%C, C_ARRAY, (/ARRAYS%LENC/) )
42     ...

```

```
1   ! Allocate an array and make it available in C
2   ARRAYS%LENF = 100
3   ALLOCATE (ETA(ARRAYS%LENF))
4   ARRAYS%F = C_LOC(ETA)
5   ...
6 END SUBROUTINE SIMULATION

7 C Struct Declaration

8   struct pass {int lenc, lenf; float *c, *f;};

9 C Function Prototype:

10  void simulation(long alpha, double *beta, long *gamma,
11  double delta[], struct pass *arrays);

12 C Calling Sequence:

13  simulation(alpha, &beta, &gamma, delta, &arrays);
```

14 The above-listed Fortran code specifies a subroutine `SIMULATION`. This subroutine corresponds to the
15 C void function `simulation`.

16 The Fortran subroutine references the intrinsic module `ISO_C_BINDING`.

17 The first Fortran dummy argument of the subroutine is `ALPHA`, which has the type `INTEGER(C_-`
18 `LONG)` and the attribute `VALUE`. This dummy argument corresponds to the C formal parameter
19 `alpha`, which is a long. The actual C parameter is also a long.

20 The second Fortran dummy argument of the subroutine is `BETA`, which has the type `REAL(C_-`
21 `DOUBLE)` and the `INTENT(INOUT)` attribute. This dummy argument corresponds to the C formal
22 parameter `beta`, which is a pointer to double. An address is passed as the actual parameter in the C
23 calling sequence.

24 The third Fortran dummy argument of the subroutine is `GAMMA`, which has the type `INTEGER(C_-`
25 `LONG)` and the `INTENT(OUT)` attribute. This dummy argument corresponds to the C formal param-
26 eter `gamma`, which is a pointer to long. An address is passed as the actual parameter in the C calling
27 sequence.

28 The fourth Fortran dummy argument is the assumed-size array `DELTA`, which has the type `REAL`
29 `(C_DOUBLE)` and the attribute `INTENT(IN)`. This dummy argument corresponds to the C formal
30 parameter `delta`, which is a double array. The actual C parameter is also a double array.

31 The fifth Fortran dummy argument is `ARRAYS`, which is a structure for accessing an array allocated
32 in C and an array allocated in Fortran. The lengths of these arrays are held in the components `LENC`
33 and `LENF`; their C addresses are held in components `C` and `F`.

34 C.11.2.3 Example of calling C functions with noninteroperable data

35 Many Fortran processors support 16-byte real numbers, which might not be supported by the C processor.
36 Assume a Fortran programmer wants to use a C procedure from a message passing library for an array
37 of these reals. The C prototype of this procedure is

1 void ProcessBuffer(void *buffer, int n_bytes);

2 with the corresponding Fortran interface

```

3 USE, INTRINSIC :: ISO_C_BINDING
4
5 INTERFACE
6     SUBROUTINE PROCESS_BUFFER(BUFFER,N_BYTES) BIND(C,NAME="ProcessBuffer")
7         IMPORT :: C_PTR, C_INT
8         TYPE(C_PTR), VALUE :: BUFFER ! The 'C address' of the array buffer
9         INTEGER(C_INT), VALUE :: N_BYTES ! Number of bytes in buffer
10    END SUBROUTINE PROCESS_BUFFER
11 END INTERFACE

```

12 This may be done using C_LOC if the particular Fortran processor specifies that C_LOC returns an
13 appropriate address:

```

14 REAL(R_QUAD), DIMENSION(:), ALLOCATABLE, TARGET :: QUAD_ARRAY
15 ...
16 CALL PROCESS_BUFFER(C_LOC(QUAD_ARRAY), INT(16*SIZE(QUAD_ARRAY)),C_INT))
17 ! One quad real takes 16 bytes on this processor

```

18 C.11.2.4 Example of opaque communication between C and Fortran

19 The following example demonstrates how a Fortran processor can make a modern OO random number
20 generator written in Fortran available to a C program:

```

21 USE, INTRINSIC :: ISO_C_BINDING
22 ! Assume this code is inside a module
23
24 TYPE RANDOM_STREAM
25 ! A (uniform) random number generator (URNG)
26 CONTAINS
27 PROCEDURE(RANDOM_UNIFORM), DEFERRED, PASS(STREAM) :: NEXT
28 ! Generates the next number from the stream
29 END TYPE RANDOM_STREAM
30
31 ABSTRACT INTERFACE
32 ! Abstract interface of Fortran URNG
33 SUBROUTINE RANDOM_UNIFORM(STREAM, NUMBER)
34 IMPORT :: RANDOM_STREAM, C_DOUBLE
35 CLASS(RANDOM_STREAM), INTENT(INOUT) :: STREAM
36 REAL(C_DOUBLE), INTENT(OUT) :: NUMBER
37 END SUBROUTINE RANDOM_UNIFORM
38 END INTERFACE

```

1 A polymorphic object of base type RANDOM_STREAM is not interoperable with C. However, we can
 2 make such a random number generator available to C by packaging it inside another nonpolymorphic,
 3 nonparameterized derived type:

```
4 TYPE :: URNG_STATE ! No BIND(C), as this type is not interoperable
5   CLASS(RANDOM_STREAM), ALLOCATABLE :: STREAM
6 END TYPE URNG_STATE
```

7 The following two procedures will enable a C program to use our Fortran uniform random number
 8 generator:

```
9 ! Initialize a uniform random number generator:
10 SUBROUTINE INITIALIZE_URNG(STATE_HANDLE, METHOD) &
11     BIND(C, NAME="InitializeURNG")
12     TYPE(C_PTR), INTENT(OUT) :: STATE_HANDLE
13     ! An opaque handle for the URNG
14     CHARACTER(C_CHAR), DIMENSION(*), INTENT(IN) :: METHOD
15     ! The algorithm to be used
16
17     TYPE(URNG_STATE), POINTER :: STATE
18     ! An actual URNG object
19
20     ALLOCATE(STATE)
21     ! There needs to be a corresponding finalization
22     ! procedure to avoid memory leaks, not shown in this example
23     ! Allocate STATE%STREAM with a dynamic type depending on METHOD
24     ...
25     STATE_HANDLE=C_LOC(STATE)
26     ! Obtain an opaque handle to return to C
27 END SUBROUTINE INITIALIZE_URNG
28
29 ! Generate a random number:
30 SUBROUTINE GENERATE_UNIFORM(STATE_HANDLE, NUMBER) &
31     BIND(C, NAME="GenerateUniform")
32     TYPE(C_PTR), INTENT(IN), VALUE :: STATE_HANDLE
33     ! An opaque handle: Obtained via a call to INITIALIZE_URNG
34     REAL(C_DOUBLE), INTENT(OUT) :: NUMBER
35
36     TYPE(URNG_STATE), POINTER :: STATE
37     ! A pointer to the actual URNG
38
39     CALL C_F_POINTER(CPTR=STATE_HANDLE, FPTR=STATE)
40     ! Convert the opaque handle into a usable pointer
41     CALL STATE%STREAM%NEXT(NUMBER)
42     ! Use the type-bound procedure NEXT to generate NUMBER
43 END SUBROUTINE GENERATE_UNIFORM
```

1 C.12 Clause 16 notes

2 C.12.1 Examples of host association (16.5.1.4)

3 The first two examples are examples of valid host association. The third example is an example of invalid
4 host association.

5 Example 1:

```
6 PROGRAM A
7     INTEGER I, J
8     ...
9 CONTAINS
10    SUBROUTINE B
11        INTEGER I ! Declaration of I hides
12                ! program A's declaration of I
13        ...
14        I = J     ! Use of variable J from program A
15                ! through host association
16    END SUBROUTINE B
17 END PROGRAM A
```

18 Example 2:

```
19 PROGRAM A
20     TYPE T
21     ...
22 END TYPE T
23     ...
24 CONTAINS
25    SUBROUTINE B
26        IMPLICIT TYPE (T) (C) ! Refers to type T declared below
27                                ! in subroutine B, not type T
28                                ! declared above in program A
29        ...
30        TYPE T
31        ...
32    END TYPE T
33        ...
34 END SUBROUTINE B
35 END PROGRAM A
```

36 Example 3:

```
37 PROGRAM Q
38     REAL (KIND = 1) :: C
```

```

1      ...
2  CONTAINS
3      SUBROUTINE R
4          REAL (KIND = KIND (C)) :: D  ! Invalid declaration
5                                          ! See below
6          REAL (KIND = 2) :: C
7      ...
8      END SUBROUTINE R
9  END PROGRAM Q

```

10 In the declaration of D in subroutine R, the use of C would refer to the declaration of C in subroutine
 11 R, not program Q. However, it is invalid because the declaration of C is required to occur before it is
 12 used in the declaration of D (7.1.7).

13 **C.12.2 Rules ensuring unambiguous generics (12.4.3.3.4)**

14 The rules in 12.4.3.3.4 are intended to ensure

- 15 • that it is possible to reference each specific procedure in the generic collection,
- 16 • that for any valid reference to the generic procedure, the determination of the specific procedure
 17 referenced is unambiguous, and
- 18 • that the determination of the specific procedure referenced can be made before execution of the
 19 program begins (during compilation).

20 Specific procedures are distinguished by fixed properties of their arguments, specifically type, kind type
 21 parameters, and rank. A valid reference to one procedure in a generic collection will differ from another
 22 because it has an argument that the other cannot accept, because it is missing an argument that the
 23 other requires, or because one of these fixed properties is different.

24 Although the declared type of a data entity is a fixed property, polymorphic variables allow for a
 25 limited degree of type mismatch between dummy arguments and actual arguments, so the requirement
 26 for distinguishing two dummy arguments is type incompatibility, not merely different types. (This is
 27 illustrated in the BAD6 example later in this note.)

28 That same limited type mismatch means that two dummy arguments that are not type incompatible
 29 can be distinguished on the basis of the values of the kind type parameters they have in common; if one
 30 of them has a kind type parameter that the other does not, that is irrelevant in distinguishing them.

31 Rank is a fixed property, but some forms of array dummy arguments allow rank mismatches when a
 32 procedure is referenced by its specific name. In order to allow rank to always be usable in distinguishing
 33 generics, such rank mismatches are disallowed for those arguments when the procedure is referenced as
 34 part of a generic. Additionally, the fact that elemental procedures can accept array arguments is not
 35 taken into account when applying these rules, so apparent ambiguity between elemental and nonelemental
 36 procedures is possible; in such cases, the reference is interpreted as being to the nonelemental procedure.

37 For procedures referenced as operators or defined-assignment, syntactically distinguished arguments are
 38 mapped to specific positions in the argument list, so the rule for distinguishing such procedures is that
 39 it be possible to distinguish the arguments at one of the argument positions.

40 For user-defined derived-type input/output procedures, only the `dtv` argument corresponds to something
 41 explicitly written in the program, so it is the `dtv` that is required to be distinguished. Because `dtv`

1 arguments are required to be scalar, they cannot differ in rank. Thus this rule effectively involves only
2 type and kind type parameters.

3 For generic procedures identified by names, the rules are more complicated because optional arguments
4 may be omitted and because arguments may be specified either positionally or by name.

5 In the special case of type-bound procedures with passed-object dummy arguments, the passed-object
6 argument is syntactically distinguished in the reference, so rule (2) can be applied. The type of passed-
7 object arguments is constrained in ways that prevent passed-object arguments in the same scoping unit
8 from being type incompatible. Thus this rule effectively involves only kind type parameters and rank.

9 The primary means of distinguishing named generics is rule (3). The most common application of that
10 rule is a single argument satisfying both (3a) and (3b):

```
11         INTERFACE GOOD1
12             FUNCTION F1A(X)
13                 REAL :: F1A,X
14             END FUNCTION F1A
15             FUNCTION F1B(X)
16                 INTEGER :: F1B,X
17             END FUNCTION F1B
18         END INTERFACE GOOD1
```

19 Whether one writes `GOOD1(1.0)` or `GOOD1(X=1.0)`, the reference is to `F1A` because `F1B` would require an
20 integer argument whereas these references provide the real constant `1.0`.

21 This example and those that follow are expressed using interface bodies, with type as the distinguishing
22 property. This was done to make it easier to write and describe the examples. The principles being
23 illustrated are equally applicable when the procedures get their explicit interfaces in some other way or
24 when kind type parameters or rank are the distinguishing property.

25 Another common variant is the argument that satisfies (3a) and (3b) by being required in one specific
26 and completely missing in the other:

```
27         INTERFACE GOOD2
28             FUNCTION F2A(X)
29                 REAL :: F2A,X
30             END FUNCTION F2A
31             FUNCTION F2B(X,Y)
32                 COMPLEX :: F2B
33                 REAL :: X,Y
34             END FUNCTION F2B
35         END INTERFACE GOOD2
```

36 Whether one writes `GOOD2(0.0,1.0)`, `GOOD2(0.0,Y=1.0)`, or `GOOD2(Y=1.0,X=0.0)`, the reference is to
37 `F2B`, because `F2A` has no argument in the second position or with the name `Y`. This approach is used as
38 an alternative to optional arguments when one wants a function to have different result type, kind type
39 parameters, or rank, depending on whether the argument is present. In many of the intrinsic functions,
40 the `DIM` argument works this way.

41 It is possible to construct cases where different arguments are used to distinguish positionally and by
42 name:

```

1      INTERFACE GOOD3
2          SUBROUTINE S3A(W,X,Y,Z)
3              REAL :: W,Y
4              INTEGER :: X,Z
5          END SUBROUTINE S3A
6          SUBROUTINE S3B(X,W,Z,Y)
7              REAL :: W,Z
8              INTEGER :: X,Y
9          END SUBROUTINE S3B
10     END INTERFACE GOOD3

```

11 If one writes `GOOD3(1.0,2,3.0,4)` to reference `S3A`, then the third and fourth arguments are consistent
12 with a reference to `S3B`, but the first and second are not. If one switches to writing the first two
13 arguments as keyword arguments in order for them to be consistent with a reference to `S3B`, the latter
14 two arguments must also be written as keyword arguments, `GOOD3(X=2,W= 1.0,Z=4,Y=3.0)`, and the
15 named arguments `Y` and `Z` are distinguished.

16 The ordering requirement in rule (3) is critical:

```

17     INTERFACE BAD4 ! this interface is invalid !
18         SUBROUTINE S4A(W,X,Y,Z)
19             REAL :: W,Y
20             INTEGER :: X,Z
21         END SUBROUTINE S4A
22         SUBROUTINE S4B(X,W,Z,Y)
23             REAL :: X,Y
24             INTEGER :: W,Z
25         END SUBROUTINE S4B
26     END INTERFACE BAD4

```

27 In this example, the positionally distinguished arguments are `Y` and `Z`, and it is `W` and `X` that are
28 distinguished by name. In this order it is possible to write `BAD4(1.0,2,Y=3.0,Z=4)`, which is a valid
29 reference for both `S4A` and `S4B`.

30 Rule (1) can be used to distinguish some cases that are not covered by rule (3):

```

31     INTERFACE GOOD5
32         SUBROUTINE S5A(X)
33             REAL :: X
34         END SUBROUTINE S5A
35         SUBROUTINE S5B(Y,X)
36             REAL :: Y,X
37         END SUBROUTINE S5B
38     END INTERFACE GOOD5

```

39 In attempting to apply rule (3), position 2 and name `Y` are distinguished, but they are in the wrong
40 order, just like the `BAD4` example. However, when we try to construct a similarly ambiguous reference,
41 we get `GOOD5(1.0,X=2.0)`, which can't be a reference to `S5A` because it would be attempting to associate

1 two different actual arguments with the dummy argument X. Rule (3) catches this case by recognizing
2 that S5B requires two real arguments, and S5A cannot possibly accept more than one.

3 The application of rule (1) becomes more complicated when extensible types are involved. If FRUIT is
4 an extensible type, PEAR and APPLE are extensions of FRUIT, and BOSC is an extension of PEAR, then

```
5     INTERFACE BAD6 ! this interface is invalid !
6         SUBROUTINE S6A(X,Y)
7             CLASS(PEAR) :: X,Y
8         END SUBROUTINE S6A
9         SUBROUTINE S6B(X,Y)
10            CLASS(FRUIT) :: X
11            CLASS(BOSC) :: Y
12        END SUBROUTINE S6B
13    END INTERFACE BAD6
```

14 might, at first glance, seem distinguishable this way, but because of the limited type mismatching allowed,
15 BAD6(A_PEAR,A_BOSC) is a valid reference to both S6A and S6B.

16 It is important to try rule (1) for each type present:

```
17     INTERFACE GOOD7
18         SUBROUTINE S7A(X,Y,Z)
19             CLASS(PEAR) :: X,Y,Z
20         END SUBROUTINE S7A
21         SUBROUTINE S7B(X,Z,W)
22             CLASS(FRUIT) :: X
23             CLASS(BOSC) :: Z
24             CLASS(APPLE),OPTIONAL :: W
25         END SUBROUTINE S7B
26     END INTERFACE GOOD7
```

27 Looking at the most general type, S7A has a minimum and maximum of 3 FRUIT arguments, while S7B
28 has a minimum of 2 and a maximum of three. Looking at the most specific, S7A has a minimum of 0
29 and a maximum of 3 BOSC arguments, while S7B has a minimum of 1 and a maximum of 2. However,
30 when we look at the intermediate, S7A has a minimum and maximum of 3 PEAR arguments, while S7B
31 has a minimum of 1 and a maximum of 2. Because S7A's minimum exceeds S7B's maximum, they can
32 be distinguished.

33 In identifying the minimum number of arguments with a particular set of properties, we exclude optional
34 arguments and test TKR compatibility, so the corresponding actual arguments are required to have
35 those properties. In identifying the maximum number of arguments with those properties, we include
36 the optional arguments and test not distinguishable, so we include actual arguments which could have
37 those properties but are not required to have them.

38 These rules are sufficient to ensure that references to procedures that meet them are unambiguous, but
39 there remain examples that fail to meet these rules but which can be shown to be unambiguous:

```
40     INTERFACE BAD8 ! this interface is invalid !
```

```

1      ! despite the fact that it is unambiguous !
2      SUBROUTINE S8A(X,Y,Z)
3          REAL,OPTIONAL :: X
4          INTEGER :: Y
5          REAL :: Z
6      END SUBROUTINE S8A
7      SUBROUTINE S8B(X,Z,Y)
8          INTEGER,OPTIONAL :: X
9          INTEGER :: Z
10         REAL :: Y
11     END SUBROUTINE S8B
12 END INTERFACE BAD8

```

13 This interface fails rule (3) because there are no required arguments that can be distinguished from the
14 positionally corresponding argument, but in order for the mismatch of the optional arguments not to
15 be relevant, the later arguments must be specified as keyword arguments, so distinguishing by name
16 does the trick. This interface is nevertheless invalid so a standard- conforming Fortran processor is not
17 required to do such reasoning. The rules to cover all cases are too complicated to be useful.

18 The real data objects that would be valid arguments for **S9A** are entirely disjoint from procedures that
19 are valid arguments to **S9B** and **S9C**, and the procedures that valid arguments for **S9B** are disjoint from
20 the procedures that are valid arguments to **S9C** because the former are required to accept real arguments
21 and the latter integer arguments. Again, this interface is invalid, so a standard-conforming Fortran
22 processor need not examine such properties when deciding whether a generic collection is valid. Again,
23 the rules to cover all cases are too complicated to be useful.

24 **C.13 Array feature notes**

25 **C.13.1 Summary of features**

26 This subclause is a summary of the principal array features.

27 **C.13.1.1 Whole array expressions and assignments (7.4.1.2, 7.4.1.3)**

28 An important feature is that whole array expressions and assignments are permitted. For example, the
29 statement

```
30 A = B + C * SIN (D)
```

31 where A, B, C, and D are arrays of the same shape, is permitted. It is interpreted element-by-element;
32 that is, the sine function is taken on each element of D, each result is multiplied by the corresponding
33 element of C, added to the corresponding element of B, and assigned to the corresponding element of
34 A. Functions, including user-written functions, may be arrays and may be generic with scalar versions.
35 All arrays in an expression or across an assignment shall conform; that is, have exactly the same shape
36 (number of dimensions and extents in each dimension), but scalars may be included freely and these are
37 interpreted as being broadcast to a conforming array. Expressions are evaluated before any assignment
38 takes place.

39 **C.13.1.2 Array sections (2.4.5, 6.2.2.3)**

40 Whenever whole arrays may be used, it is also possible to use subarrays called “sections”. For example:

1 A (:, 1:N, 2, 3:1:-1)

2 consists of a subarray containing the whole of the first dimension, positions 1 to N of the second dimension, position 2 of the third dimension and positions 1 to 3 in reverse order of the fourth dimension.
3 This is an artificial example chosen to illustrate the different forms. Of course, a common use may be
4 to select a row or column of an array, for example:
5

6 A (:, J)

7 **C.13.1.3 WHERE statement (7.4.3)**

8 The WHERE statement applies a conforming logical array as a mask on the individual operations in the
9 expression and in the assignment. For example:

10 WHERE (A > 0) B = LOG (A)

11 takes the logarithm only for positive components of A and makes assignments only in these positions.

12 The WHERE statement also has a block form (WHERE construct).

13 **C.13.1.4 Automatic arrays and allocatable variables (5.2, 5.3.7.4)**

14 Two features useful for writing modular software are automatic arrays, created on entry to a subprogram
15 and destroyed on return, and allocatable variables, including arrays whose rank is fixed but whose actual
16 size and lifetime is fully under the programmer's control through explicit ALLOCATE and DEALLO-
17 CATE statements. The declarations

18 SUBROUTINE X (N, A, B)

19 REAL WORK (N, N); REAL, ALLOCATABLE :: HEAP (:, :)

20 specify an automatic array WORK and an allocatable array HEAP. Note that a stack is an adequate
21 storage mechanism for the implementation of automatic arrays, but a heap will be needed for some
22 allocatable variables.

23 **C.13.1.5 Array constructors (4.7)**

24 Arrays, and in particular array constants, may be constructed with array constructors exemplified by:

25 (/ 1.0, 3.0, 7.2 /)

26 which is a rank-one array of size 3,

27 (/ (1.3, 2.7, L = 1, 10), 7.1 /)

28 which is a rank-one array of size 21 and contains the pair of real constants 1.3 and 2.7 repeated 10 times
29 followed by 7.1, and

30 (/ (I, I = 1, N) /)

31 which contains the integers 1, 2, ..., N. Only rank-one arrays may be constructed in this way, but higher
32 dimensional arrays may be made from them by means of the intrinsic function RESHAPE.

33 **C.13.2 Examples**

34 The array features have the potential to simplify the way that almost any array-using program is con-
35 ceived and written. Many algorithms involving arrays can now be written conveniently as a series of
36 computations with whole arrays.

1 C.13.2.1 Unconditional array computations

2 At the simplest level, statements such as

3 $A = B + C$

4 or

5 $S = \text{SUM} (A)$

6 can take the place of entire DO loops. The loops were required to perform array addition or to sum all
7 the elements of an array.

8 Further examples of unconditional operations on arrays that are simple to write are:

matrix multiply	$P = \text{MATMUL} (Q, R)$
largest array element	$L = \text{MAXVAL} (P)$
factorial N	$F = \text{PRODUCT} ((/ (K, K = 2, N) /))$

9 The Fourier sum $F = \sum_{i=1}^N a_i \times \cos x_i$ may also be computed without writing a DO loop if one makes
10 use of the element-by-element definition of array expressions as described in Clause 7. Thus, we can
11 write

12 $F = \text{SUM} (A * \text{COS} (X))$

13 The successive stages of calculation of F would then involve the arrays:

A	$= (/ A (1), \dots, A (N) /)$
X	$= (/ X (1), \dots, X (N) /)$
$\text{COS} (X)$	$= (/ \text{COS} (X (1)), \dots, \text{COS} (X (N)) /)$
$A * \text{COS} (X)$	$= (/ A (1) * \text{COS} (X (1)), \dots, A (N) * \text{COS} (X (N)) /)$

14 The final scalar result is obtained simply by summing the elements of the last of these arrays. Thus, the
15 processor is dealing with arrays at every step of the calculation.

16 C.13.2.2 Conditional array computations

17 Suppose we wish to compute the Fourier sum in the above example, but to include only those terms
18 $a(i) \cos x(i)$ that satisfy the condition that the coefficient $a(i)$ is less than 0.01 in absolute value. More
19 precisely, we are now interested in evaluating the conditional Fourier sum

$$CF = \sum_{|a_i| < 0.01} a_i \times \cos x_i$$

20 where the index runs from 1 to N as before.

21 This can be done by using the MASK parameter of the SUM function, which restricts the summation
22 of the elements of the array $A * \text{COS} (X)$ to those elements that correspond to true elements of MASK.
23 Clearly, the mask required is the logical array expression $\text{ABS} (A) < 0.01$. Note that the stages of
24 evaluation of this expression are:

A	$= (/ A (1), \dots, A (N) /)$
$\text{ABS} (A)$	$= (/ \text{ABS} (A (1)), \dots, \text{ABS} (A (N)) /)$
$\text{ABS} (A) < 0.01$	$= (/ \text{ABS} (A (1)) < 0.01, \dots, \text{ABS} (A (N)) < 0.01 /)$

1 The conditional Fourier sum we arrive at is:

2 `CF = SUM (A * COS (X), MASK = ABS (A) < 0.01)`

3 If the mask is all false, the value of CF is zero.

4 The use of a mask to define a subset of an array is crucial to the action of the WHERE statement. Thus
5 for example, to zero an entire array, we may write simply `A = 0`; but to set only the negative elements
6 to zero, we need to write the conditional assignment

7 `WHERE (A .LT. 0) A = 0`

8 The WHERE statement complements ordinary array assignment by providing array assignment to any
9 subset of an array that can be restricted by a logical expression.

10 In the Ising model described below, the WHERE statement predominates in use over the ordinary array
11 assignment statement.

12 **C.13.2.3 A simple program: the Ising model**

13 The Ising model is a well-known Monte Carlo simulation in 3-dimensional Euclidean space which is
14 useful in certain physical studies. We will consider in some detail how this might be programmed. The
15 model may be described in terms of a logical array of shape N by N by N. Each gridpoint is a single
16 logical variable which is to be interpreted as either an up-spin (true) or a down-spin (false).

17 The Ising model operates by passing through many successive states. The transition to the next state is
18 governed by a local probabilistic process. At each transition, all gridpoints change state simultaneously.
19 Every spin either flips to its opposite state or not according to a rule that depends only on the states
20 of its 6 nearest neighbors in the surrounding grid. The neighbors of gridpoints on the boundary faces of
21 the model cube are defined by assuming cubic periodicity. In effect, this extends the grid periodically
22 by replicating it in all directions throughout space.

23 The rule states that a spin is flipped to its opposite parity for certain gridpoints where a mere 3 or
24 fewer of the 6 nearest neighbors have the same parity as it does. Also, the flip is executed only with
25 probability P (4), P (5), or P (6) if as many as 4, 5, or 6 of them have the same parity as it does. (The
26 rule seems to promote neighborhood alignments that may presumably lead to equilibrium in the long
27 run.)

28 **C.13.2.3.1 Problems to be solved**

29 Some of the programming problems that we will need to solve in order to translate the Ising model into
30 Fortran statements using entire arrays are

- 31 (1) counting nearest neighbors that have the same spin,
- 32 (2) providing an array function to return an array of random numbers, and
- 33 (3) determining which gridpoints are to be flipped.

34 **C.13.2.3.2 Solutions in Fortran**

35 The arrays needed are:

36 `LOGICAL ISING (N, N, N), FLIPS (N, N, N)`

37 `INTEGER ONES (N, N, N), COUNT (N, N, N)`

38 `REAL THRESHOLD (N, N, N)`

The array function needed is:

```

1
2 FUNCTION RAND (N)
3 REAL RAND (N, N, N)

```

4 The transition probabilities are specified in the array

```
5 REAL P (6)
```

6 The first task is to count the number of nearest neighbors of each gridpoint g that have the same spin
7 as g .

8 Assuming that ISING is given to us, the statements

```
9 ONES = 0
10 WHERE (ISING) ONES = 1

```

11 make the array ONES into an exact analog of ISING in which 1 stands for an up-spin and 0 for a
12 down-spin.

13 The next array we construct, COUNT, will record for every gridpoint of ISING the number of spins to
14 be found among the 6 nearest neighbors of that gridpoint. COUNT will be computed by adding together
15 6 arrays, one for each of the 6 relative positions in which a nearest neighbor is found. Each of the 6
16 arrays is obtained from the ONES array by shifting the ONES array one place circularly along one of
17 its dimensions. This use of circular shifting imparts the cubic periodicity.

```
18 COUNT = CSHIFT (ONES, SHIFT = -1, DIM = 1) &
19         + CSHIFT (ONES, SHIFT = 1, DIM = 1) &
20         + CSHIFT (ONES, SHIFT = -1, DIM = 2) &
21         + CSHIFT (ONES, SHIFT = 1, DIM = 2) &
22         + CSHIFT (ONES, SHIFT = -1, DIM = 3) &
23         + CSHIFT (ONES, SHIFT = 1, DIM = 3)

```

24 At this point, COUNT contains the count of nearest neighbor up-spins even at the gridpoints where
25 the Ising model has a down-spin. But we want a count of down-spins at those gridpoints, so we correct
26 COUNT at the down (false) points of ISING by writing:

```
27 WHERE (.NOT. ISING) COUNT = 6 - COUNT

```

28 Our object now is to use these counts of what may be called the “like-minded nearest neighbors” to
29 decide which gridpoints are to be flipped. This decision will be recorded as the true elements of an array
30 FLIPS. The decision to flip will be based on the use of uniformly distributed random numbers from the
31 interval $0 \leq p < 1$. These will be provided at each gridpoint by the array function RAND. The flip will
32 occur at a given point if and only if the random number at that point is less than a certain threshold
33 value. In particular, by making the threshold value equal to 1 at the points where there are 3 or fewer
34 like-minded nearest neighbors, we guarantee that a flip occurs at those points (because p is always less
35 than 1). Similarly, the threshold values corresponding to counts of 4, 5, and 6 are assigned P (4), P (5),
36 and P (6) in order to achieve the desired probabilities of a flip at those points (P (4), P (5), and P (6)
37 are input parameters in the range 0 to 1).

38 The thresholds are established by the statements:

```
39 THRESHOLD = 1.0
   WHERE (COUNT == 4) THRESHOLD = P (4)

```

```

1
2 WHERE (COUNT == 5) THRESHOLD = P (5)
3 WHERE (COUNT == 6) THRESHOLD = P (6)

```

4 and the spins that are to be flipped are located by the statement:

```
5 FLIPS = RAND (N) <= THRESHOLD
```

6 All that remains to complete one transition to the next state of the ISING model is to reverse the spins
7 in ISING wherever FLIPS is true:

```
8 WHERE (FLIPS) ISING = .NOT. ISING
```

9 **C.13.2.3.3 The complete Fortran subroutine**

10 The complete code, enclosed in a subroutine that performs a sequence of transitions, is as follows:

```

11 SUBROUTINE TRANSITION (N, ISING, ITERATIONS, P)
12
13     LOGICAL ISING (N, N, N), FLIPS (N, N, N)
14     INTEGER ONES (N, N, N), COUNT (N, N, N)
15     REAL THRESHOLD (N, N, N), P (6)
16
17     DO I = 1, ITERATIONS
18         ONES = 0
19         WHERE (ISING) ONES = 1
20         COUNT = CSHIFT (ONES, -1, 1) + CSHIFT (ONES, 1, 1) &
21             + CSHIFT (ONES, -1, 2) + CSHIFT (ONES, 1, 2) &
22             + CSHIFT (ONES, -1, 3) + CSHIFT (ONES, 1, 3)
23         WHERE (.NOT. ISING) COUNT = 6 - COUNT
24         THRESHOLD = 1.0
25         WHERE (COUNT == 4) THRESHOLD = P (4)
26         WHERE (COUNT == 5) THRESHOLD = P (5)
27         WHERE (COUNT == 6) THRESHOLD = P (6)
28         FLIPS = RAND (N) <= THRESHOLD
29         WHERE (FLIPS) ISING = .NOT. ISING
30     END DO
31
32 CONTAINS
33     FUNCTION RAND (N)
34         REAL RAND (N, N, N)
35         CALL RANDOM_NUMBER (HARVEST = RAND)
36         RETURN
37     END FUNCTION RAND
38 END

```

39 **C.13.2.3.4 Reduction of storage**

1 The array ISING could be removed (at some loss of clarity) by representing the model in ONES all the
2 time. The array FLIPS can be avoided by combining the two statements that use it as:

3 WHERE (RAND (N) <= THRESHOLD) ISING = .NOT. ISING

4 but an extra temporary array would probably be needed. Thus, the scope for saving storage while
5 performing whole array operations is limited. If N is small, this will not matter and the use of whole
6 array operations is likely to lead to good execution speed. If N is large, storage may be very important
7 and adequate efficiency will probably be available by performing the operations plane by plane. The
8 resulting code is not as elegant, but all the arrays except ISING will have size of order N^2 instead of N^3 .

9 **C.13.3 FORMula TRANslation and array processing**

10 Many mathematical formulas can be translated directly into Fortran by use of the array processing
11 features.

12 We assume the following array declarations:

13 REAL X (N), A (M, N)

14 Some examples of mathematical formulas and corresponding Fortran expressions follow.

15 **C.13.3.1 A sum of products**

The expression

$$\sum_{j=1}^N \prod_{i=1}^M a_{ij}$$

16 can be formed using the Fortran expression

17 SUM (PRODUCT (A, DIM=1))

18 The argument DIM=1 means that the product is to be computed down each column of A. If A had the
19 value $\begin{bmatrix} B & C & D \\ E & F & G \end{bmatrix}$ the result of this expression is BE + CF + DG.

20 **C.13.3.2 A product of sums**

The expression

$$\prod_{i=1}^M \sum_{j=1}^N a_{ij}$$

21 can be formed using the Fortran expression

22 PRODUCT (SUM (A, DIM = 2))

23 The argument DIM = 2 means that the sum is to be computed along each row of A. If A had the
24 value $\begin{bmatrix} B & C & D \\ E & F & G \end{bmatrix}$ the result of this expression is (B+C+D)(E+F+G).

25 **C.13.3.3 Addition of selected elements**

The expression

$$\sum_{x_i > 0.0} x_i$$

1 can be formed using the Fortran expression

2 `SUM (X, MASK = X > 0.0)`

3 The mask locates the positive elements of the array of rank one. If X has the vector value (0.0, -0.1,
4 0.2, 0.3, 0.2, -0.1, 0.0), the result of this expression is 0.7.

5 **C.13.4 Sum of squared residuals**

The expression

$$\sum_{i=1}^N (x_i - x_{\text{mean}})^2$$

6 can be formed using the Fortran statements

7 `XMEAN = SUM (X) / SIZE (X)`

8 `SS = SUM ((X - XMEAN) ** 2)`

9 Thus, SS is the sum of the squared residuals.

10 **C.13.5 Vector norms: infinity-norm and one-norm**

11 The infinity-norm of vector $X = (X(1), \dots, X(N))$ is defined as the largest of the numbers $ABS(X(1)),$
12 $\dots, ABS(X(N))$ and therefore has the value `MAXVAL (ABS(X))`.

13 The one-norm of vector X is defined as the *sum* of the numbers $ABS(X(1)), \dots, ABS(X(N))$ and
14 therefore has the value `SUM (ABS(X))`.

15 **C.13.6 Matrix norms: infinity-norm and one-norm**

16 The infinity-norm of the matrix $A = (A(I, J))$ is the largest row-sum of the matrix $ABS(A(I, J))$ and
17 therefore has the value `MAXVAL (SUM (ABS(A), DIM = 2))`.

18 The one-norm of the matrix $A = (A(I, J))$ is the largest column-sum of the matrix $ABS(A(I, J))$ and
19 therefore has the value `MAXVAL (SUM (ABS(A), DIM = 1))`.

20 **C.13.7 Logical queries**

21 The intrinsic functions allow quite complicated questions about tabular data to be answered without
22 use of loops or conditional constructs. Consider, for example, the questions asked below about a simple
23 tabulation of students' test scores.

24 Suppose the rectangular table T (M, N) contains the test scores of M students who have taken N different
25 tests. T is an integer matrix with entries in the range 0 to 100.

26 Example: The scores on 4 tests made by 3 students are held as the table

$$T = \begin{bmatrix} 85 & 76 & 90 & 60 \\ 71 & 45 & 50 & 80 \\ 66 & 45 & 21 & 55 \end{bmatrix}$$

27 Question: What is each student's top score?

28 Answer: `MAXVAL (T, DIM = 2)`; in the example: [90, 80, 66].

1 Question: What is the average of all the scores?

2 Answer: $\text{SUM}(T) / \text{SIZE}(T)$; in the example: 62.

3 Question: How many of the scores in the table are above average?

4 Answer: $\text{ABOVE} = T > \text{SUM}(T) / \text{SIZE}(T)$; $N = \text{COUNT}(\text{ABOVE})$; in the example: ABOVE is the
 5 logical array ($t = \text{true}$, $. = \text{false}$): $\begin{bmatrix} t & t & t & . \\ t & . & . & t \\ t & . & . & . \end{bmatrix}$ and $\text{COUNT}(\text{ABOVE})$ is 6.

6 Question: What was the lowest score in the above-average group of scores?

7 Answer: $\text{MINVAL}(T, \text{MASK} = \text{ABOVE})$, where ABOVE is as defined previously; in the example: 66.

8 Question: Was there a student whose scores were all above average?

9 Answer: With ABOVE as previously defined, the answer is yes or no according as the value of the
 10 expression $\text{ANY}(\text{ALL}(\text{ABOVE}, \text{DIM} = 2))$ is true or false; in the example, the answer is no.

11 C.13.8 Parallel computations

12 The most straightforward kind of parallel processing is to do the same thing at the same time to many
 13 operands. Matrix addition is a good example of this very simple form of parallel processing. Thus, the
 14 array assignment $A = B + C$ specifies that corresponding elements of the identically-shaped arrays B
 15 and C be added together in parallel and that the resulting sums be assigned in parallel to the array A.

16 The process being done in parallel in the example of matrix addition is of course the process of addi-
 17 tion; the array feature that implements matrix addition as a parallel process is the element-by-element
 18 evaluation of array expressions.

19 These observations lead us to look to element-by-element computation as a means of implementing other
 20 simple parallel processing algorithms.

21 C.13.9 Example of element-by-element computation

22 Several polynomials of the same degree may be evaluated at the same point by arranging their coefficients
 23 as the rows of a matrix and applying Horner's method for polynomial evaluation to the columns of the
 24 matrix so formed.

25 The procedure is illustrated by the code to evaluate the three cubic polynomials

$$P(t) = 1 + 2t - 3t^2 + 4t^3$$

$$Q(t) = 2 - 3t + 4t^2 - 5t^3$$

$$R(t) = 3 + 4t - 5t^2 + 6t^3$$

26 in parallel at the point $t = X$ and to place the resulting vector of numbers $[P(X), Q(X), R(X)]$ in the
 27 real array RESULT (3).

28 The code to compute RESULT is just the one statement

29 $\text{RESULT} = \text{M}(:, 1) + X * (\text{M}(:, 2) + X * (\text{M}(:, 3) + X * \text{M}(:, 4)))$

where M represents the matrix M (3, 4) with value $\begin{bmatrix} 1 & 2 & -3 & 4 \\ 2 & -3 & 4 & -5 \\ 3 & 4 & -5 & 6 \end{bmatrix}$.

1 C.13.10 Bit manipulation and inquiry procedures

2 The procedures IOR, IAND, NOT, IEXOR, ISHFT, ISHFTC, IBITS, MVBITS, BTEST, IBSET, and
3 IBCLR are defined by MIL-STD 1753 for scalar arguments and are extended in this standard to accept
4 array arguments and to return array results.

Annex D

(Informative)

Syntax rules

D.1 Extract of all syntax rules

Clause 1:

- R101 *xyz-list* is *xyz* [, *xyz*] ...
 R102 *xyz-name* is *name*
 R103 *scalar-xyz* is *xyz*
 C101 (R103) *scalar-xyz* shall be scalar.

Clause 2:

- R201 *program* is *program-unit*
 [*program-unit*] ...
 R202 *program-unit* is *main-program*
 or *external-subprogram*
 or *module*
 or *submodule*
 or *block-data*
 R203 *external-subprogram* is *function-subprogram*
 or *subroutine-subprogram*
 R204 *specification-part* is [*use-stmt*] ...
 [*import-stmt*] ...
 [*implicit-part*]
 [*declaration-construct*] ...
 R205 *implicit-part* is [*implicit-part-stmt*] ...
 implicit-stmt
 R206 *implicit-part-stmt* is *implicit-stmt*
 or *parameter-stmt*
 or *format-stmt*
 or *entry-stmt*
 R207 *declaration-construct* is *derived-type-def*
 or *entry-stmt*
 or *enum-def*
 or *format-stmt*
 or *interface-block*
 or *macro-definition*
 or *parameter-stmt*
 or *procedure-declaration-stmt*
 or *specification-stmt*
 or *type-declaration-stmt*
 or *stmt-function-stmt*

R208	<i>execution-part</i>	is <i>executable-construct</i> [<i>execution-part-construct</i>] ...
R209	<i>execution-part-construct</i>	is <i>executable-construct</i> or <i>format-stmt</i> or <i>entry-stmt</i> or <i>data-stmt</i>
R210	<i>internal-subprogram-part</i>	is <i>contains-stmt</i> [<i>internal-subprogram</i>] ...
R211	<i>internal-subprogram</i>	is <i>function-subprogram</i> or <i>subroutine-subprogram</i>
R212	<i>specification-stmt</i>	is <i>access-stmt</i> or <i>allocatable-stmt</i> or <i>asynchronous-stmt</i> or <i>bind-stmt</i> or <i>common-stmt</i> or <i>data-stmt</i> or <i>dimension-stmt</i> or <i>equivalence-stmt</i> or <i>external-stmt</i> or <i>intent-stmt</i> or <i>intrinsic-stmt</i> or <i>namelist-stmt</i> or <i>optional-stmt</i> or <i>pointer-stmt</i> or <i>protected-stmt</i> or <i>save-stmt</i> or <i>target-stmt</i> or <i>volatile-stmt</i> or <i>value-stmt</i>
R213	<i>executable-construct</i>	is <i>action-stmt</i> or <i>associate-construct</i> or <i>block-construct</i> or <i>case-construct</i> or <i>critical-construct</i> or <i>do-construct</i> or <i>forall-construct</i> or <i>if-construct</i> or <i>select-type-construct</i> or <i>where-construct</i>
R214	<i>action-stmt</i>	is <i>allocate-stmt</i> or <i>assignment-stmt</i> or <i>backspace-stmt</i> or <i>call-stmt</i> or <i>close-stmt</i> or <i>continue-stmt</i> or <i>cycle-stmt</i>

or *deallocate-stmt*
 or *endfile-stmt*
 or *end-function-stmt*
 or *end-program-stmt*
 or *end-subroutine-stmt*
 or *exit-stmt*
 or *flush-stmt*
 or *forall-stmt*
 or *goto-stmt*
 or *if-stmt*
 or *inquire-stmt*
 or *notify-stmt*
 or *nullify-stmt*
 or *open-stmt*
 or *pointer-assignment-stmt*
 or *print-stmt*
 or *query-stmt*
 or *read-stmt*
 or *return-stmt*
 or *rewind-stmt*
 or *stop-stmt*
 or *sync-all-stmt*
 or *sync-images-stmt*
 or *sync-memory-stmt*
 or *sync-team-stmt*
 or *wait-stmt*
 or *where-stmt*
 or *write-stmt*
 or *arithmetic-if-stmt*
 or *computed-goto-stmt*

C201 (R208) An *execution-part* shall not contain an *end-function-stmt*, *end-program-stmt*, or *end-subroutine-stmt*.

R215 *keyword* is *name*

Clause 3:

R301 *character* is *alphanumeric-character*
 or *special-character*

R302 *alphanumeric-character* is *letter*
 or *digit*
 or *underscore*

R303 *underscore* is *-*

R304 *name* is *letter* [*alphanumeric-character*] ...

C301 (R304) The maximum length of a *name* is 63 characters.

R305 *constant* is *literal-constant*
 or *named-constant*

R306 *literal-constant* is *int-literal-constant*
 or *real-literal-constant*

		or <i>complex-literal-constant</i>
		or <i>logical-literal-constant</i>
		or <i>char-literal-constant</i>
		or <i>boz-literal-constant</i>
R307	<i>named-constant</i>	is <i>name</i>
R308	<i>int-constant</i>	is <i>constant</i>
C302	(R308) <i>int-constant</i> shall be of type integer.	
R309	<i>char-constant</i>	is <i>constant</i>
C303	(R309) <i>char-constant</i> shall be of type character.	
R310	<i>intrinsic-operator</i>	is <i>power-op</i>
		or <i>mult-op</i>
		or <i>add-op</i>
		or <i>concat-op</i>
		or <i>rel-op</i>
		or <i>not-op</i>
		or <i>and-op</i>
		or <i>or-op</i>
		or <i>equiv-op</i>
R311	<i>defined-operator</i>	is <i>defined-unary-op</i>
		or <i>defined-binary-op</i>
		or <i>extended-intrinsic-op</i>
R312	<i>extended-intrinsic-op</i>	is <i>intrinsic-operator</i>
R313	<i>label</i>	is <i>digit</i> [<i>digit</i> [<i>digit</i> [<i>digit</i> [<i>digit</i>]]]]
C304	(R313) At least one digit in a <i>label</i> shall be nonzero.	
R314	<i>macro-definition</i>	is <i>define-macro-stmt</i>
		[<i>macro-declaration-stmt</i>] ...
		<i>macro-body-block</i>
		<i>end-macro-stmt</i>
R315	<i>define-macro-stmt</i>	is DEFINE MACRO [, <i>macro-attribute-list</i>] :: <i>macro-name</i> ■ ■ [([<i>macro-dummy-arg-name-list</i>])]
C305	(R315) A <i>macro-dummy-arg-name</i> shall not appear more than once in a <i>macro-dummy-arg-name-list</i> .	
R316	<i>macro-attribute</i>	is <i>access-spec</i>
R317	<i>macro-declaration-stmt</i>	is <i>macro-type-declaration-stmt</i>
		or <i>macro-optional-decl-stmt</i>
R318	<i>macro-type-declaration-stmt</i>	is MACRO <i>macro-type-spec</i> :: <i>macro-local-variable-name-list</i>
R319	<i>macro-optional-decl-stmt</i>	is MACRO OPTIONAL :: <i>macro-dummy-arg-name-list</i>
R320	<i>macro-type-spec</i>	is INTEGER [([KIND=] <i>macro-expr</i>)]
C306	(R318) A <i>macro-local-variable-name</i> shall not be the same as the name of a dummy argument of the macro being defined.	
C307	(R319) A <i>macro-dummy-arg-name</i> shall be the name of a dummy argument of the macro being defined.	
C308	(R320) If <i>macro-expr</i> appears, when the macro is expanded <i>macro-expr</i> shall be of type integer, and have a non-negative value that specifies a representation method that exists on the processor.	
R321	<i>macro-body-block</i>	is [<i>macro-body-construct</i>] ...
R322	<i>macro-body-construct</i>	is <i>macro-definition</i>
		or <i>expand-stmt</i>

- or** *macro-body-stmt*
or *macro-do-construct*
or *macro-if-construct*
- C309 A statement in a macro definition that is not a *macro-body-construct* or *macro-definition* shall not appear on a line with any other statement.
- R323 *macro-do-construct* **is** *macro-do-stmt*
macro-body-block
macro-end-do-stmt
- R324 *macro-do-stmt* **is** MACRO DO *macro-do-variable-name* = *macro-do-limit* , ■
■ *macro-do-limit* [, *macro-do-limit*]
- C310 (R324) A *macro-do-variable-name* shall be a local variable of the macro being defined, and shall not be a macro dummy argument.
- R325 *macro-do-limit* **is** *macro-expr*
- C311 (R325) A *macro-do-limit* shall expand to an expression of type integer.
- R326 *macro-end-do-stmt* **is** MACRO END DO
- R327 *macro-if-construct* **is** *macro-if-then-stmt*
macro-body-block
[*macro-else-if-stmt*
macro-body-block] ...
[*macro-else-stmt*
macro-body-block]
macro-end-if-stmt
- R328 *macro-if-then-stmt* **is** MACRO IF (*macro-condition*) THEN
- R329 *macro-else-if-stmt* **is** MACRO ELSE IF (*macro-condition*) THEN
- R330 *macro-else-stmt* **is** MACRO ELSE
- R331 *macro-end-if-stmt* **is** MACRO END IF
- R332 *macro-condition* **is** *macro-expr*
- C312 (R332) A macro condition shall expand to an expression of type logical.
- R333 *macro-body-stmt* **is** *result-token* [*result-token*] ... [&&]
- C313 (R333) The first *result-token* shall not be MACRO unless the second *result-token* is not a keyword or name.
- R334 *result-token* **is** *token* [%% *token*] ...
- C314 (R334) The concatenated textual *tokens* in a *result-token* shall have the form of a lexical token.
- R335 *token* **is** any lexical token including labels, keywords, and semi-colon.
- C315 && shall not appear in the last *macro-body-stmt* of a macro definition.
- C316 When a macro is expanded, the last *macro-body-stmt* processed shall not end with &&.
- R336 *end-macro-stmt* **is** END MACRO [*macro-name*]
- C317 (R314) The *macro-name* in the END MACRO statement shall be the same as the *macro-name* in the DEFINE MACRO statement.
- R337 *macro-expr* **is** *basic-token-sequence*
- C318 (R337) A *macro-expr* shall expand to a scalar initialization expression.
- R338 *expand-stmt* **is** EXPAND *macro-name* [(*macro-actual-arg-list*)]
- C319 (R338) *macro-name* shall be the name of a macro that was previously defined or accessed via use or host association.
- C320 (R338) The macro shall expand to a sequence or zero or more complete Fortran statements.
- C321 (R338) The statements produced by a macro expansion shall conform to the syntax rules and

constraints as if they replaced the EXPAND statement prior to program processing.

- C322 (R338) The statements produced by a macro expansion shall not include a statement which ends the scoping unit containing the EXPAND statement.
- C323 (R338) If a macro expansion produces a statement which begins a new scoping unit, it shall also produce a statement which ends that scoping unit.
- C324 (R338) If the EXPAND statement appears as the *action-stmt* of an *if-stmt*, it shall expand to exactly one *action-stmt* that is not an *if-stmt*, *end-program-stmt*, *end-function-stmt*, or *end-subroutine-stmt*.
- C325 (R338) If the EXPAND statement appears as a *do-term-action-stmt*, it shall expand to exactly one *action-stmt* that is not a *continue-stmt*, a *goto-stmt*, a *return-stmt*, a *stop-stmt*, an *exit-stmt*, a *cycle-stmt*, an *end-function-stmt*, an *end-subroutine-stmt*, an *end-program-stmt*, or an *arithmetic-if-stmt*.
- C326 (R338) If the EXPAND statement has a label, the expansion of the macro shall produce at least one statement, and the first statement produced shall not have a label.
- C327 (R338) A *macro-actual-arg* shall appear corresponding to each nonoptional macro dummy argument.
- C328 (R338) At most one *macro-actual-arg* shall appear corresponding to each optional macro dummy argument.
- R339 *macro-actual-arg* **is** [*macro-dummy-name* =] *macro-actual-arg-value*
- C329 (R339) *macro-dummy-name* shall be the name of a macro dummy argument of the macro being expanded.
- C330 (R338) The *macro-dummy-name*= shall not be omitted unless it has been omitted from each preceding *macro-actual-arg* in the *expand-stmt*.
- R340 *macro-actual-arg-value* **is** *basic-token-sequence*
- R341 *basic-token-sequence* **is** *basic-token*
 or [*basic-token-sequence*] *nested-token-sequence* ■
 ■ [*basic-token-sequence*]
 or *basic-token basic-token-sequence*
- R342 *basic-token* **is** any lexical token except comma, parentheses, array ■
 ■ constructor delimiters, and semi-colon.
- R343 *nested-token-sequence* **is** ([*arg-token*] ...)
 or (/ [*arg-token*] ... /)
 or *lbracket* [*arg-token*] ... *rbracket*
- R344 *arg-token* **is** *basic-token*
 or ,

Clause 4:

- R401 *type-param-value* **is** *scalar-int-expr*
 or *
 or :
- C401 (R401) The *type-param-value* for a kind type parameter shall be an initialization expression.
- C402 (R401) A colon may be used as a *type-param-value* only in the declaration of an entity or component that has the POINTER or ALLOCATABLE attribute.
- R402 *type-spec* **is** *intrinsic-type-spec*
 or *derived-type-spec*
- C403 (R402) The *derived-type-spec* shall not specify an abstract type (4.5.7).
- R403 *declaration-type-spec* **is** *intrinsic-type-spec*
 or TYPE (*intrinsic-type-spec*)
 or TYPE (*derived-type-spec*)
 or CLASS (*derived-type-spec*)

- or** CLASS (*)
- C404 (R403) In a *declaration-type-spec*, every *type-param-value* that is not a colon or an asterisk shall be a *specification-expr*.
- C405 (R403) In a *declaration-type-spec* that uses the CLASS keyword, *derived-type-spec* shall specify an extensible type (4.5.7).
- C406 (R403) The TYPE(*derived-type-spec*) shall not specify an abstract type (4.5.7).
- C407 An entity declared with the CLASS keyword shall be a dummy argument or have the ALLOCATABLE or POINTER attribute. It shall not have the VALUE attribute.
- R404 *intrinsic-type-spec* **is** INTEGER [*kind-selector*]
 or REAL [*kind-selector*]
 or DOUBLE PRECISION
 or COMPLEX [*kind-selector*]
 or CHARACTER [*char-selector*]
 or LOGICAL [*kind-selector*]
 or BITS [*kind-selector*]
- R405 *kind-selector* **is** ([KIND =] *scalar-int-initialization-expr*)
- C408 (R405) The value of *scalar-int-initialization-expr* shall be nonnegative and shall specify a representation method that exists on the processor.
- R406 *signed-int-literal-constant* **is** [*sign*] *int-literal-constant*
- R407 *int-literal-constant* **is** *digit-string* [_ *kind-param*]
- R408 *kind-param* **is** *digit-string*
 or *scalar-int-constant-name*
- R409 *signed-digit-string* **is** [*sign*] *digit-string*
- R410 *digit-string* **is** *digit* [*digit*] ...
- R411 *sign* **is** +
 or -
- C409 (R408) A *scalar-int-constant-name* shall be a named constant of type integer.
- C410 (R408) The value of *kind-param* shall be nonnegative.
- C411 (R407) The value of *kind-param* shall specify a representation method that exists on the processor.
- R412 *signed-real-literal-constant* **is** [*sign*] *real-literal-constant*
- R413 *real-literal-constant* **is** *significand* [*exponent-letter* *exponent*] [_ *kind-param*]
 or *digit-string* *exponent-letter* *exponent* [_ *kind-param*]
- R414 *significand* **is** *digit-string* . [*digit-string*]
 or . *digit-string*
- R415 *exponent-letter* **is** E
 or D
- R416 *exponent* **is** *signed-digit-string*
- C412 (R413) If both *kind-param* and *exponent-letter* are present, *exponent-letter* shall be E.
- C413 (R413) The value of *kind-param* shall specify an approximation method that exists on the processor.
- R417 *complex-literal-constant* **is** (*real-part* , *imag-part*)
- R418 *real-part* **is** *signed-int-literal-constant*
 or *signed-real-literal-constant*
 or *named-constant*
- R419 *imag-part* **is** *signed-int-literal-constant*
 or *signed-real-literal-constant*

- or** *named-constant*
- C414 (R417) Each named constant in a complex literal constant shall be of type integer or real.
- R420 *char-selector* **is** *length-selector*
or (*LEN* = *type-param-value* , ■
 ■ *KIND* = *scalar-int-initialization-expr*)
or (*type-param-value* , ■
 ■ [*KIND* =] *scalar-int-initialization-expr*)
or (*KIND* = *scalar-int-initialization-expr* ■
 ■ [, *LEN* = *type-param-value*])
- R421 *length-selector* **is** ([*LEN* =] *type-param-value*)
or * *char-length* [,]
- R422 *char-length* **is** (*type-param-value*)
or *scalar-int-literal-constant*
- C415 (R420) The value of *scalar-int-initialization-expr* shall be nonnegative and shall specify a representation method that exists on the processor.
- C416 (R422) The *scalar-int-literal-constant* shall not include a *kind-param*.
- C417 (R422) A *type-param-value* in a *char-length* shall be a colon, asterisk, or *specification-expr*.
- C418 (R420 R421 R422) A *type-param-value* of * shall be used only
- C419 A function name shall not be declared with an asterisk *type-param-value* unless it is of type CHARACTER and is the name of the result of an external function or the name of a dummy function.
- C420 A function name declared with an asterisk *type-param-value* shall not be an array, a pointer, recursive, or pure.
- C421 (R421) The optional comma in a *length-selector* is permitted only in a *declaration-type-spec* in a *type-declaration-stmt*.
- C422 (R421) The optional comma in a *length-selector* is permitted only if no double-colon separator appears in the *type-declaration-stmt*.
- C423 (R420) The length specified for a character statement function or for a statement function dummy argument of type character shall be an initialization expression.
- R423 *char-literal-constant* **is** [*kind-param* _] ' [*rep-char*] ... '
or [*kind-param* _] " [*rep-char*] ... "
- C424 (R423) The value of *kind-param* shall specify a representation method that exists on the processor.
- R424 *logical-literal-constant* **is** .TRUE. [_ *kind-param*]
or .FALSE. [_ *kind-param*]
- C425 (R424) The value of *kind-param* shall specify a representation method that exists on the processor.
- R425 *boz-literal-constant* **is** *binary-constant* [_ *kind-param*]
or *octal-constant* [_ *kind-param*]
or *hex-constant* [_ *kind-param*]
- R426 *binary-constant* **is** B ' *digit* [*digit*] ... '
or B " *digit* [*digit*] ... "
- C426 (R426) *digit* shall have one of the values 0 or 1.
- R427 *octal-constant* **is** O ' *digit* [*digit*] ... '
or O " *digit* [*digit*] ... "
- C427 (R427) *digit* shall have one of the values 0 through 7.
- R428 *hex-constant* **is** Z ' *hex-digit* [*hex-digit*] ... '
or Z " *hex-digit* [*hex-digit*] ... "
- R429 *hex-digit* **is** *digit*

- or** A
or B
or C
or D
or E
or F
- R430 *derived-type-def* **is** *derived-type-stmt*
 [*type-param-def-stmt*] ...
 [*private-or-sequence*] ...
 [*component-part*]
 [*type-bound-procedure-part*]
 end-type-stmt
- R431 *derived-type-stmt* **is** TYPE [[, *type-attr-spec-list*] ::] *type-name* ■
 ■ [(*type-param-name-list*)]
- R432 *type-attr-spec* **is** ABSTRACT
or *access-spec*
or BIND (C)
or EXTENDS (*parent-type-name*)
- C428 (R431) A derived type *type-name* shall not be DOUBLEPRECISION or the same as the name of any intrinsic type defined in this part of ISO/IEC 1539.
- C429 (R431) The same *type-attr-spec* shall not appear more than once in a given *derived-type-stmt*.
- C430 (R432) A *parent-type-name* shall be the name of a previously defined extensible type (4.5.7).
- C431 (R430) If the type definition contains or inherits (4.5.7.2) a deferred binding (4.5.5), ABSTRACT shall appear.
- C432 (R430) If ABSTRACT appears, the type shall be extensible.
- C433 (R430) If EXTENDS appears, SEQUENCE shall not appear.
- C434 (R430) If EXTENDS appears and the type being defined has a co-array ultimate component, its parent type shall have a co-array ultimate component.
- R433 *private-or-sequence* **is** *private-components-stmt*
or *sequence-stmt*
- C435 (R430) The same *private-or-sequence* shall not appear more than once in a given *derived-type-def*.
- R434 *end-type-stmt* **is** END TYPE [*type-name*]
- C436 (R434) If END TYPE is followed by a *type-name*, the *type-name* shall be the same as that in the corresponding *derived-type-stmt*.
- R435 *sequence-stmt* **is** SEQUENCE
- C437 (R439) If SEQUENCE appears, each data component shall be declared to be of an intrinsic type or of a sequence derived type.
- C438 (R430) If SEQUENCE appears, a *type-bound-procedure-part* shall not appear.
- R436 *type-param-def-stmt* **is** INTEGER [*kind-selector*] , *type-param-attr-spec* :: ■
 ■ *type-param-decl-list*
- R437 *type-param-decl* **is** *type-param-name* [= *scalar-int-initialization-expr*]
- C439 (R436) A *type-param-name* in a *type-param-def-stmt* in a *derived-type-def* shall be one of the *type-param-names* in the *derived-type-stmt* of that *derived-type-def*.
- C440 (R436) Each *type-param-name* in the *derived-type-stmt* in a *derived-type-def* shall appear as a *type-param-name* in a *type-param-def-stmt* in that *derived-type-def*.
- R438 *type-param-attr-spec* **is** KIND

- or LEN
- R439 *component-part* is [*component-def-stmt*] ...
- R440 *component-def-stmt* is *data-component-def-stmt*
or *proc-component-def-stmt*
- R441 *data-component-def-stmt* is *declaration-type-spec* [[, *component-attr-spec-list*] ::] ■
■ *component-decl-list*
- R442 *component-attr-spec* is *access-spec*
or ALLOCATABLE
or DIMENSION (*component-array-spec*)
or DIMENSION [(*deferred-shape-spec-list*)] ■
■ *lbracket co-array-spec rbracket*
or CONTIGUOUS
or POINTER
- R443 *component-decl* is *component-name* [(*component-array-spec*)] ■
■ [*lbracket co-array-spec rbracket*] ■
■ [* *char-length*] [*component-initialization*]
- R444 *component-array-spec* is *explicit-shape-spec-list*
or *deferred-shape-spec-list*
- C441 (R441) No *component-attr-spec* shall appear more than once in a given *component-def-stmt*.
- C442 (R441) If neither the POINTER nor ALLOCATABLE attribute is specified, the *declaration-type-spec* in the *component-def-stmt* shall specify an intrinsic type or a previously defined derived type.
- C443 (R441) If the POINTER or ALLOCATABLE attribute is specified, the *declaration-type-spec* in the *component-def-stmt* shall be CLASS(*) or shall specify an intrinsic type or any accessible derived type including the type being defined.
- C444 (R441) If the POINTER or ALLOCATABLE attribute is specified, each *component-array-spec* shall be a *deferred-shape-spec-list*.
- C445 (R441) If a *co-array-spec* appears, it shall be a *deferred-co-shape-spec-list* and the component shall have the ALLOCATABLE attribute.
- C446 (R441) If a *co-array-spec* appears, the component shall not be of type IMAGE.TEAM (13.8.3.7), C_PTR, or C_FUNPTR (15.3.3).
- C447 A data component whose type has a co-array ultimate component shall be a nonpointer nonallocatable scalar and shall not be a co-array.
- C448 (R441) If neither the POINTER attribute nor the ALLOCATABLE attribute is specified, each *component-array-spec* shall be an *explicit-shape-spec-list*.
- C449 (R444) Each bound in the *explicit-shape-spec* shall either be an initialization expression or be a specification expression that does not contain references to specification functions or any object designators other than named constants or subobjects thereof.
- C450 (R441) A component shall not have both the ALLOCATABLE and the POINTER attribute.
- C451 (R441) If the CONTIGUOUS attribute is specified, the component shall be an array with the POINTER attribute.
- C452 (R443) The * *char-length* option is permitted only if the component is of type character.
- C453 (R440) Each *type-param-value* within a *component-def-stmt* shall either be a colon, be an initialization expression, or be a specification expression that contains neither references to specification functions nor any object designators other than named constants or subobjects thereof.
- R445 *proc-component-def-stmt* is PROCEDURE ([*proc-interface*]) , ■
■ *proc-component-attr-spec-list* :: *proc-decl-list*
- R446 *proc-component-attr-spec* is POINTER

- or** PASS [(*arg-name*)]
or NOPASS
or *access-spec*
- C454 (R445) The same *proc-component-attr-spec* shall not appear more than once in a given *proc-component-def-stmt*.
- C455 (R445) POINTER shall appear in each *proc-component-attr-spec-list*.
- C456 (R445) If the procedure pointer component has an implicit interface or has no arguments, NOPASS shall be specified.
- C457 (R445) If PASS (*arg-name*) appears, the interface shall have a dummy argument named *arg-name*.
- C458 (R445) PASS and NOPASS shall not both appear in the same *proc-component-attr-spec-list*.
- C459 The passed-object dummy argument shall be a scalar, nonpointer, nonallocatable dummy data object with the same declared type as the type being defined; all of its length type parameters shall be assumed; it shall be polymorphic (4.3.1.3) if and only if the type being defined is extensible (4.5.7). It shall not have the VALUE attribute.
- R447 *component-initialization* **is** = *initialization-expr*
 or => *null-init*
 or => *initial-data-target*
- R448 *initial-data-target* **is** *designator*
- C460 (R441) If *component-initialization* appears, a double-colon separator shall appear before the *component-decl-list*.
- C461 (R441) If *component-initialization* appears, every type parameter and array bound of the component shall be a colon or initialization expression.
- C462 (R441) If => appears in *component-initialization*, POINTER shall appear in the *component-attr-spec-list*. If = appears in *component-initialization*, neither POINTER nor ALLOCATABLE shall appear in the *component-attr-spec-list*.
- C463 (R447) If *initial-data-target* appears, *component-name* shall be data-pointer-initialization compatible with it.
- C464 (R448) The *designator* shall designate a variable that is an initialization target. Every subscript, section subscript, substring starting point, and substring ending point in *designator* shall be an initialization expression.
- R449 *private-components-stmt* **is** PRIVATE
- C465 (R449) A *private-components-stmt* is permitted only if the type definition is within the specification part of a module.
- R450 *type-bound-procedure-part* **is** *contains-stmt*
 [*binding-private-stmt*]
 [*proc-binding-stmt*] ...
- R451 *binding-private-stmt* **is** PRIVATE
- C466 (R450) A *binding-private-stmt* is permitted only if the type definition is within the specification part of a module.
- R452 *proc-binding-stmt* **is** *specific-binding*
 or *generic-binding*
 or *final-binding*
- R453 *specific-binding* **is** PROCEDURE [(*interface-name*)] ■
 ■ [[, *binding-attr-list*] ::] ■

- *binding-name* [=> *procedure-name*]
- C467 (R453) If => *procedure-name* appears, the double-colon separator shall appear.
- C468 (R453) If => *procedure-name* appears, *interface-name* shall not appear.
- C469 (R453) The *procedure-name* shall be the name of an accessible module procedure or an external procedure that has an explicit interface.
- R454 *generic-binding* **is** **GENERIC** ■
 ■ [, *access-spec*] :: *generic-spec* => *binding-name-list*
- C470 (R454) Within the *specification-part* of a module, each *generic-binding* shall specify, either implicitly or explicitly, the same accessibility as every other *generic-binding* with that *generic-spec* in the same derived type.
- C471 (R454) Each *binding-name* in *binding-name-list* shall be the name of a specific binding of the type.
- C472 (R454) If *generic-spec* is not *generic-name*, each of its specific bindings shall have a passed-object dummy argument (4.5.4.4).
- C473 (R454) If *generic-spec* is OPERATOR (*defined-operator*), the interface of each binding shall be as specified in 12.4.3.3.1.
- C474 (R454) If *generic-spec* is ASSIGNMENT (=), the interface of each binding shall be as specified in 12.4.3.3.2.
- C475 (R454) If *generic-spec* is *dtio-generic-spec*, the interface of each binding shall be as specified in 9.5.4.7. The type of the *dtv* argument shall be *type-name*.
- R455 *binding-attr* **is** **PASS** [(*arg-name*)]
 or **NOPASS**
 or **NON_OVERRIDABLE**
 or **DEFERRED**
 or *access-spec*
- C476 (R455) The same *binding-attr* shall not appear more than once in a given *binding-attr-list*.
- C477 (R453) If the interface of the binding has no dummy argument of the type being defined, NOPASS shall appear.
- C478 (R453) If PASS (*arg-name*) appears, the interface of the binding shall have a dummy argument named *arg-name*.
- C479 (R455) PASS and NOPASS shall not both appear in the same *binding-attr-list*.
- C480 (R455) NON_OVERRIDABLE and DEFERRED shall not both appear in the same *binding-attr-list*.
- C481 (R455) DEFERRED shall appear if and only if *interface-name* appears.
- C482 (R453) An overriding binding (4.5.7.3) shall have the DEFERRED attribute only if the binding it overrides is deferred.
- C483 (R453) A binding shall not override an inherited binding (4.5.7.2) that has the NON_OVERRIDABLE attribute.
- R456 *final-binding* **is** **FINAL** [::] *final-subroutine-name-list*
- C484 (R456) A *final-subroutine-name* shall be the name of a module procedure with exactly one dummy argument. That argument shall be nonoptional and shall be a nonpointer, nonallocatable, nonpolymorphic variable of the derived type being defined. All length type parameters of the dummy argument shall be assumed. The dummy argument shall not be INTENT(OUT).
- C485 (R456) A *final-subroutine-name* shall not be one previously specified as a final subroutine for that type.
- C486 (R456) A final subroutine shall not have a dummy argument with the same kind type parameters and rank as the dummy argument of another final subroutine of that type.
- R457 *derived-type-spec* **is** *type-name* [(*type-param-spec-list*)]

- R458 *type-param-spec* **is** [*keyword* =] *type-param-value*
- C487 (R457) *type-name* shall be the name of an accessible derived type.
- C488 (R457) *type-param-spec-list* shall appear only if the type is parameterized.
- C489 (R457) There shall be at most one *type-param-spec* corresponding to each parameter of the type. If a type parameter does not have a default value, there shall be a *type-param-spec* corresponding to that type parameter.
- C490 (R458) The *keyword*= may be omitted from a *type-param-spec* only if the *keyword*= has been omitted from each preceding *type-param-spec* in the *type-param-spec-list*.
- C491 (R458) Each *keyword* shall be the name of a parameter of the type.
- C492 (R458) An asterisk may be used as a *type-param-value* in a *type-param-spec* only in the declaration of a dummy argument or associate name or in the allocation of a dummy argument.
- R459 *structure-constructor* **is** *derived-type-spec* ([*component-spec-list*])
- R460 *component-spec* **is** [*keyword* =] *component-data-source*
- R461 *component-data-source* **is** *expr*
or *data-target*
or *proc-target*
- C493 (R459) The *derived-type-spec* shall not specify an abstract type (4.5.7).
- C494 (R459) At most one *component-spec* shall be provided for a component.
- C495 (R459) If a *component-spec* is provided for an ancestor component, a *component-spec* shall not be provided for any component that is inheritance associated with a subcomponent of that ancestor component.
- C496 (R459) A *component-spec* shall be provided for a nonallocatable component unless it has default initialization or is inheritance associated with a subcomponent of another component for which a *component-spec* is provided.
- C497 (R460) The *keyword*= may be omitted from a *component-spec* only if the *keyword*= has been omitted from each preceding *component-spec* in the constructor.
- C498 (R460) Each *keyword* shall be the name of a component of the type.
- C499 (R459) The type name and all components of the type for which a *component-spec* appears shall be accessible in the scoping unit containing the structure constructor.
- C4100 (R459) If *derived-type-spec* is a type name that is the same as a generic name, the *component-spec-list* shall not be a valid *actual-arg-spec-list* for a function reference that is resolvable as a generic reference to that name (12.5.5.2).
- C4101 (R461) A *data-target* shall correspond to a nonprocedure pointer component; a *proc-target* shall correspond to a procedure pointer component.
- C4102 (R461) A *data-target* shall have the same rank as its corresponding component.
- R462 *enum-def* **is** *enum-def-stmt*
enumerator-def-stmt
[*enumerator-def-stmt*] ...
end-enum-stmt
- R463 *enum-def-stmt* **is** ENUM, BIND(C)
- R464 *enumerator-def-stmt* **is** ENUMERATOR [::] *enumerator-list*
- R465 *enumerator* **is** *named-constant* [= *scalar-int-initialization-expr*]
- R466 *end-enum-stmt* **is** END ENUM
- C4103 (R464) If = appears in an *enumerator*, a double-colon separator shall appear before the *enumerator-list*.
- R467 *array-constructor* **is** (/ *ac-spec* /)
or *lbracket ac-spec rbracket*
- R468 *ac-spec* **is** *type-spec* ::

- R469 *lbracket* **or** [*type-spec* ::] *ac-value-list*
 R470 *rbracket* **is** [
 R471 *ac-value* **is** *expr*
or *ac-implied-do*
 R472 *ac-implied-do* **is** (*ac-value-list* , *ac-implied-do-control*)
 R473 *ac-implied-do-control* **is** *ac-do-variable* = *scalar-int-expr* , *scalar-int-expr* ■
 ■ [, *scalar-int-expr*]
 R474 *ac-do-variable* **is** *do-variable*
 C4104 (R468) If *type-spec* is omitted, each *ac-value* expression in the *array-constructor* shall have the same type and kind type parameters.
 C4105 (R468) If *type-spec* specifies an intrinsic type, each *ac-value* expression in the *array-constructor* shall be of an intrinsic type that is in type conformance with a variable of type *type-spec* as specified in Table 7.11.
 C4106 (R468) If *type-spec* specifies a derived type, all *ac-value* expressions in the *array-constructor* shall be of that derived type and shall have the same kind type parameter values as specified by *type-spec*.
 C4107 (R472) The *ac-do-variable* of an *ac-implied-do* that is in another *ac-implied-do* shall not appear as the *ac-do-variable* of the containing *ac-implied-do*.

Clause 5:

- R501 *type-declaration-stmt* **is** *declaration-type-spec* [[, *attr-spec*] ... ::] *entity-decl-list*
 R502 *attr-spec* **is** *access-spec*
or ALLOCATABLE
or ASYNCHRONOUS
or CONTIGUOUS
or dimension-spec
or EXTERNAL
or INTENT (*intent-spec*)
or INTRINSIC
or *language-binding-spec*
or OPTIONAL
or PARAMETER
or POINTER
or PROTECTED
or SAVE
or TARGET
or VALUE
or VOLATILE
 C501 (R501) The same *attr-spec* shall not appear more than once in a given *type-declaration-stmt*.
 C502 (R501) If a *language-binding-spec* with a NAME= specifier appears, the *entity-decl-list* shall consist of a single *entity-decl*.
 C503 (R501) If a *language-binding-spec* is specified, the *entity-decl-list* shall not contain any procedure names.
 R503 *entity-decl* **is** *object-name* [(*array-spec*)] ■
 ■ [*lbracket* *co-array-spec* *rbracket*] ■
 ■ [* *char-length*] [*initialization*]

- or** *function-name* [* *char-length*]
- C504 (R503) If the entity is not of type character, * *char-length* shall not appear.
- C505 (R501) If *initialization* appears, a double-colon separator shall appear before the *entity-decl-list*.
- C506 (R503) An *initialization* shall not appear if *object-name* is a dummy argument, a function result, an object in a named common block unless the type declaration is in a block data program unit, an object in blank common, an allocatable variable, an external function, an intrinsic function, or an automatic object.
- C507 (R503) An *initialization* shall appear if the entity is a named constant (5.3.12).
- C508 (R503) The *function-name* shall be the name of an external function, an intrinsic function, a function dummy procedure, a procedure pointer, or a statement function.
- R504 *object-name* **is** *name*
- C509 (R504) The *object-name* shall be the name of a data object.
- R505 *initialization* **is** = *initialization-expr*
 or => *null-init*
 or => *initial-data-target*
- R506 *null-init* **is** *function-reference*
- C510 (R503) If => appears in *initialization*, the entity shall have the POINTER attribute. If = appears in *initialization*, the entity shall not have the POINTER attribute.
- C511 (R503) If *initial-data-target* appears, *object-name* shall be data-pointer-initialization compatible with it (4.5.4.5).
- C512 (R506) The *function-reference* shall be a reference to the NULL intrinsic function with no arguments.
- C513 An automatic object shall not be a local variable of a main program, module, or submodule.
- C514 An entity shall not be explicitly given any attribute more than once in a scoping unit.
- C515 An *array-spec* for a function result that does not have the ALLOCATABLE or POINTER attribute shall be an *explicit-shape-spec-list*.
- C516 The ALLOCATABLE, POINTER, or OPTIONAL attribute shall not be specified for a dummy argument of a procedure that has a *proc-language-binding-spec*.
- R507 *access-spec* **is** PUBLIC
 or PRIVATE
- C517 (R507) An *access-spec* shall appear only in the *specification-part* of a module.
- R508 *language-binding-spec* **is** BIND (C [, NAME = *scalar-char-initialization-expr*])
- C518 An entity with the BIND attribute shall be a common block, variable, or procedure.
- C519 A variable with the BIND attribute shall be declared in the specification part of a module.
- C520 A variable with the BIND attribute shall be interoperable (15.3).
- C521 Each variable of a common block with the BIND attribute shall be interoperable.
- C522 (R508) The *scalar-char-initialization-expr* shall be of default character kind.
- C523 An entity that has the CONTIGUOUS attribute shall be an array pointer or an assumed-shape array.
- R509 *dimension-spec* **is** DIMENSION (*array-spec*)
 or DIMENSION [(*array-spec*)] *lbracket co-array-spec rbracket*
- C524 (R501) A co-array that has the ALLOCATABLE attribute shall be specified with a *co-array-spec* that is a *deferred-co-shape-spec-list*.
- C525 A co-array shall be a component or a variable that is not a function result.
- C526 A co-array shall not be of type IMAGE_TEAM (13.8.3.7), C_PTR, or C_FUNPTR (15.3.3).
- C527 An entity whose type has a co-array ultimate component shall be a nonpointer nonallocatable scalar, shall not be a co-array, and shall not be a function result.
- C528 The SAVE attribute shall be specified for a co-array unless it is a dummy argument, declared

- in the main program, or allocatable.
- R510 *array-spec* **is** *explicit-shape-spec-list*
 or *assumed-shape-spec-list*
 or *deferred-shape-spec-list*
 or *assumed-size-spec*
 or *implied-shape-spec-list*
- R511 *co-array-spec* **is** *deferred-co-shape-spec-list*
 or *explicit-co-shape-spec*
- C529 The sum of the rank and co-rank of an entity shall not exceed fifteen.
- R512 *explicit-shape-spec* **is** [*lower-bound* :] *upper-bound*
- R513 *lower-bound* **is** *specification-expr*
- R514 *upper-bound* **is** *specification-expr*
- C530 (R512) An *explicit-shape-spec* whose bounds are not initialization expressions shall appear only in a subprogram or interface body.
- R515 *assumed-shape-spec* **is** [*lower-bound*] :
- R516 *deferred-shape-spec* **is** :
- C531 An array that has the POINTER or ALLOCATABLE attribute shall have an *array-spec* that is a *deferred-shape-spec-list*.
- R517 *assumed-size-spec* **is** [*explicit-shape-spec* ,]... [*lower-bound* :] *
- C532 An *assumed-size-spec* shall not appear except as the declaration of the array bounds of a dummy data argument.
- C533 An assumed-size array with INTENT (OUT) shall not be polymorphic, of a finalizable type, of a type with an ultimate allocatable component, or of a type for which default initialization is specified.
- R518 *implied-shape-spec* **is** [*lower-bound* :] *
- C534 An implied-shape array shall be a named constant.
- C535 An entity shall not have both the EXTERNAL attribute and the INTRINSIC attribute.
- R519 *deferred-co-shape-spec* **is** :
- C536 A co-array that has the ALLOCATABLE attribute shall have a *co-array-spec* that is a *deferred-co-shape-spec-list*.
- R520 *explicit-co-shape-spec* **is** [[*lower-co-bound* :] *upper-co-bound* ,]... [*lower-co-bound* :] *
- C537 A co-array that does not have the ALLOCATABLE attribute shall have a *co-array-spec* that is an *explicit-co-shape-spec*.
- R521 *lower-co-bound* **is** *specification-expr*
- R522 *upper-co-bound* **is** *specification-expr*
- C538 (R520) A *lower-co-bound* or *upper-co-bound* that is not an initialization expression shall appear only in a subprogram or interface body.
- R523 *intent-spec* **is** IN
 or OUT
 or INOUT
- C539 An entity with the INTENT attribute shall be a dummy data object or a dummy procedure pointer.
- C540 (R523) A nonpointer object with the INTENT (IN) attribute shall not appear in a variable definition context (16.6.7).
- C541 A pointer with the INTENT (IN) attribute shall not appear in a pointer association context (16.6.8).
- C542 If the name of a generic intrinsic procedure is explicitly declared to have the INTRINSIC attribute, and it is also the generic name of one or more generic interfaces (12.4.3.2) accessible in

the same scoping unit, the procedures in the interfaces and the specific intrinsic procedures shall all be functions or all be subroutines, and the characteristics of the specific intrinsic procedures and the procedures in the interfaces shall differ as specified in 12.4.3.3.4.

- C543 An entity with the OPTIONAL attribute shall be a dummy argument.
- C544 An entity with the PARAMETER attribute shall not be a variable, a co-array, or a procedure.
- C545 An entity with the POINTER attribute shall not have the ALLOCATABLE, INTRINSIC, or TARGET attribute.
- C546 A procedure with the POINTER attribute shall have the EXTERNAL attribute.
- C547 A co-array shall not have the POINTER attribute.
- C548 The PROTECTED attribute shall be specified only in the specification part of a module.
- C549 An entity with the PROTECTED attribute shall be a procedure pointer or variable.
- C550 An entity with the PROTECTED attribute shall not be in a common block.
- C551 A nonpointer object that has the PROTECTED attribute and is accessed by use association shall not appear in a variable definition context (16.6.7) or as the *data-target* or *proc-target* in a *pointer-assignment-stmt*.
- C552 A pointer that has the PROTECTED attribute and is accessed by use association shall not appear in a pointer association context (16.6.8).
- (1) A *pointer-object* in a *pointer-assignment-stmt* or *nullify-stmt*,
 - (2) An *allocate-object* in an *allocate-stmt* or *deallocate-stmt*, or
 - (3) An actual argument in a reference to a procedure if the associated dummy argument is a pointer with the INTENT(OUT) or INTENT(INOUT) attribute.
- C553 An entity with the SAVE attribute shall be a common block, variable, or procedure pointer.
- C554 The SAVE attribute shall not be specified for a dummy argument, a function result, an automatic data object, or an object that is in a common block.
- C555 An entity with the TARGET attribute shall be a variable.
- C556 An entity with the TARGET attribute shall not have the POINTER attribute.
- C557 An entity with the VALUE attribute shall be a scalar dummy data object.
- C558 An entity with the VALUE attribute shall not have the ALLOCATABLE, INTENT(INOUT), INTENT(OUT), POINTER, or VOLATILE attributes.
- C559 If an entity has the VALUE attribute, any length type parameter value in its declaration shall be omitted or specified by an initialization expression.
- C560 An entity with the VOLATILE attribute shall be a variable that is not an INTENT(IN) dummy argument.
- R524 *access-stmt* **is** *access-spec* [[*::*] *access-id-list*]
- R525 *access-id* **is** *use-name*
 or *generic-spec*
- C561 (R524) An *access-stmt* shall appear only in the *specification-part* of a module. Only one accessibility statement with an omitted *access-id-list* is permitted in the *specification-part* of a module.
- C562 (R525) Each *use-name* shall be the name of a named variable, procedure, derived type, named constant, namelist group, or macro.
- R526 *allocatable-stmt* **is** ALLOCATABLE [*::*] *allocatable-decl-list*
- R527 *allocatable-decl* **is** *object-name* [(*array-spec*)] ■
 ■ [*lbracket co-array-spec rbracket*]
- R528 *asynchronous-stmt* **is** ASYNCHRONOUS [*::*] *object-name-list*
- R529 *bind-stmt* **is** *language-binding-spec* [*::*] *bind-entity-list*
- R530 *bind-entity* **is** *entity-name*

- or** / *common-block-name* /
- C563 (R529) If the *language-binding-spec* has a NAME= specifier, the *bind-entity-list* shall consist of a single *bind-entity*.
- R531 *contiguous-stmt* **is** CONTIGUOUS [::] *object-name-list*
- R532 *data-stmt* **is** DATA *data-stmt-set* [[,] *data-stmt-set*] ...
- R533 *data-stmt-set* **is** *data-stmt-object-list* / *data-stmt-value-list* /
- R534 *data-stmt-object* **is** *variable*
or *data-implied-do*
- R535 *data-implied-do* **is** (*data-i-do-object-list* , *data-i-do-variable* = ■
■ *scalar-int-initialization-expr* , ■
■ *scalar-int-initialization-expr* ■
■ [, *scalar-int-initialization-expr*])
- R536 *data-i-do-object* **is** *array-element*
or *scalar-structure-component*
or *data-implied-do*
- R537 *data-i-do-variable* **is** *do-variable*
- C564 A *data-stmt-object* or *data-i-do-object* shall not be a co-indexed variable.
- C565 (R534) In a *variable* that is a *data-stmt-object*, any subscript, section subscript, substring starting point, and substring ending point shall be an initialization expression.
- C566 (R534) A variable whose designator appears as a *data-stmt-object* or a *data-i-do-object* shall not be a dummy argument, made accessible by use association or host association, in a named common block unless the DATA statement is in a block data program unit, in a blank common block, a function name, a function result name, an automatic object, or an allocatable variable.
- C567 (R534) A *data-i-do-object* or a *variable* that appears as a *data-stmt-object* shall not be an object designator in which a pointer appears other than as the entire rightmost *part-ref*.
- C568 (R536) The *array-element* shall be a variable.
- C569 (R536) The *scalar-structure-component* shall be a variable.
- C570 (R536) The *scalar-structure-component* shall contain at least one *part-ref* that contains a *subscript-list*.
- C571 (R536) In an *array-element* or *scalar-structure-component* that is a *data-i-do-object*, any subscript shall be an initialization expression, and any primary within that subscript that is a *data-i-do-variable* shall be a DO variable of this *data-implied-do* or of a containing *data-implied-do*.
- R538 *data-stmt-value* **is** [*data-stmt-repeat* *] *data-stmt-constant*
- R539 *data-stmt-repeat* **is** *scalar-int-constant*
or *scalar-int-constant-subobject*
- C572 (R539) The *data-stmt-repeat* shall be positive or zero. If the *data-stmt-repeat* is a named constant, it shall have been declared previously in the scoping unit or made accessible by use association or host association.
- R540 *data-stmt-constant* **is** *scalar-constant*
or *scalar-constant-subobject*
or *signed-int-literal-constant*
or *signed-real-literal-constant*
or *null-init*
or *initial-data-target*
or *structure-constructor*
- C573 (R540) If a DATA statement constant value is a named constant or a structure constructor, the named constant or derived type shall have been declared previously in the scoping unit or made

- accessible by use or host association.
- C574 (R540) If a *data-stmt-constant* is a *structure-constructor*, it shall be an initialization expression.
- R541 *int-constant-subobject* **is** *constant-subobject*
- C575 (R541) *int-constant-subobject* shall be of type integer.
- R542 *constant-subobject* **is** *designator*
- C576 (R542) *constant-subobject* shall be a subobject of a constant.
- C577 (R542) Any subscript, substring starting point, or substring ending point shall be an initialization expression.
- R543 *dimension-stmt* **is** DIMENSION [::] *dimension-decl-list*
- R544 *dimension-decl* **is** *array-name* (*array-spec*)
or *co-name* [(*array-spec*)] *lbracket co-array-spec rbracket*
- R545 *intent-stmt* **is** INTENT (*intent-spec*) [::] *dummy-arg-name-list*
- R546 *optional-stmt* **is** OPTIONAL [::] *dummy-arg-name-list*
- R547 *parameter-stmt* **is** PARAMETER (*named-constant-def-list*)
- R548 *named-constant-def* **is** *named-constant* = *initialization-expr*
- R549 *pointer-stmt* **is** POINTER [::] *pointer-decl-list*
- R550 *pointer-decl* **is** *object-name* [(*deferred-shape-spec-list*)]
or *proc-entity-name*
- R551 *protected-stmt* **is** PROTECTED [::] *entity-name-list*
- R552 *save-stmt* **is** SAVE [[::] *saved-entity-list*]
- R553 *saved-entity* **is** *object-name*
or *proc-pointer-name*
or / *common-block-name* /
- R554 *proc-pointer-name* **is** *name*
- C578 (R552) If a SAVE statement with an omitted saved entity list appears in a scoping unit, no other appearance of the SAVE *attr-spec* or SAVE statement is permitted in that scoping unit.
- R555 *target-stmt* **is** TARGET [::] *target-decl-list*
- R556 *target-decl* **is** *object-name* [(*array-spec*)] ■
■ [*lbracket co-array-spec rbracket*]
- R557 *value-stmt* **is** VALUE [::] *dummy-arg-name-list*
- R558 *volatile-stmt* **is** VOLATILE [::] *object-name-list*
- R559 *implicit-stmt* **is** IMPLICIT *implicit-spec-list*
or IMPLICIT NONE
- R560 *implicit-spec* **is** *declaration-type-spec* (*letter-spec-list*)
- R561 *letter-spec* **is** *letter* [- *letter*]
- C579 (R559) If IMPLICIT NONE is specified in a scoping unit, it shall precede any PARAMETER statements that appear in the scoping unit and there shall be no other IMPLICIT statements in the scoping unit.
- C580 (R561) If the minus and second *letter* appear, the second letter shall follow the first letter alphabetically.
- R562 *namelist-stmt* **is** NAMELIST ■
■ / *namelist-group-name* / *namelist-group-object-list* ■
■ [[,] / *namelist-group-name* / ■
■ *namelist-group-object-list*] ...
- C581 (R562) The *namelist-group-name* shall not be a name accessed by use association.

- R563 *namelist-group-object* **is** *variable-name*
- C582 (R563) A *namelist-group-object* shall not be an assumed-size array.
- C583 (R562) A *namelist-group-object* shall not have the PRIVATE attribute if the *namelist-group-name* has the PUBLIC attribute.
- R564 *equivalence-stmt* **is** EQUIVALENCE *equivalence-set-list*
- R565 *equivalence-set* **is** (*equivalence-object* , *equivalence-object-list*)
- R566 *equivalence-object* **is** *variable-name*
or *array-element*
or *substring*
- C584 (R566) An *equivalence-object* shall not be a designator with a base object that is a dummy argument, a pointer, an allocatable variable, a derived-type object that has an allocatable ultimate component, an object of a nonsequence derived type, an object of a derived type that has a pointer at any level of component selection, an automatic object, a function name, an entry name, a result name, a variable with the BIND attribute, a variable in a common block that has the BIND attribute, or a named constant.
- C585 (R566) An *equivalence-object* shall not be a designator that has more than one *part-ref*.
- C586 (R566) An *equivalence-object* shall not be a co-array or a subobject thereof.
- C587 (R566) An *equivalence-object* shall not have the TARGET attribute.
- C588 (R566) Each subscript or substring range expression in an *equivalence-object* shall be an integer initialization expression (7.1.7).
- C589 (R565) If an *equivalence-object* is of type default integer, default real, double precision real, default complex, default logical, default bits, or numeric sequence type, all of the objects in the equivalence set shall be of these types.
- C590 (R565) If an *equivalence-object* is of type default character or character sequence type, all of the objects in the equivalence set shall be of these types.
- C591 (R565) If an *equivalence-object* is of a sequence derived type that is not a numeric sequence or character sequence type, all of the objects in the equivalence set shall be of the same type with the same type parameter values.
- C592 (R565) If an *equivalence-object* is of an intrinsic type other than default integer, default real, double precision real, default complex, default logical, or default character, all of the objects in the equivalence set shall be of the same type with the same kind type parameter value.
- C593 (R566) If an *equivalence-object* has the PROTECTED attribute, all of the objects in the equivalence set shall have the PROTECTED attribute.
- C594 (R566) The name of an *equivalence-object* shall not be a name made accessible by use association.
- C595 (R566) A *substring* shall not have length zero.
- R567 *common-stmt* **is** COMMON ■
 ■ [/ [*common-block-name*] /] *common-block-object-list* ■
 ■ [[,] / [*common-block-name*] / ■
 ■ *common-block-object-list*] ...
- R568 *common-block-object* **is** *variable-name* [(*array-spec*)]
or *proc-pointer-name*
- C596 (R568) An *array-spec* in a *common-block-object* shall be an *explicit-shape-spec-list*.
- C597 (R568) Only one appearance of a given *variable-name* or *proc-pointer-name* is permitted in all *common-block-object-lists* within a scoping unit.
- C598 (R568) A *common-block-object* shall not be a dummy argument, an allocatable variable, a derived-type object with an ultimate component that is allocatable, an automatic object, a function name, an entry name, a variable with the BIND attribute, a co-array, or a result name.
- C599 (R568) If a *common-block-object* is of a derived type, it shall be a sequence type (4.5.2.3) or a

type with the BIND attribute and it shall have no default initialization.

C5100 (R568) A *variable-name* or *proc-pointer-name* shall not be a name made accessible by use association.

Clause 6:

- R601 *variable* **is** *designator*
or *expr*
- C601 (R601) *designator* shall not be a constant or a subobject of a constant.
- C602 (R601) *expr* shall be a reference to a function that has a pointer result.
- R602 *variable-name* **is** *name*
- C603 (R602) *variable-name* shall be the name of a variable.
- R603 *designator* **is** *object-name*
or *array-element*
or *array-section*
or *complex-part-designator*
or *structure-component*
or *substring*
- R604 *logical-variable* **is** *variable*
- C604 (R604) *logical-variable* shall be of type logical.
- R605 *default-logical-variable* **is** *variable*
- C605 (R605) *default-logical-variable* shall be of type default logical.
- R606 *char-variable* **is** *variable*
- C606 (R606) *char-variable* shall be of type character.
- R607 *default-char-variable* **is** *variable*
- C607 (R607) *default-char-variable* shall be of type default character.
- R608 *int-variable* **is** *variable*
- C608 (R608) *int-variable* shall be of type integer.
- R609 *substring* **is** *parent-string* (*substring-range*)
- R610 *parent-string* **is** *scalar-variable-name*
or *array-element*
or *scalar-structure-component*
or *scalar-constant*
- R611 *substring-range* **is** [*scalar-int-expr*] : [*scalar-int-expr*]
- C609 (R610) *parent-string* shall be of type character.
- R612 *data-ref* **is** *part-ref* [% *part-ref*] ...
- R613 *part-ref* **is** *part-name* [(*section-subscript-list*)] [*image-selector*]
- C610 (R612) Each *part-name* except the rightmost shall be of derived type.
- C611 (R612) Each *part-name* except the leftmost shall be the name of a component of the declared type of the preceding *part-name*.
- C612 (R612) If the rightmost *part-name* is of abstract type, *data-ref* shall be polymorphic.
- C613 (R612) The leftmost *part-name* shall be the name of a data object.
- C614 (R613) If a *section-subscript-list* appears, the number of *section-subscripts* shall equal the rank of *part-name*.
- C615 (R613) If *image-selector* appears and *part-name* is an array, *section-subscript-list* shall appear.
- C616 (R612) A *data-ref* that is a co-indexed object shall not be of a type that has a pointer ultimate component.
- C617 (R612) If *image-selector* appears, *data-ref* shall not be, or have a direct component, of type

- IMAGE_TEAM (13.8.3.7), C_PTR, or C_FUNPTR (15.3.3).
- C618 (R612) There shall not be more than one *part-ref* with nonzero rank. A *part-name* to the right of a *part-ref* with nonzero rank shall not have the ALLOCATABLE or POINTER attribute.
- R614 *structure-component* is *data-ref*
- C619 (R614) There shall be more than one *part-ref* and the rightmost *part-ref* shall be of the form *part-name*.
- R615 *complex-part-designator* is *designator* % RE
or *designator* % IM
- C620 (R615) The *designator* shall be of complex type.
- R616 *type-param-inquiry* is *designator* % *type-param-name*
- C621 (R616) The *type-param-name* shall be the name of a type parameter of the declared type of the object designated by the *designator*.
- R617 *array-element* is *data-ref*
- C622 (R617) Every *part-ref* shall have rank zero and the last *part-ref* shall contain a *subscript-list*.
- R618 *array-section* is *data-ref* [(*substring-range*)]
or *complex-part-designator*
- C623 (R618) Exactly one *part-ref* shall have nonzero rank, and either the final *part-ref* shall have a *section-subscript-list* with nonzero rank, the *array-section* is a *complex-part-designator* that is an array, or another *part-ref* shall have nonzero rank.
- C624 (R618) If a *substring-range* appears, the rightmost *part-name* shall be of type character.
- R619 *subscript* is *scalar-int-expr*
- R620 *section-subscript* is *subscript*
or *subscript-triplet*
or *vector-subscript*
- R621 *subscript-triplet* is [*subscript*] : [*subscript*] [: *stride*]
- R622 *stride* is *scalar-int-expr*
- R623 *vector-subscript* is *int-expr*
- C625 (R623) A *vector-subscript* shall be an integer array expression of rank one.
- C626 (R621) The second subscript shall not be omitted from a *subscript-triplet* in the last dimension of an assumed-size array.
- R624 *image-selector* is *lbracket co-subscript-list rbracket*
- R625 *co-subscript* is *scalar-int-expr*
- R626 *allocate-stmt* is ALLOCATE ([*type-spec* ::] *allocation-list* ■
■ [*alloc-opt-list*])
- R627 *alloc-opt* is ERRMSG = *errmsg-variable*
or MOLD = *source-expr*
or SOURCE = *source-expr*
or STAT = *stat-variable*
- R628 *stat-variable* is *scalar-int-variable*
- R629 *errmsg-variable* is *scalar-default-char-variable*
- R630 *source-expr* is *expr*
- R631 *allocation* is *allocate-object* [(*allocate-shape-spec-list*)] ■
■ [*lbracket allocate-co-array-spec rbracket*]
- R632 *allocate-object* is *variable-name*
or *structure-component*
- R633 *allocate-shape-spec* is [*lower-bound-expr* :] *upper-bound-expr*
- R634 *lower-bound-expr* is *scalar-int-expr*

- R635 *upper-bound-expr* **is** *scalar-int-expr*
- R636 *allocate-co-array-spec* **is** [*allocate-co-shape-spec-list* ,] [*lower-bound-expr* :] *
- R637 *allocate-co-shape-spec* **is** [*lower-bound-expr* :] *upper-bound-expr*
- C627 (R632) Each *allocate-object* shall be a nonprocedure pointer or an allocatable variable.
- C628 (R626) If any *allocate-object* in the statement has a deferred type parameter, either *type-spec* or *source-expr* shall appear.
- C629 (R626) If a *type-spec* appears, it shall specify a type with which each *allocate-object* is type compatible.
- C630 (R626) If any *allocate-object* is unlimited polymorphic or is of abstract type, either *type-spec* or *source-expr* shall appear.
- C631 (R626) A *type-param-value* in a *type-spec* shall be an asterisk if and only if each *allocate-object* is a dummy argument for which the corresponding type parameter is assumed.
- C632 (R626) If a *type-spec* appears, the kind type parameter values of each *allocate-object* shall be the same as the corresponding type parameter values of the *type-spec*.
- C633 (R631) An *allocate-shape-spec-list* shall appear if and only if the *allocate-object* is an array. An *allocate-co-array-spec* shall appear if and only if the *allocate-object* is a co-array.
- C634 (R631) The number of *allocate-shape-specs* in an *allocate-shape-spec-list* shall be the same as the rank of the *allocate-object*. The number of *allocate-co-shape-specs* in an *allocate-co-array-spec* shall be one less than the co-rank of the *allocate-object*.
- C635 (R627) No *alloc-opt* shall appear more than once in a given *alloc-opt-list*.
- C636 (R626) At most one of *source-expr* and *type-spec* shall appear.
- C637 (R626) Each *allocate-object* shall be type compatible (4.3.1.3) with *source-expr*. If SOURCE= appears, *source-expr* shall be a scalar or have the same rank as each *allocate-object*.
- C638 (R626) Corresponding kind type parameters of *allocate-object* and *source-expr* shall have the same values.
- C639 (R626) *type-spec* shall not specify a type that has a co-array ultimate component.
- C640 (R630) The declared type of *source-expr* shall not have a co-array ultimate component.
- C641 (R632) An *allocate-object* shall not be a co-indexed object.
- R638 *nullify-stmt* **is** NULLIFY (*pointer-object-list*)
- R639 *pointer-object* **is** *variable-name*
or *structure-component*
or *proc-pointer-name*
- C642 (R639) Each *pointer-object* shall have the POINTER attribute.
- R640 *deallocate-stmt* **is** DEALLOCATE (*allocate-object-list* [, *dealloc-opt-list*])
- C643 (R640) Each *allocate-object* shall be a nonprocedure pointer or an allocatable variable.
- R641 *dealloc-opt* **is** STAT = *stat-variable*
or ERRMSG = *errmsg-variable*
- C644 (R641) No *dealloc-opt* shall appear more than once in a given *dealloc-opt-list*.

Clause 7:

- R701 *primary* **is** *constant*
or *designator*
or *array-constructor*
or *structure-constructor*
or *function-reference*
or *type-param-inquiry*
or *type-param-name*

- or** (*expr*)
- C701 (R701) The *type-param-name* shall be the name of a type parameter.
- C702 (R701) The *designator* shall not be a whole assumed-size array.
- R702 *level-1-expr* **is** [*defined-unary-op*] *primary*
- R703 *defined-unary-op* **is** . *letter* [*letter*]
- C703 (R703) A *defined-unary-op* shall not contain more than 63 letters and shall not be the same as any *intrinsic-operator* or *logical-literal-constant*.
- R704 *mult-operand* **is** *level-1-expr* [*power-op mult-operand*]
- R705 *add-operand* **is** [*add-operand mult-op*] *mult-operand*
- R706 *level-2-expr* **is** [[*level-2-expr*] *add-op*] *add-operand*
- R707 *power-op* **is** **
- R708 *mult-op* **is** *
- or** /
- R709 *add-op* **is** +
- or** -
- R710 *level-3-expr* **is** [*level-3-expr concat-op*] *level-2-expr*
- R711 *concat-op* **is** //
- R712 *level-4-expr* **is** [*level-3-expr rel-op*] *level-3-expr*
- R713 *rel-op* **is** .EQ.
- or** .NE.
- or** .LT.
- or** .LE.
- or** .GT.
- or** .GE.
- or** ==
- or** /=
- or** <
- or** <=
- or** >
- or** >=
- R714 *and-operand* **is** [*not-op*] *level-4-expr*
- R715 *or-operand* **is** [*or-operand and-op*] *and-operand*
- R716 *equiv-operand* **is** [*equiv-operand or-op*] *or-operand*
- R717 *level-5-expr* **is** [*level-5-expr equiv-op*] *equiv-operand*
- R718 *not-op* **is** .NOT.
- R719 *and-op* **is** .AND.
- R720 *or-op* **is** .OR.
- R721 *equiv-op* **is** .EQV.
- or** .NEQV.
- or** .XOR.
- R722 *expr* **is** [*expr defined-binary-op*] *level-5-expr*
- R723 *defined-binary-op* **is** . *letter* [*letter*]
- C704 (R723) A *defined-binary-op* shall not contain more than 63 letters and shall not be the same as any *intrinsic-operator* or *logical-literal-constant*.
- R724 *logical-expr* **is** *expr*
- C705 (R724) *logical-expr* shall be of type logical.

- R725 *char-expr* **is** *expr*
- C706 (R725) *char-expr* shall be of type character.
- R726 *default-char-expr* **is** *expr*
- C707 (R726) *default-char-expr* shall be of type default character.
- R727 *int-expr* **is** *expr*
- C708 (R727) *int-expr* shall be of type integer.
- R728 *numeric-expr* **is** *expr*
- C709 (R728) *numeric-expr* shall be of type integer, real, or complex.
- C710 The kind type parameter for the result of a bits concatenation operation expression shall be a bits kind type parameter value supported by the processor.
- R729 *specification-expr* **is** *scalar-int-expr*
- C711 (R729) The *scalar-int-expr* shall be a restricted expression.
- R730 *initialization-expr* **is** *expr*
- C712 (R730) *initialization-expr* shall be an initialization expression.
- R731 *char-initialization-expr* **is** *char-expr*
- C713 (R731) *char-initialization-expr* shall be an initialization expression.
- R732 *int-initialization-expr* **is** *int-expr*
- C714 (R732) *int-initialization-expr* shall be an initialization expression.
- R733 *logical-initialization-expr* **is** *logical-expr*
- C715 (R733) *logical-initialization-expr* shall be an initialization expression.
- R734 *assignment-stmt* **is** *variable = expr*
- C716 (R734) The *variable* shall not be a whole assumed-size array.
- R735 *pointer-assignment-stmt* **is** *data-pointer-object* [(*bounds-spec-list*)] => *data-target*
or *data-pointer-object* (*bounds-remapping-list*) => *data-target*
or *proc-pointer-object* => *proc-target*
- R736 *data-pointer-object* **is** *variable-name*
or *scalar-variable* % *data-pointer-component-name*
- C717 (R735) If *data-target* is not unlimited polymorphic, *data-pointer-object* shall be type compatible (4.3.1.3) with it and the corresponding kind type parameters shall be equal.
- C718 (R735) If *data-target* is unlimited polymorphic, *data-pointer-object* shall be unlimited polymorphic, of a sequence derived type, or of a type with the BIND attribute.
- C719 (R735) If *bounds-spec-list* is specified, the number of *bounds-specs* shall equal the rank of *data-pointer-object*.
- C720 (R735) If *bounds-remapping-list* is specified, the number of *bounds-remappings* shall equal the rank of *data-pointer-object*.
- C721 (R735) If *bounds-remapping-list* is not specified, the ranks of *data-pointer-object* and *data-target* shall be the same.
- C722 (R736) A *variable-name* shall have the POINTER attribute.
- C723 (R736) A *scalar-variable* shall be a *data-ref*.
- C724 (R736) A *data-pointer-component-name* shall be the name of a component of *scalar-variable* that is a data pointer.
- C725 (R736) A *data-pointer-object* shall not be a co-indexed object.
- R737 *bounds-spec* **is** *lower-bound-expr* :
- R738 *bounds-remapping* **is** *lower-bound-expr* : *upper-bound-expr*
- R739 *data-target* **is** *variable*
or *expr*
- C726 (R739) A *variable* shall have either the TARGET or POINTER attribute, and shall not be an

- array section with a vector subscript.
- C727 (R739) A *data-target* shall not be a co-indexed object.
- C728 (R739) An *expr* shall be a reference to a function whose result is a data pointer.
- R740 *proc-pointer-object* **is** *proc-pointer-name*
 or *proc-component-ref*
- R741 *proc-component-ref* **is** *scalar-variable* % *procedure-component-name*
- C729 (R741) The *scalar-variable* shall be a *data-ref*.
- C730 (R741) The *procedure-component-name* shall be the name of a procedure pointer component of the declared type of *scalar-variable*.
- R742 *proc-target* **is** *expr*
 or *procedure-name*
 or *proc-component-ref*
- C731 (R742) An *expr* shall be a reference to a function whose result is a procedure pointer.
- C732 (R742) A *procedure-name* shall be the name of an external, internal, module, or dummy procedure, a procedure pointer, or a specific intrinsic function listed in 13.6 and not marked with a bullet (●).
- C733 (R742) The *proc-target* shall not be a nonintrinsic elemental procedure.
- R743 *where-stmt* **is** WHERE (*mask-expr*) *where-assignment-stmt*
- R744 *where-construct* **is** *where-construct-stmt*
 [*where-body-construct*] ...
 [*masked-elsewhere-stmt*
 [*where-body-construct*] ...] ...
 [*elsewhere-stmt*
 [*where-body-construct*] ...]
 end-where-stmt
- R745 *where-construct-stmt* **is** [*where-construct-name*:] WHERE (*mask-expr*)
- R746 *where-body-construct* **is** *where-assignment-stmt*
 or *where-stmt*
 or *where-construct*
- R747 *where-assignment-stmt* **is** *assignment-stmt*
- R748 *mask-expr* **is** *logical-expr*
- R749 *masked-elsewhere-stmt* **is** ELSEWHERE (*mask-expr*) [*where-construct-name*]
- R750 *elsewhere-stmt* **is** ELSEWHERE [*where-construct-name*]
- R751 *end-where-stmt* **is** END WHERE [*where-construct-name*]
- C734 (R747) A *where-assignment-stmt* that is a defined assignment shall be elemental.
- C735 (R744) If the *where-construct-stmt* is identified by a *where-construct-name*, the corresponding *end-where-stmt* shall specify the same *where-construct-name*. If the *where-construct-stmt* is not identified by a *where-construct-name*, the corresponding *end-where-stmt* shall not specify a *where-construct-name*. If an *elsewhere-stmt* or a *masked-elsewhere-stmt* is identified by a *where-construct-name*, the corresponding *where-construct-stmt* shall specify the same *where-construct-name*.
- C736 (R746) A statement that is part of a *where-body-construct* shall not be a branch target statement.
- R752 *forall-construct* **is** *forall-construct-stmt*
 [*forall-body-construct*] ...
 end-forall-stmt
- R753 *forall-construct-stmt* **is** [*forall-construct-name* :] FORALL *forall-header*
- R754 *forall-header* **is** (*forall-triplet-spec-list* [, *scalar-mask-expr*])

- R755 *forall-triplet-spec* **is** *index-name = subscript : subscript [: stride]*
- R756 *forall-body-construct* **is** *forall-assignment-stmt*
 or *where-stmt*
 or *where-construct*
 or *forall-construct*
 or *forall-stmt*
- R757 *forall-assignment-stmt* **is** *assignment-stmt*
 or *pointer-assignment-stmt*
- R758 *end-forall-stmt* **is** END FORALL [*forall-construct-name*]
- C737 (R758) If the *forall-construct-stmt* has a *forall-construct-name*, the *end-forall-stmt* shall have the same *forall-construct-name*. If the *end-forall-stmt* has a *forall-construct-name*, the *forall-construct-stmt* shall have the same *forall-construct-name*.
- C738 (R754) The *scalar-mask-expr* shall be scalar and of type logical.
- C739 (R754) Any procedure referenced in the *scalar-mask-expr*, including one referenced by a defined operation, shall be a pure procedure (12.7).
- C740 (R755) The *index-name* shall be a named scalar variable of type integer.
- C741 (R755) A *subscript* or *stride* in a *forall-triplet-spec* shall not contain a reference to any *index-name* in the *forall-triplet-spec-list* in which it appears.
- C742 (R756) A statement in a *forall-body-construct* shall not define an *index-name* of the *forall-construct*.
- C743 (R756) Any procedure referenced in a *forall-body-construct*, including one referenced by a defined operation, assignment, or finalization, shall be a pure procedure.
- C744 (R756) A *forall-body-construct* shall not be a branch target.
- C745 (R757) The *variable* in an *assignment-stmt* that is a *forall-assignment-stmt* shall be a *designator*.
- R759 *forall-stmt* **is** FORALL *forall-header forall-assignment-stmt*

Clause 8:

- R801 *block* **is** [*execution-part-construct*] ...
- R802 *associate-construct* **is** *associate-stmt*
 block
 end-associate-stmt
- R803 *associate-stmt* **is** [*associate-construct-name* :] ASSOCIATE ■
 ■ (*association-list*)
- R804 *association* **is** *associate-name => selector*
- R805 *selector* **is** *expr*
 or *variable*
- C801 (R804) If *selector* is not a *variable* or is a *variable* that has a vector subscript, *associate-name* shall not appear in a variable definition context (16.6.7).
- C802 (R804) An *associate-name* shall not be the same as another *associate-name* in the same *associate-stmt*.
- C803 (R805) *expr* shall not be a reference to a function that has a pointer result.
- R806 *end-associate-stmt* **is** END ASSOCIATE [*associate-construct-name*]
- C804 (R806) If the *associate-stmt* of an *associate-construct* specifies an *associate-construct-name*, the corresponding *end-associate-stmt* shall specify the same *associate-construct-name*. If the *associate-stmt* of an *associate-construct* does not specify an *associate-construct-name*, the corresponding *end-associate-stmt* shall not specify an *associate-construct-name*.
- R807 *block-construct* **is** *block-stmt*
 [*specification-part*]

- block*
end-block-stmt
- R808 *block-stmt* **is** [*block-construct-name* :] BLOCK
- R809 *end-block-stmt* **is** END BLOCK [*block-construct-name*]
- C805 (R807) The *specification-part* of a BLOCK construct shall not contain a COMMON statement, IMPLICIT statement, INTENT statement, OPTIONAL statement, or USE statement.
- C806 (R807) If the *block-stmt* of a *block-construct* specifies a *block-construct-name*, the corresponding *end-block-stmt* shall specify the same *block-construct-name*. If the *block-stmt* does not specify a *block-construct-name*, the corresponding *end-block-stmt* shall not specify a *block-construct-name*.
- R810 *case-construct* **is** *select-case-stmt*
 [*case-stmt*
 block] ...
 end-select-stmt
- R811 *select-case-stmt* **is** [*case-construct-name* :] SELECT CASE (*case-expr*)
- R812 *case-stmt* **is** CASE *case-selector* [*case-construct-name*]
- R813 *end-select-stmt* **is** END SELECT [*case-construct-name*]
- C807 (R810) If the *select-case-stmt* of a *case-construct* specifies a *case-construct-name*, the corresponding *end-select-stmt* shall specify the same *case-construct-name*. If the *select-case-stmt* of a *case-construct* does not specify a *case-construct-name*, the corresponding *end-select-stmt* shall not specify a *case-construct-name*. If a *case-stmt* specifies a *case-construct-name*, the corresponding *select-case-stmt* shall specify the same *case-construct-name*.
- R814 *case-expr* **is** *scalar-int-expr*
 or *scalar-char-expr*
 or *scalar-logical-expr*
- R815 *case-selector* **is** (*case-value-range-list*)
 or DEFAULT
- C808 (R810) No more than one of the selectors of one of the CASE statements shall be DEFAULT.
- R816 *case-value-range* **is** *case-value*
 or *case-value* :
 or : *case-value*
 or *case-value* : *case-value*
- R817 *case-value* **is** *scalar-int-initialization-expr*
 or *scalar-char-initialization-expr*
 or *scalar-logical-initialization-expr*
- C809 (R810) For a given *case-construct*, each *case-value* shall be of the same type as *case-expr*. For character type, the kind type parameters shall be the same; character length differences are allowed.
- C810 (R810) A *case-value-range* using a colon shall not be used if *case-expr* is of type logical.
- C811 (R810) For a given *case-construct*, the *case-value-ranges* shall not overlap; that is, there shall be no possible value of the *case-expr* that matches more than one *case-value-range*.
- R818 *critical-construct* **is** *critical-stmt*
 block
 end-critical-stmt
- R819 *critical-stmt* **is** [*critical-construct-name* :] CRITICAL
- R820 *end-critical-stmt* **is** END CRITICAL [*critical-construct-name*]
- C812 (R818) If the *critical-stmt* of a *critical-construct* specifies a *critical-construct-name*, the corresponding *end-critical-stmt* shall specify the same *critical-construct-name*. If the *critical-stmt* of a

- critical-construct* does not specify a *critical-construct-name*, the corresponding *end-critical-stmt* shall not specify a *critical-construct-name*.
- C813 (R818) The *block* of a *critical-construct* shall not contain an image control statement.
- R821 *do-construct* **is** *block-do-construct*
 or *nonblock-do-construct*
- R822 *block-do-construct* **is** *do-stmt*
 do-block
 end-do
- R823 *do-stmt* **is** *label-do-stmt*
 or *nonlabel-do-stmt*
- R824 *label-do-stmt* **is** [*do-construct-name* :] DO *label* [*loop-control*]
- R825 *nonlabel-do-stmt* **is** [*do-construct-name* :] DO [*loop-control*]
- R826 *loop-control* **is** [,] *do-variable* = *scalar-int-expr*, *scalar-int-expr* ■
 ■ [, *scalar-int-expr*]
 or [,] WHILE (*scalar-logical-expr*)
 or [,] CONCURRENT *forall-header*
- R827 *do-variable* **is** *scalar-int-variable-name*
- C814 (R827) The *do-variable* shall be a variable of type integer.
- R828 *do-block* **is** *block*
- R829 *end-do* **is** *end-do-stmt*
 or *continue-stmt*
- R830 *end-do-stmt* **is** END DO [*do-construct-name*]
- C815 (R822) If the *do-stmt* of a *block-do-construct* specifies a *do-construct-name*, the corresponding *end-do* shall be an *end-do-stmt* specifying the same *do-construct-name*. If the *do-stmt* of a *block-do-construct* does not specify a *do-construct-name*, the corresponding *end-do* shall not specify a *do-construct-name*.
- C816 (R822) If the *do-stmt* is a *nonlabel-do-stmt*, the corresponding *end-do* shall be an *end-do-stmt*.
- C817 (R822) If the *do-stmt* is a *label-do-stmt*, the corresponding *end-do* shall be identified with the same *label*.
- R831 *nonblock-do-construct* **is** *action-term-do-construct*
 or *outer-shared-do-construct*
- R832 *action-term-do-construct* **is** *label-do-stmt*
 do-body
 do-term-action-stmt
- R833 *do-body* **is** [*execution-part-construct*] ...
- R834 *do-term-action-stmt* **is** *action-stmt*
- C818 (R834) A *do-term-action-stmt* shall not be a *continue-stmt*, a *goto-stmt*, a *return-stmt*, a *stop-stmt*, an *exit-stmt*, a *cycle-stmt*, an *end-function-stmt*, an *end-subroutine-stmt*, an *end-program-stmt*, or an *arithmetic-if-stmt*.
- C819 (R831) The *do-term-action-stmt* shall be identified with a label and the corresponding *label-do-stmt* shall refer to the same label.
- R835 *outer-shared-do-construct* **is** *label-do-stmt*
 do-body
 shared-term-do-construct
- R836 *shared-term-do-construct* **is** *outer-shared-do-construct*
 or *inner-shared-do-construct*
- R837 *inner-shared-do-construct* **is** *label-do-stmt*

- do-body*
do-term-shared-stmt
- R838 *do-term-shared-stmt* **is** *action-stmt*
- C820 (R838) A *do-term-shared-stmt* shall not be a *goto-stmt*, a *return-stmt*, a *stop-stmt*, an *exit-stmt*, a *cycle-stmt*, an *end-function-stmt*, an *end-subroutine-stmt*, an *end-program-stmt*, or an *arithmetic-if-stmt*.
- C821 (R836) The *do-term-shared-stmt* shall be identified with a label and all of the *label-do-stmts* of the *inner-shared-do-construct* and *outer-shared-do-construct* shall refer to the same label.
- R839 *cycle-stmt* **is** CYCLE [*do-construct-name*]
- C822 (R839) If a *do-construct-name* appears, the CYCLE statement shall be within the range of that *do-construct*; otherwise, it shall be within the range of at least one *do-construct*.
- C823 (R839) A *cycle-stmt* shall not appear within the range of a DO CONCURRENT construct if it belongs to an outer construct.
- C824 A RETURN statement shall not appear within a DO CONCURRENT construct.
- C825 A branch (8.2) within a DO CONCURRENT construct shall not have a branch target that is outside the construct.
- C826 A reference to a nonpure procedure shall not appear within a DO CONCURRENT construct.
- C827 A reference to the procedure IEEE.GET_FLAG, IEEE.SET_HALTING_MODE, or IEEE.GET_-HALTING_MODE from the intrinsic module IEEE.EXCEPTIONS, shall not be appear within a DO CONCURRENT construct.
- R840 *if-construct* **is** *if-then-stmt*
 block
 [*else-if-stmt*
 block] ...
 [*else-stmt*
 block]
 end-if-stmt
- R841 *if-then-stmt* **is** [*if-construct-name* :] IF (*scalar-logical-expr*) THEN
- R842 *else-if-stmt* **is** ELSE IF (*scalar-logical-expr*) THEN [*if-construct-name*]
- R843 *else-stmt* **is** ELSE [*if-construct-name*]
- R844 *end-if-stmt* **is** END IF [*if-construct-name*]
- C828 (R840) If the *if-then-stmt* of an *if-construct* specifies an *if-construct-name*, the corresponding *end-if-stmt* shall specify the same *if-construct-name*. If the *if-then-stmt* of an *if-construct* does not specify an *if-construct-name*, the corresponding *end-if-stmt* shall not specify an *if-construct-name*. If an *else-if-stmt* or *else-stmt* specifies an *if-construct-name*, the corresponding *if-then-stmt* shall specify the same *if-construct-name*.
- R845 *if-stmt* **is** IF (*scalar-logical-expr*) *action-stmt*
- C829 (R845) The *action-stmt* in the *if-stmt* shall not be an *if-stmt*, *end-program-stmt*, *end-function-stmt*, or *end-subroutine-stmt*.
- R846 *select-type-construct* **is** *select-type-stmt*
 [*type-guard-stmt*
 block] ...
 end-select-type-stmt
- R847 *select-type-stmt* **is** [*select-construct-name* :] SELECT TYPE ■
 ■ ([*associate-name* =>] *selector*)
- C830 (R847) If *selector* is not a named *variable*, *associate-name* => shall appear.
- C831 (R847) If *selector* is not a *variable* or is a *variable* that has a vector subscript, *associate-name*

- shall not appear in a variable definition context (16.6.7).
- C832 (R847) The *selector* in a *select-type-stmt* shall be polymorphic.
- R848 *type-guard-stmt* **is** TYPE IS (*type-spec*) [*select-construct-name*]
 or CLASS IS (*derived-type-spec*) [*select-construct-name*]
 or CLASS DEFAULT [*select-construct-name*]
- C833 (R848) The *type-spec* or *derived-type-spec* shall specify that each length type parameter is assumed.
- C834 (R848) The *type-spec* or *derived-type-spec* shall not specify a sequence derived type or a type with the BIND attribute.
- C835 (R848) If *selector* is not unlimited polymorphic, the *type-spec* or *derived-type-spec* shall specify an extension of the declared type of *selector*.
- C836 (R848) For a given *select-type-construct*, the same type and kind type parameter values shall not be specified in more than one TYPE IS *type-guard-stmt* and shall not be specified in more than one CLASS IS *type-guard-stmt*.
- C837 (R848) For a given *select-type-construct*, there shall be at most one CLASS DEFAULT *type-guard-stmt*.
- R849 *end-select-type-stmt* **is** END SELECT [*select-construct-name*]
- C838 (R846) If the *select-type-stmt* of a *select-type-construct* specifies a *select-construct-name*, the corresponding *end-select-type-stmt* shall specify the same *select-construct-name*. If the *select-type-stmt* of a *select-type-construct* does not specify a *select-construct-name*, the corresponding *end-select-type-stmt* shall not specify a *select-construct-name*. If a *type-guard-stmt* specifies a *select-construct-name*, the corresponding *select-type-stmt* shall specify the same *select-construct-name*.
- R850 *exit-stmt* **is** EXIT [*construct-name*]
- C839 If a *construct-name* appears, the EXIT statement shall be within that construct; otherwise, it shall be within the range of at least one *do-construct*.
- C840 An *exit-stmt* shall not belong to a DO CONCURRENT construct, nor shall it appear within the range of a DO CONCURRENT construct if it belongs to a construct that contains that DO CONCURRENT construct.
- R851 *goto-stmt* **is** GO TO *label*
- C841 (R851) The *label* shall be the statement label of a branch target statement that appears in the same scoping unit as the *goto-stmt*.
- R852 *computed-goto-stmt* **is** GO TO (*label-list*) [,] *scalar-int-expr*
- C842 (R852) Each *label* in *label-list* shall be the statement label of a branch target statement that appears in the same scoping unit as the *computed-goto-stmt*.
- R853 *arithmetic-if-stmt* **is** IF (*scalar-numeric-expr*) *label* , *label* , *label*
- C843 (R853) Each *label* shall be the label of a branch target statement that appears in the same scoping unit as the *arithmetic-if-stmt*.
- C844 (R853) The *scalar-numeric-expr* shall not be of type complex.
- R854 *continue-stmt* **is** CONTINUE
- R855 *stop-stmt* **is** STOP [*stop-code*]
- R856 *stop-code* **is** *scalar-char-initialization-expr*
 or *scalar-int-initialization-expr*
- C845 (R856) The *scalar-char-initialization-expr* shall be of default kind.
- C846 (R856) The *scalar-int-initialization-expr* shall be of default kind.
- R857 *sync-all-stmt* **is** SYNC ALL [([*sync-stat-list*])]
- R858 *sync-stat* **is** STAT = *stat-variable*
 or ERRMSG = *errmsg-variable*
- C847 No specifier shall appear more than once in a given *sync-stat-list*.

- R859 *sync-team-stmt* is SYNC TEAM (*image-team* [, *sync-stat-list*])
- R860 *image-team* is *scalar-variable*
- C848 The *image-team* shall be a scalar variable of type IMAGE_TEAM from the intrinsic module ISO_FORTRAN_ENV.
- R861 *sync-images-stmt* is SYNC IMAGES (*image-set* [, *sync-stat-list*])
- R862 *image-set* is *int-expr*
or *
- C849 An *image-set* that is an *int-expr* shall be scalar or of rank one.
- R863 *notify-stmt* is NOTIFY (*image-set* [, *sync-stat-list*])
- R864 *query-stmt* is QUERY (*image-set* [, *query-spec-list*])
- R865 *query-spec* is READY = *scalar-logical-variable*
or *sync-stat*
- C850 (R864) No specifier shall appear more than once in a given *query-spec-list*.
- R866 *sync-memory-stmt* is SYNC MEMORY [([*sync-stat-list*])]

Clause 9:

- R901 *io-unit* is *file-unit-number*
or *
or *internal-file-variable*
- R902 *file-unit-number* is *scalar-int-expr*
- R903 *internal-file-variable* is *char-variable*
- C901 (R903) The *char-variable* shall not be an array section with a vector subscript.
- C902 (R903) The *char-variable* shall be of type default character, ASCII character, or ISO 10646 character.
- R904 *open-stmt* is OPEN (*connect-spec-list*)
- R905 *connect-spec* is [UNIT =] *file-unit-number*
or ACCESS = *scalar-default-char-expr*
or ACTION = *scalar-default-char-expr*
or ASYNCHRONOUS = *scalar-default-char-expr*
or BLANK = *scalar-default-char-expr*
or DECIMAL = *scalar-default-char-expr*
or DELIM = *scalar-default-char-expr*
or ENCODING = *scalar-default-char-expr*
or ERR = *label*
or FILE = *file-name-expr*
or FORM = *scalar-default-char-expr*
or IOMSG = *iomsg-variable*
or IOSTAT = *scalar-int-variable*
or NEWUNIT = *scalar-int-variable*
or PAD = *scalar-default-char-expr*
or POSITION = *scalar-default-char-expr*
or RECL = *scalar-int-expr*
or ROUND = *scalar-default-char-expr*
or SIGN = *scalar-default-char-expr*
or STATUS = *scalar-default-char-expr*
or TEAM = *image-team*
- R906 *file-name-expr* is *scalar-default-char-expr*

- R907 *iomsg-variable* **is** *scalar-default-char-variable*
- C903 No specifier shall appear more than once in a given *connect-spec-list*.
- C904 (R904) If the NEWUNIT= specifier does not appear, a *file-unit-number* shall be specified; if the optional characters UNIT= are omitted, the *file-unit-number* shall be the first item in the *connect-spec-list*.
- C905 (R904) The *label* used in the ERR= specifier shall be the statement label of a branch target statement that appears in the same scoping unit as the OPEN statement.
- C906 (R904) If a NEWUNIT= specifier appears, a *file-unit-number* shall not appear.
- R908 *close-stmt* **is** CLOSE (*close-spec-list*)
- R909 *close-spec* **is** [UNIT =] *file-unit-number*
or IOSTAT = *scalar-int-variable*
or IOMSG = *iomsg-variable*
or ERR = *label*
or STATUS = *scalar-default-char-expr*
- C907 No specifier shall appear more than once in a given *close-spec-list*.
- C908 A *file-unit-number* shall be specified in a *close-spec-list*; if the optional characters UNIT= are omitted, the *file-unit-number* shall be the first item in the *close-spec-list*.
- C909 (R909) The *label* used in the ERR= specifier shall be the statement label of a branch target statement that appears in the same scoping unit as the CLOSE statement.
- R910 *read-stmt* **is** READ (*io-control-spec-list*) [*input-item-list*]
or READ *format* [, *input-item-list*]
- R911 *write-stmt* **is** WRITE (*io-control-spec-list*) [*output-item-list*]
- R912 *print-stmt* **is** PRINT *format* [, *output-item-list*]
- R913 *io-control-spec* **is** [UNIT =] *io-unit*
or [FMT =] *format*
or [NML =] *namelist-group-name*
or ADVANCE = *scalar-default-char-expr*
or ASYNCHRONOUS = *scalar-char-initialization-expr*
or BLANK = *scalar-default-char-expr*
or DECIMAL = *scalar-default-char-expr*
or DELIM = *scalar-default-char-expr*
or END = *label*
or EOR = *label*
or ERR = *label*
or ID = *scalar-int-variable*
or IOMSG = *iomsg-variable*
or IOSTAT = *scalar-int-variable*
or PAD = *scalar-default-char-expr*
or POS = *scalar-int-expr*
or REC = *scalar-int-expr*
or ROUND = *scalar-default-char-expr*
or SIGN = *scalar-default-char-expr*
or SIZE = *scalar-int-variable*
- C910 No specifier shall appear more than once in a given *io-control-spec-list*.
- C911 An *io-unit* shall be specified in an *io-control-spec-list*; if the optional characters UNIT= are

- omitted, the *io-unit* shall be the first item in the *io-control-spec-list*.
- C912 (R913) A DELIM= or SIGN= specifier shall not appear in a *read-stmt*.
- C913 (R913) A BLANK=, PAD=, END=, EOR=, or SIZE= specifier shall not appear in a *write-stmt*.
- C914 (R913) The *label* in the ERR=, EOR=, or END= specifier shall be the statement label of a branch target statement that appears in the same scoping unit as the data transfer statement.
- C915 (R913) A *namelist-group-name* shall be the name of a namelist group.
- C916 (R913) A *namelist-group-name* shall not appear if an *input-item-list* or an *output-item-list* appears in the data transfer statement.
- C917 (R913) An *io-control-spec-list* shall not contain both a *format* and a *namelist-group-name*.
- C918 (R913) If *format* appears without a preceding FMT=, it shall be the second item in the *io-control-spec-list* and the first item shall be *io-unit*.
- C919 (R913) If *namelist-group-name* appears without a preceding NML=, it shall be the second item in the *io-control-spec-list* and the first item shall be *io-unit*.
- C920 (R913) If *io-unit* is not a *file-unit-number*, the *io-control-spec-list* shall not contain a REC= specifier or a POS= specifier.
- C921 (R913) If the REC= specifier appears, an END= specifier shall not appear, a *namelist-group-name* shall not appear, and the *format*, if any, shall not be an asterisk.
- C922 (R913) An ADVANCE= specifier may appear only in a formatted sequential or stream input/output statement with explicit format specification (10.2) whose control information list does not contain an *internal-file-variable* as the *io-unit*.
- C923 (R913) If an EOR= specifier appears, an ADVANCE= specifier also shall appear.
- C924 (R913) If a SIZE= specifier appears, an ADVANCE= specifier also shall appear.
- C925 (R913) The *scalar-char-initialization-expr* in an ASYNCHRONOUS= specifier shall be of type default character and shall have the value YES or NO.
- C926 (R913) An ASYNCHRONOUS= specifier with a value YES shall not appear unless *io-unit* is a *file-unit-number*.
- C927 (R913) If an ID= specifier appears, an ASYNCHRONOUS= specifier with the value YES shall also appear.
- C928 (R913) If a POS= specifier appears, the *io-control-spec-list* shall not contain a REC= specifier.
- C929 (R913) If a DECIMAL=, BLANK=, PAD=, SIGN=, or ROUND= specifier appears, a *format* or *namelist-group-name* shall also appear.
- C930 (R913) If a DELIM= specifier appears, either *format* shall be an asterisk or *namelist-group-name* shall appear.
- R914 *format* **is** *default-char-expr*
 or *label*
 or *
- C931 (R914) The *label* shall be the label of a FORMAT statement that appears in the same scoping unit as the statement containing the FMT= specifier.
- R915 *input-item* **is** *variable*
 or *io-implied-do*
- R916 *output-item* **is** *expr*
 or *io-implied-do*
- R917 *io-implied-do* **is** (*io-implied-do-object-list* , *io-implied-do-control*)
- R918 *io-implied-do-object* **is** *input-item*
 or *output-item*
- R919 *io-implied-do-control* **is** *do-variable = scalar-int-expr* , ■

■ *scalar-int-expr* [, *scalar-int-expr*]

- C932 (R915) A variable that is an *input-item* shall not be a whole assumed-size array.
- C933 (R915) A variable that is an *input-item* shall not be a procedure pointer.
- C934 (R919) The *do-variable* shall be a named scalar variable of type integer.
- C935 (R918) In an *input-item-list*, an *io-implied-do-object* shall be an *input-item*. In an *output-item-list*, an *io-implied-do-object* shall be an *output-item*.
- C936 (R916) An expression that is an *output-item* shall not have a value that is a procedure pointer.
- R920 *dtv-type-spec* **is** TYPE(*derived-type-spec*)
 or CLASS(*derived-type-spec*)
- C937 (R920) If *derived-type-spec* specifies an extensible type, the CLASS keyword shall be used; otherwise, the TYPE keyword shall be used.
- C938 (R920) All length type parameters of *derived-type-spec* shall be assumed.
- R921 *wait-stmt* **is** WAIT (*wait-spec-list*)
- R922 *wait-spec* **is** [UNIT =] *file-unit-number*
 or END = *label*
 or EOR = *label*
 or ERR = *label*
 or ID = *scalar-int-expr*
 or IOMSG = *iomsg-variable*
 or IOSTAT = *scalar-int-variable*
- C939 No specifier shall appear more than once in a given *wait-spec-list*.
- C940 A *file-unit-number* shall be specified in a *wait-spec-list*; if the optional characters UNIT= are omitted, the *file-unit-number* shall be the first item in the *wait-spec-list*.
- C941 (R922) The *label* in the ERR=, EOR=, or END= specifier shall be the statement label of a branch target statement that appears in the same scoping unit as the WAIT statement.
- R923 *backspace-stmt* **is** BACKSPACE *file-unit-number*
 or BACKSPACE (*position-spec-list*)
- R924 *endfile-stmt* **is** ENDFILE *file-unit-number*
 or ENDFILE (*position-spec-list*)
- R925 *rewind-stmt* **is** REWIND *file-unit-number*
 or REWIND (*position-spec-list*)
- R926 *position-spec* **is** [UNIT =] *file-unit-number*
 or IOMSG = *iomsg-variable*
 or IOSTAT = *scalar-int-variable*
 or ERR = *label*
- C942 No specifier shall appear more than once in a given *position-spec-list*.
- C943 A *file-unit-number* shall be specified in a *position-spec-list*; if the optional characters UNIT= are omitted, the *file-unit-number* shall be the first item in the *position-spec-list*.
- C944 (R926) The *label* in the ERR= specifier shall be the statement label of a branch target statement that appears in the same scoping unit as the file positioning statement.
- R927 *flush-stmt* **is** FLUSH *file-unit-number*
 or FLUSH (*flush-spec-list*)
- R928 *flush-spec* **is** [UNIT =] *file-unit-number*
 or IOSTAT = *scalar-int-variable*
 or IOMSG = *iomsg-variable*

- or** ERR = *label*
- C945 No specifier shall appear more than once in a given *flush-spec-list*.
- C946 A *file-unit-number* shall be specified in a *flush-spec-list*; if the optional characters UNIT= are omitted from the unit specifier, the *file-unit-number* shall be the first item in the *flush-spec-list*.
- C947 (R928) The *label* in the ERR= specifier shall be the statement label of a branch target statement that appears in the same scoping unit as the FLUSH statement.
- R929 *inquire-stmt* **is** INQUIRE (*inquire-spec-list*)
or INQUIRE (IOLENGTH = *scalar-int-variable*) ■
 ■ *output-item-list*
- R930 *inquire-spec* **is** [UNIT =] *file-unit-number*
or FILE = *file-name-expr*
or ACCESS = *scalar-default-char-variable*
or ACTION = *scalar-default-char-variable*
or ASYNCHRONOUS = *scalar-default-char-variable*
or BLANK = *scalar-default-char-variable*
or DECIMAL = *scalar-default-char-variable*
or DELIM = *scalar-default-char-variable*
or DIRECT = *scalar-default-char-variable*
or ENCODING = *scalar-default-char-variable*
or ERR = *label*
or EXIST = *scalar-default-logical-variable*
or FORM = *scalar-default-char-variable*
or FORMATTED = *scalar-default-char-variable*
or ID = *scalar-int-expr*
or IOMSG = *iomsg-variable*
or IOSTAT = *scalar-int-variable*
or NAME = *scalar-default-char-variable*
or NAMED = *scalar-default-logical-variable*
or NEXTREC = *scalar-int-variable*
or NUMBER = *scalar-int-variable*
or OPENED = *scalar-default-logical-variable*
or PAD = *scalar-default-char-variable*
or PENDING = *scalar-default-logical-variable*
or POS = *scalar-int-variable*
or POSITION = *scalar-default-char-variable*
or READ = *scalar-default-char-variable*
or READWRITE = *scalar-default-char-variable*
or RECL = *scalar-int-variable*
or ROUND = *scalar-default-char-variable*
or SEQUENTIAL = *scalar-default-char-variable*
or SIGN = *scalar-default-char-variable*
or SIZE = *scalar-int-variable*
or STREAM = *scalar-default-char-variable*
or TEAM = *image-team*
or UNFORMATTED = *scalar-default-char-variable*

or WRITE = *scalar-default-char-variable*

- C948 No specifier shall appear more than once in a given *inquire-spec-list*.
 C949 An *inquire-spec-list* shall contain one FILE= specifier or one UNIT= specifier, but not both.
 C950 In the inquire by unit form of the INQUIRE statement, if the optional characters UNIT= are omitted, the *file-unit-number* shall be the first item in the *inquire-spec-list*.
 C951 If an ID= specifier appears in an *inquire-spec-list*, a PENDING= specifier shall also appear.
 C952 (R928) The *label* in the ERR= specifier shall be the statement label of a branch target statement that appears in the same scoping unit as the INQUIRE statement.

Clause 10:

- R1001 *format-stmt* is FORMAT *format-specification*
 R1002 *format-specification* is ([*format-item-list*])
 or ([*format-item-list*,] *unlimited-format-item*)
 C1001 (R1001) The *format-stmt* shall be labeled.
 C1002 (R1002) The comma used to separate *format-items* in a *format-item-list* may be omitted
 R1003 *format-item* is [*r*] *data-edit-desc*
 or *control-edit-desc*
 or *char-string-edit-desc*
 or [*r*] (*format-item-list*)
 R1004 *unlimited-format-item* is * (*format-item-list*)
 R1005 *r* is *int-literal-constant*
 C1003 (R1005) *r* shall be positive.
 C1004 (R1005) *r* shall not have a kind parameter specified for it.
 R1006 *data-edit-desc* is I *w* [. *m*]
 or B *w* [. *m*]
 or O *w* [. *m*]
 or Z *w* [. *m*]
 or F *w* . *d*
 or E *w* . *d* [E *e*]
 or EN *w* . *d* [E *e*]
 or ES *w* . *d* [E *e*]
 or G *w* [. *d* [E *e*]]
 or L *w*
 or A [*w*]
 or D *w* . *d*
 or DT [*char-literal-constant*] [(*v-list*)]
 R1007 *w* is *int-literal-constant*
 R1008 *m* is *int-literal-constant*
 R1009 *d* is *int-literal-constant*
 R1010 *e* is *int-literal-constant*
 R1011 *v* is *signed-int-literal-constant*
 C1005 (R1010) *e* shall be positive.
 C1006 (R1007) *w* shall be zero or positive for the I, B, O, Z, F, and G edit descriptors. *w* shall be positive for all other edit descriptors.
 C1007 (R1006) For the G edit descriptor, *d* shall be specified if and only if *w* is not zero.
 C1008 (R1006) *w*, *m*, *d*, *e*, and *v* shall not have kind parameters specified for them.
 C1009 (R1006) The *char-literal-constant* in the DT edit descriptor shall not have a kind parameter

specified for it.

- R1012 *control-edit-desc* **is** *position-edit-desc*
 or [*r*] /
 or :
 or *sign-edit-desc*
 or *k P*
 or *blank-interp-edit-desc*
 or *round-edit-desc*
 or *decimal-edit-desc*
- R1013 *k* **is** *signed-int-literal-constant*
- C1010 (R1013) *k* shall not have a kind parameter specified for it.
- R1014 *position-edit-desc* **is** *T n*
 or *TL n*
 or *TR n*
 or *n X*
- R1015 *n* **is** *int-literal-constant*
- C1011 (R1015) *n* shall be positive.
- C1012 (R1015) *n* shall not have a kind parameter specified for it.
- R1016 *sign-edit-desc* **is** *SS*
 or *SP*
 or *S*
- R1017 *blank-interp-edit-desc* **is** *BN*
 or *BZ*
- R1018 *round-edit-desc* **is** *RU*
 or *RD*
 or *RZ*
 or *RN*
 or *RC*
 or *RP*
- R1019 *decimal-edit-desc* **is** *DC*
 or *DP*
- R1020 *char-string-edit-desc* **is** *char-literal-constant*
- C1013 (R1020) The *char-literal-constant* shall not have a kind parameter specified for it.
- R1021 *hex-digit-string* **is** *hex-digit* [*hex-digit*] ...

Clause 11:

- R1101 *main-program* **is** [*program-stmt*]
 [*specification-part*]
 [*execution-part*]
 [*internal-subprogram-part*]
 end-program-stmt
- R1102 *program-stmt* **is** *PROGRAM program-name*
- R1103 *end-program-stmt* **is** *END* [*PROGRAM* [*program-name*]]
- C1101 (R1101) In a *main-program*, the *execution-part* shall not contain a RETURN statement or an ENTRY statement.
- C1102 (R1101) The *program-name* may be included in the *end-program-stmt* only if the optional *program-stmt* is used and, if included, shall be identical to the *program-name* specified in the

- program-stmt.*
- R1104 *module* **is** *module-stmt*
 [*specification-part*]
 [*module-subprogram-part*]
 end-module-stmt
- R1105 *module-stmt* **is** MODULE *module-name*
- R1106 *end-module-stmt* **is** END [MODULE [*module-name*]]
- R1107 *module-subprogram-part* **is** *contains-stmt*
 [*module-subprogram*] ...
- R1108 *module-subprogram* **is** *function-subprogram*
 or *subroutine-subprogram*
 or *separate-module-subprogram*
- C1103 (R1104) If the *module-name* is specified in the *end-module-stmt*, it shall be identical to the *module-name* specified in the *module-stmt*.
- C1104 (R1104) A *module specification-part* shall not contain a *stmt-function-stmt*, an *entry-stmt*, or a *format-stmt*.
- C1105 (R1104) If an object that has a default-initialized direct component is declared in the *specification-part* of a module it shall have the ALLOCATABLE, POINTER, or SAVE attribute.
- C1106 If an object that has a co-array ultimate component is declared in the *specification-part* of a module it shall have the SAVE attribute.
- R1109 *use-stmt* **is** USE [[, *module-nature*] ::] *module-name* [, *rename-list*]
 or USE [[, *module-nature*] ::] *module-name* , ■
 ■ ONLY : [*only-list*]
- R1110 *module-nature* **is** INTRINSIC
 or NON_INTRINSIC
- R1111 *rename* **is** *local-name* => *use-name*
 or OPERATOR (*local-defined-operator*) => ■
 ■ OPERATOR (*use-defined-operator*)
- R1112 *only* **is** *generic-spec*
 or *only-use-name*
 or *rename*
- R1113 *only-use-name* **is** *use-name*
- C1107 (R1109) If *module-nature* is INTRINSIC, *module-name* shall be the name of an intrinsic module.
- C1108 (R1109) If *module-nature* is NON_INTRINSIC, *module-name* shall be the name of a nonintrinsic module.
- C1109 (R1109) A scoping unit shall not access an intrinsic module and a nonintrinsic module of the same name.
- C1110 (R1111) OPERATOR(*use-defined-operator*) shall not identify a *generic-binding*.
- C1111 (R1112) The *generic-spec* shall not identify a *generic-binding*.
- C1112 (R1112) Each *generic-spec* shall be a public entity in the module.
- C1113 (R1113) Each *use-name* shall be the name of a public entity in the module.
- R1114 *local-defined-operator* **is** *defined-unary-op*
 or *defined-binary-op*
- R1115 *use-defined-operator* **is** *defined-unary-op*
 or *defined-binary-op*
- C1114 (R1115) Each *use-defined-operator* shall be a public entity in the module.
- R1116 *submodule* **is** *submodule-stmt*

- [*specification-part*]
[*module-subprogram-part*]
end-submodule-stmt
- R1117 *submodule-stmt* **is** SUBMODULE (*parent-identifier*) *submodule-name*
R1118 *parent-identifier* **is** *ancestor-module-name* [: *parent-submodule-name*]
R1119 *end-submodule-stmt* **is** END [SUBMODULE [*submodule-name*]]
C1115 (R1116) A submodule *specification-part* shall not contain a *format-stmt*, *entry-stmt*, or *stmt-function-stmt*.
C1116 (R1116) An object with a default-initialized direct component that is declared in the *specification-part* of a submodule shall have the ALLOCATABLE, POINTER, or SAVE attribute.
C1117 (R1118) The *ancestor-module-name* shall be the name of a nonintrinsic module; the *parent-submodule-name* shall be the name of a descendant of that module.
C1118 (R1116) If a *submodule-name* appears in the *end-submodule-stmt*, it shall be identical to the one in the *submodule-stmt*.
R1120 *block-data* **is** *block-data-stmt*
 [*specification-part*]
 end-block-data-stmt
R1121 *block-data-stmt* **is** BLOCK DATA [*block-data-name*]
R1122 *end-block-data-stmt* **is** END [BLOCK DATA [*block-data-name*]]
C1119 (R1120) The *block-data-name* shall be included in the *end-block-data-stmt* only if it was provided in the *block-data-stmt* and, if included, shall be identical to the *block-data-name* in the *block-data-stmt*.
C1120 (R1120) A *block-data specification-part* shall contain only derived-type definitions and ASYNCHRONOUS, BIND, COMMON, DATA, DIMENSION, EQUIVALENCE, IMPLICIT, INTRINSIC, PARAMETER, POINTER, SAVE, TARGET, USE, VOLATILE, and type declaration statements.
C1121 (R1120) A type declaration statement in a *block-data specification-part* shall not contain ALLOCATABLE, EXTERNAL, or BIND attribute specifiers.

Clause 12:

- R1201 *interface-block* **is** *interface-stmt*
 [*interface-specification*] ...
 end-interface-stmt
R1202 *interface-specification* **is** *interface-body*
 or *procedure-stmt*
R1203 *interface-stmt* **is** INTERFACE [*generic-spec*]
 or ABSTRACT INTERFACE
R1204 *end-interface-stmt* **is** END INTERFACE [*generic-spec*]
R1205 *interface-body* **is** *function-stmt*
 [*specification-part*]
 end-function-stmt
 or *subroutine-stmt*
 [*specification-part*]
 end-subroutine-stmt
R1206 *procedure-stmt* **is** [MODULE] PROCEDURE *procedure-name-list*
R1207 *generic-spec* **is** *generic-name*
 or OPERATOR (*defined-operator*)
 or ASSIGNMENT (=)

- R1208 *dtio-generic-spec* **or** *dtio-generic-spec*
is READ (FORMATTED)
or READ (UNFORMATTED)
or WRITE (FORMATTED)
or WRITE (UNFORMATTED)
- R1209 *import-stmt* **is** IMPORT [[::] *import-name-list*
- C1201 (R1201) An *interface-block* in a subprogram shall not contain an *interface-body* for a procedure defined by that subprogram.
- C1202 (R1201) The *generic-spec* shall be included in the *end-interface-stmt* only if it is provided in the *interface-stmt*. If the *end-interface-stmt* includes *generic-name*, the *interface-stmt* shall specify the same *generic-name*. If the *end-interface-stmt* includes ASSIGNMENT(=), the *interface-stmt* shall specify ASSIGNMENT(=). If the *end-interface-stmt* includes *dtio-generic-spec*, the *interface-stmt* shall specify the same *dtio-generic-spec*. If the *end-interface-stmt* includes OPERATOR(*defined-operator*), the *interface-stmt* shall specify the same *defined-operator*. If one *defined-operator* is .LT., .LE., .GT., .GE., .EQ., or .NE., the other is permitted to be the corresponding operator <, <=, >, >=, ==, or /=.
- C1203 (R1203) If the *interface-stmt* is ABSTRACT INTERFACE, then the *function-name* in the *function-stmt* or the *subroutine-name* in the *subroutine-stmt* shall not be the same as a keyword that specifies an intrinsic type.
- C1204 (R1202) A *procedure-stmt* is allowed only in an interface block that has a *generic-spec*.
- C1205 (R1205) An *interface-body* of a pure procedure shall specify the intents of all dummy arguments except pointer, alternate return, and procedure arguments.
- C1206 (R1205) An *interface-body* shall not contain an *entry-stmt*, *data-stmt*, *format-stmt*, or *stmt-function-stmt*.
- C1207 (R1206) A *procedure-name* shall have an explicit interface and shall refer to an accessible procedure pointer, external procedure, dummy procedure, or module procedure.
- C1208 (R1206) If MODULE appears in a *procedure-stmt*, each *procedure-name* in that statement shall be accessible in the current scope as a module procedure.
- C1209 (R1206) A *procedure-name* shall not specify a procedure that is specified previously in any *procedure-stmt* in any accessible interface with the same generic identifier.
- C1210 (R1209) The IMPORT statement is allowed only in an *interface-body* that is not a module procedure interface body.
- C1211 (R1209) Each *import-name* shall be the name of an entity in the host scoping unit.
- C1212 (R1205) A module procedure interface body shall not appear in an abstract interface block.
- R1210 *external-stmt* **is** EXTERNAL [::] *external-name-list*
- R1211 *procedure-declaration-stmt* **is** PROCEDURE ([*proc-interface*]) ■
 ■ [[, *proc-attr-spec*] ... ::] *proc-decl-list*
- R1212 *proc-interface* **is** *interface-name*
or *declaration-type-spec*
- R1213 *proc-attr-spec* **is** *access-spec*
or *proc-language-binding-spec*
or INTENT (*intent-spec*)
or OPTIONAL
or POINTER
or SAVE
- R1214 *proc-decl* **is** *procedure-entity-name*[=> *proc-pointer-init*]
- R1215 *interface-name* **is** *name*
- R1216 *proc-pointer-init* **is** *null-init*
or *initial-proc-target*

- R1217 *initial-proc-target* **is** *procedure-name*
- C1213 (R1215) The *name* shall be the name of an abstract interface or of a procedure that has an explicit interface. If *name* is declared by a *procedure-declaration-stmt* it shall be previously declared. If *name* denotes an intrinsic procedure it shall be one that is listed in 13.6 and not marked with a bullet (●).
- C1214 (R1215) The *name* shall not be the same as a keyword that specifies an intrinsic type.
- C1215 If a procedure entity has the INTENT attribute or SAVE attribute, it shall also have the POINTER attribute.
- C1216 (R1211) If a *proc-interface* describes an elemental procedure, each *procedure-entity-name* shall specify an external procedure.
- C1217 (R1214) If => appears in *proc-decl*, the procedure entity shall have the POINTER attribute.
- C1218 (R1217) The *procedure-name* shall be the name of an initialization target.
- C1219 (R1211) If *proc-language-binding-spec* with a NAME= is specified, then *proc-decl-list* shall contain exactly one *proc-decl*, which shall neither have the POINTER attribute nor be a dummy procedure.
- C1220 (R1211) If *proc-language-binding-spec* is specified, the *proc-interface* shall appear, it shall be an *interface-name*, and *interface-name* shall be declared with a *proc-language-binding-spec*.
- R1218 *intrinsic-stmt* **is** INTRINSIC [::] *intrinsic-procedure-name-list*
- C1221 (R1218) Each *intrinsic-procedure-name* shall be the name of an intrinsic procedure.
- R1219 *function-reference* **is** *procedure-designator* ([*actual-arg-spec-list*])
- C1222 (R1219) The *procedure-designator* shall designate a function.
- C1223 (R1219) The *actual-arg-spec-list* shall not contain an *alt-return-spec*.
- R1220 *call-stmt* **is** CALL *procedure-designator* [([*actual-arg-spec-list*])]
- C1224 (R1220) The *procedure-designator* shall designate a subroutine.
- R1221 *procedure-designator* **is** *procedure-name*
 or *proc-component-ref*
 or *data-ref* % *binding-name*
- C1225 (R1221) A *procedure-name* shall be the name of a procedure or procedure pointer.
- C1226 (R1221) A *binding-name* shall be a binding name (4.5.5) of the declared type of *data-ref*.
- C1227 (R1221) If *data-ref* is an array, the referenced type-bound procedure shall have the PASS attribute.
- R1222 *actual-arg-spec* **is** [*keyword* =] *actual-arg*
- R1223 *actual-arg* **is** *expr*
 or *variable*
 or *procedure-name*
 or *proc-component-ref*
 or *alt-return-spec*
- R1224 *alt-return-spec* **is** * *label*
- C1228 (R1222) The *keyword* = shall not appear if the interface of the procedure is implicit in the scoping unit.
- C1229 (R1222) The *keyword* = shall not be omitted from an *actual-arg-spec* unless it has been omitted from each preceding *actual-arg-spec* in the argument list.
- C1230 (R1222) Each *keyword* shall be the name of a dummy argument in the explicit interface of the procedure.
- C1231 (R1223) A nonintrinsic elemental procedure shall not be used as an actual argument.
- C1232 (R1223) A *procedure-name* shall be the name of an external, internal, module, or dummy procedure, a specific intrinsic function listed in 13.6 and not marked with a bullet (●), or a procedure

- pointer.
- C1233 (R1224) The *label* shall be the statement label of a branch target statement that appears in the same scoping unit as the *call-stmt*.
- C1234 An actual argument that is a co-indexed object shall not be associated with a dummy argument that has the ASYNCHRONOUS attribute.
- C1235 (R1223) If an actual argument is an array section or an assumed-shape array, and the corresponding dummy argument has either the VOLATILE or ASYNCHRONOUS attribute, that dummy argument shall be an assumed-shape array.
- C1236 (R1223) If an actual argument is a pointer array, and the corresponding dummy argument has either the VOLATILE or ASYNCHRONOUS attribute, that dummy argument shall be an assumed-shape array that does not have the CONTIGUOUS attribute or a pointer array.
- R1225 *function-subprogram* **is** *function-stmt*
 [*specification-part*]
 [*execution-part*]
 [*internal-subprogram-part*]
 end-function-stmt
- R1226 *function-stmt* **is** [*prefix*] FUNCTION *function-name* ■
 ■ ([*dummy-arg-name-list*]) [*suffix*]
- C1237 (R1226) If RESULT appears, *result-name* shall not be the same as *function-name* and shall not be the same as the *entry-name* in any ENTRY statement in the subprogram.
- C1238 (R1226) If RESULT appears, the *function-name* shall not appear in any specification statement in the scoping unit of the function subprogram.
- R1227 *proc-language-binding-spec* **is** *language-binding-spec*
- C1239 (R1227) A *proc-language-binding-spec* with a NAME= specifier shall not be specified in the *function-stmt* or *subroutine-stmt* of an interface body for an abstract interface or a dummy procedure.
- C1240 (R1227) A *proc-language-binding-spec* shall not be specified for an internal procedure.
- C1241 (R1227) If *proc-language-binding-spec* is specified for a procedure, each of the procedure's dummy arguments shall be a nonoptional interoperable variable (15.3.5, 15.3.6) or a nonoptional interoperable procedure (15.3.7). If *proc-language-binding-spec* is specified for a function, the function result shall be an interoperable scalar variable.
- R1228 *dummy-arg-name* **is** *name*
- C1242 (R1228) A *dummy-arg-name* shall be the name of a dummy argument.
- R1229 *prefix* **is** *prefix-spec* [*prefix-spec*] ...
- R1230 *prefix-spec* **is** *declaration-type-spec*
 or ELEMENTAL
 or IMPURE
 or MODULE
 or PURE
 or RECURSIVE
- C1243 (R1229) A *prefix* shall contain at most one of each *prefix-spec*.
- C1244 (R1229) A *prefix* shall not specify both PURE and IMPURE.
- C1245 (R1229) A *prefix* shall not specify both ELEMENTAL and RECURSIVE.
- C1246 (R1229) A *prefix* shall not specify ELEMENTAL if *proc-language-binding-spec* appears in the *function-stmt* or *subroutine-stmt*.
- C1247 (R1229) MODULE shall appear only within the *function-stmt* or *subroutine-stmt* of a module subprogram or of an interface body that is declared in the scoping unit of a module or submodule.
- C1248 (R1229) If MODULE appears within the *prefix* in a module subprogram, an accessible module procedure interface having the same name as the subprogram shall be declared in the module

or submodule in which the subprogram is defined, or shall be declared in an ancestor of that program unit.

C1249 (R1229) If MODULE appears within the *prefix* in a module subprogram, the subprogram shall specify the same characteristics and dummy argument names as its corresponding (12.6.2.4) module procedure interface body.

C1250 (R1229) If MODULE appears within the *prefix* in a module subprogram and a binding label is specified, it shall be the same as the binding label specified in the corresponding module procedure interface body.

C1251 (R1229) If MODULE appears within the *prefix* in a module subprogram, RECURSIVE shall appear if and only if RECURSIVE appears in the *prefix* in the corresponding module procedure interface body.

R1231 *suffix* **is** *proc-language-binding-spec* [RESULT (*result-name*)]
or RESULT (*result-name*) [*proc-language-binding-spec*]

R1232 *end-function-stmt* **is** END [FUNCTION [*function-name*]]

C1252 (R1225) An internal function subprogram shall not contain an ENTRY statement.

C1253 (R1225) An internal function subprogram shall not contain an *internal-subprogram-part*.

C1254 (R1232) If a *function-name* appears in the *end-function-stmt*, it shall be identical to the *function-name* specified in the *function-stmt*.

R1233 *subroutine-subprogram* **is** *subroutine-stmt*
[*specification-part*]
[*execution-part*]
[*internal-subprogram-part*]
end-subroutine-stmt

R1234 *subroutine-stmt* **is** [*prefix*] SUBROUTINE *subroutine-name* ■
■ [([*dummy-arg-list*]) [*proc-language-binding-spec*]]

C1255 (R1234) The *prefix* of a *subroutine-stmt* shall not contain a *declaration-type-spec*.

R1235 *dummy-arg* **is** *dummy-arg-name*
or *

R1236 *end-subroutine-stmt* **is** END [SUBROUTINE [*subroutine-name*]]

C1256 (R1233) An internal subroutine subprogram shall not contain an ENTRY statement.

C1257 (R1233) An internal subroutine subprogram shall not contain an *internal-subprogram-part*.

C1258 (R1236) If a *subroutine-name* appears in the *end-subroutine-stmt*, it shall be identical to the *subroutine-name* specified in the *subroutine-stmt*.

R1237 *separate-module-subprogram* **is** *mp-subprogram-stmt*
[*specification-part*]
[*execution-part*]
[*internal-subprogram-part*]
end-mp-subprogram-stmt

R1238 *mp-subprogram-stmt* **is** MODULE PROCEDURE *procedure-name*

R1239 *end-mp-subprogram-stmt* **is** END [PROCEDURE [*procedure-name*]]

C1259 (R1237) The *procedure-name* shall be the same as the name of an accessible module procedure interface that is declared in the module or submodule in which the *separate-module-subprogram* is defined, or is declared in an ancestor of that program unit.

C1260 (R1239) If a *procedure-name* appears in the *end-mp-subprogram-stmt*, it shall be identical to the *procedure-name* in the MODULE PROCEDURE statement.

R1240 *entry-stmt* **is** ENTRY *entry-name* [([*dummy-arg-list*]) [*suffix*]]

C1261 (R1240) If RESULT appears, the *entry-name* shall not appear in any specification or type-

declaration statement in the scoping unit of the function program.

- C1262 (R1240) An *entry-stmt* shall appear only in an *external-subprogram* or a *module-subprogram* that does not define a separate module procedure. An *entry-stmt* shall not appear within an *executable-construct*.
- C1263 (R1240) RESULT shall appear only if the *entry-stmt* is in a function subprogram.
- C1264 (R1240) Within the subprogram containing the *entry-stmt*, the *entry-name* shall not appear as a dummy argument in the FUNCTION or SUBROUTINE statement or in another ENTRY statement nor shall it appear in an EXTERNAL, INTRINSIC, or PROCEDURE statement.
- C1265 (R1240) A *dummy-arg* shall not be an alternate return indicator if the ENTRY statement is in a function subprogram.
- C1266 (R1240) If RESULT appears, *result-name* shall not be the same as the *function-name* in the FUNCTION statement and shall not be the same as the *entry-name* in any ENTRY statement in the subprogram.
- R1241 *return-stmt* **is** RETURN [*scalar-int-expr*]
- C1267 (R1241) The *return-stmt* shall be in the scoping unit of a function or subroutine subprogram.
- C1268 (R1241) The *scalar-int-expr* is allowed only in the scoping unit of a subroutine subprogram.
- R1242 *contains-stmt* **is** CONTAINS
- R1243 *stmt-function-stmt* **is** *function-name* ([*dummy-arg-name-list*]) = *scalar-expr*
- C1269 (R1243) The *primaries* of the *scalar-expr* shall be constants (literal and named), references to variables, references to functions and function dummy procedures, and intrinsic operations. If *scalar-expr* contains a reference to a function or a function dummy procedure, the reference shall not require an explicit interface, the function shall not require an explicit interface unless it is an intrinsic function, the function shall not be a transformational intrinsic, and the result shall be scalar. If an argument to a function or a function dummy procedure is an array, it shall be an array name. If a reference to a statement function appears in *scalar-expr*, its definition shall have been provided earlier in the scoping unit and shall not be the name of the statement function being defined.
- C1270 (R1243) Named constants in *scalar-expr* shall have been declared earlier in the scoping unit or made accessible by use or host association. If array elements appear in *scalar-expr*, the array shall have been declared as an array earlier in the scoping unit or made accessible by use or host association.
- C1271 (R1243) If a *dummy-arg-name*, variable, function reference, or dummy function reference is typed by the implicit typing rules, its appearance in any subsequent type declaration statement shall confirm this implied type and the values of any implied type parameters.
- C1272 (R1243) The *function-name* and each *dummy-arg-name* shall be specified, explicitly or implicitly, to be scalar.
- C1273 (R1243) A given *dummy-arg-name* shall not appear more than once in any *dummy-arg-name-list*.
- C1274 (R1243) Each variable reference in *scalar-expr* may be either a reference to a dummy argument of the statement function or a reference to a variable accessible in the same scoping unit as the statement function statement.
- C1275 The *specification-part* of a pure function subprogram shall specify that all its nonpointer dummy data objects have INTENT(IN).
- C1276 The *specification-part* of a pure subroutine subprogram shall specify the intents of all its non-pointer dummy data objects.
- C1277 A local variable declared in the *specification-part* of a pure subprogram, or within the *specification-part* of a BLOCK construct within a pure subprogram, shall not have the SAVE attribute.
- C1278 The *specification-part* of a pure subprogram shall specify that all its dummy procedures are pure.
- C1279 If a procedure that is neither an intrinsic procedure nor a statement function is used in a context that requires it to be pure, then its interface shall be explicit in the scope of that use. The interface shall specify that the procedure is pure.
- C1280 All internal subprograms in a pure subprogram shall be pure.
- C1281 In a pure subprogram any designator with a base object that is in common or accessed by host or use association, is a dummy argument of a pure function, is a dummy argument with INTENT (IN) of a pure subroutine, or an object that is storage associated with any such variable,

shall not be used

- C1282 Any procedure referenced in a pure subprogram, including one referenced via a defined operation, defined assignment, user-defined derived-type input/output, or finalization, shall be pure.
- C1283 A pure subprogram shall not contain a *print-stmt*, *open-stmt*, *close-stmt*, *backspace-stmt*, *endfile-stmt*, *rewind-stmt*, *flush-stmt*, *wait-stmt*, or *inquire-stmt*.
- C1284 A pure subprogram shall not contain a *read-stmt* or *write-stmt* whose *io-unit* is a *file-unit-number* or ***.
- C1285 A pure subprogram shall not contain a *stop-stmt*.
- C1286 A co-indexed object shall not appear in a variable definition context in a pure subprogram.
- C1287 A pure subprogram shall not contain an image control statement (8.5.1).
- C1288 All dummy arguments of an elemental procedure shall be scalar dummy data objects and shall not have the POINTER or ALLOCATABLE attribute.
- C1289 The result variable of an elemental function shall be scalar and shall not have the POINTER or ALLOCATABLE attribute.
- C1290 In the scoping unit of an elemental subprogram, an object designator with a dummy argument as the base object shall not appear in a *specification-expr* except as the argument to one of the intrinsic functions BIT_SIZE, KIND, LEN, or the numeric inquiry functions (13.5.6).

Clause 13:

Clause 14:

Clause 15:

- C1501 (R430) A derived type with the BIND attribute shall not be a SEQUENCE type.
- C1502 (R430) A derived type with the BIND attribute shall not have type parameters.
- C1503 (R430) A derived type with the BIND attribute shall not have the EXTENDS attribute.
- C1504 (R430) A derived type with the BIND attribute shall not have a *type-bound-procedure-part*.
- C1505 (R430) Each component of a derived type with the BIND attribute shall be a nonpointer, nonallocatable data component with interoperable type and type parameters.
- C1506 A procedure defined in a submodule shall not have a binding label unless its interface is declared in the ancestor module.

Clause 16:

D.2 Syntax rule cross-reference

R474	<i>ac-do-variable</i>	R473, C510
R472	<i>ac-implied-do</i>	R471, C510
R473	<i>ac-implied-do-control</i>	R472
R468	<i>ac-spec</i>	R467
R471	<i>ac-value</i>	R468, R472, C507, C508, C509
R525	<i>access-id</i>	R524, C562
R507	<i>access-spec</i>	R316, R432, R442, R446, R454, R455, R502, C517, R524, R1213
R524	<i>access-stmt</i>	R212, C562
R214	<i>action-stmt</i>	R213, C324, C325, R834, R838, R845, C829
R832	<i>action-term-do-construct</i>	R831
R1223	<i>actual-arg</i>	R1222
R1222	<i>actual-arg-spec</i>	C503, R1219, C1223, R1220, C1229
R709	<i>add-op</i>	R310, R706
R705	<i>add-operand</i>	R705, R706

R627	<i>alloc-opt</i>	R626, C635
R527	<i>allocatable-decl</i>	R526
R526	<i>allocatable-stmt</i>	R212
R636	<i>allocate-co-array-spec</i>	R631, C633, C634
R637	<i>allocate-co-shape-spec</i>	R636, C634
R632	<i>allocate-object</i>	R631, C627, C628, C629, C630, C631, C632, C633, C634, C637, C638, C641, R640, C643
R633	<i>allocate-shape-spec</i>	R631, C633, C634
R626	<i>allocate-stmt</i>	R214
R631	<i>allocation</i>	R626
R302	<i>alphanumeric-character</i>	R301, R304
R1224	<i>alt-return-spec</i>	C1223, R1223
—	<i>ancestor-module-name</i>	R1118, C1117
R719	<i>and-op</i>	R310, R715
R714	<i>and-operand</i>	R715
—	<i>arg-name</i>	R446, C460, R455, C481
R344	<i>arg-token</i>	R343
R853	<i>arithmetic-if-stmt</i>	R214, C325, C818, C820, C843
R467	<i>array-constructor</i>	C507, C508, C509, R701
R617	<i>array-element</i>	R536, C569, C572, R566, R603, R610
—	<i>array-name</i>	R543, R544
R618	<i>array-section</i>	R603, C623
R510	<i>array-spec</i>	R502, R503, C515, R509, C532, R526, R527, R543, R544, R555, R556, R568, C597
R734	<i>assignment-stmt</i>	R214, R747, R757, C745
R802	<i>associate-construct</i>	R213, C804
—	<i>associate-construct-name</i>	R803, R806, C804
—	<i>associate-name</i>	R804, C801, C802, R847, C830, C831
R803	<i>associate-stmt</i>	R802, C802, C804
R804	<i>association</i>	R803
R515	<i>assumed-shape-spec</i>	R510
R517	<i>assumed-size-spec</i>	R510
R528	<i>asynchronous-stmt</i>	R212
R502	<i>attr-spec</i>	R501, C501, C579
R923	<i>backspace-stmt</i>	R214, C1283
R342	<i>basic-token</i>	R341, R344
R341	<i>basic-token-sequence</i>	R337, R340, R341
R426	<i>binary-constant</i>	R425
R530	<i>bind-entity</i>	R529, C564
R529	<i>bind-stmt</i>	R212
R455	<i>binding-attr</i>	R453, C479, C482, C483
—	<i>binding-name</i>	R453, R454, C474, R1221, C1226
R451	<i>binding-private-stmt</i>	R450, C469
R1017	<i>blank-interp-edit-desc</i>	R1012
R801	<i>block</i>	R802, R807, R810, R818, C813, R828, R840, R846
R807	<i>block-construct</i>	R213, C806
—	<i>block-construct-name</i>	R808, R809, C806

R1120	<i>block-data</i>	R202, C1120, C1121
—	<i>block-data-name</i>	R1121, R1122, C1119
R1121	<i>block-data-stmt</i>	R1120, C1119
R822	<i>block-do-construct</i>	R821, C815
R808	<i>block-stmt</i>	R807, C806
R738	<i>bounds-remapping</i>	R735, C720, C721
R737	<i>bounds-spec</i>	R735, C719
R425	<i>boz-literal-constant</i>	R306, C414
R1220	<i>call-stmt</i>	R214, C1233
R810	<i>case-construct</i>	R213, C807, C809, C811
—	<i>case-construct-name</i>	R811, R812, R813, C807
R814	<i>case-expr</i>	R811, C809, C810, C811
R815	<i>case-selector</i>	R812
R812	<i>case-stmt</i>	R810, C807
R817	<i>case-value</i>	R816, C809
R816	<i>case-value-range</i>	R815, C810, C811
R309	<i>char-constant</i>	C303
R725	<i>char-expr</i>	C706, R731, R814
R731	<i>char-initialization-expr</i>	R508, C522, C713, R817, R856, C845, R913, C925
R422	<i>char-length</i>	R421, C420, C455, R503, C504
R423	<i>char-literal-constant</i>	R306, R1006, C1009, R1020, C1013
R420	<i>char-selector</i>	R404
R1020	<i>char-string-edit-desc</i>	R1003
R606	<i>char-variable</i>	C606, R903, C901, C902
R909	<i>close-spec</i>	R908, C907, C908
R908	<i>close-stmt</i>	R214, C1283
R511	<i>co-array-spec</i>	R442, C448, C449, R503, R509, C524, C537, C538, R527, R544, R556
—	<i>co-name</i>	R544
R625	<i>co-subscript</i>	R624
—	<i>common-block-name</i>	R530, R553, R567
R568	<i>common-block-object</i>	R567, C597, C598, C599, C600
R567	<i>common-stmt</i>	R212
R417	<i>complex-literal-constant</i>	R306
R615	<i>complex-part-designator</i>	R603, R618, C623
R444	<i>component-array-spec</i>	R442, C447, C451
R442	<i>component-attr-spec</i>	R441, C444, C465
R461	<i>component-data-source</i>	R460
R443	<i>component-decl</i>	R441, C463
R440	<i>component-def-stmt</i>	R439, C444, C445, C446, C456
R447	<i>component-initialization</i>	C463, C464, C465
—	<i>component-name</i>	C466
R439	<i>component-part</i>	R430
R460	<i>component-spec</i>	R459, C497, C498, C499, C500, C502, C503
R852	<i>computed-goto-stmt</i>	R214, C842
R711	<i>concat-op</i>	R310, R710
R905	<i>connect-spec</i>	R904, C903, C904

R305	<i>constant</i>	R308, R309, R540, R610, R701
R542	<i>constant-subobject</i>	R540, R541, C577
—	<i>construct-name</i>	R850, C839
R1242	<i>contains-stmt</i>	R210, R450, R1107
R854	<i>continue-stmt</i>	R214, C325, R829, C818
R1012	<i>control-edit-desc</i>	R1003
R818	<i>critical-construct</i>	R213, C812, C813
—	<i>critical-construct-name</i>	R819, R820, C812
R819	<i>critical-stmt</i>	R818, C812
R839	<i>cycle-stmt</i>	R214, C325, C818, C820, C823
R1009	<i>d</i>	R1006, C1007, C1008
R441	<i>data-component-def-stmt</i>	R440
R1006	<i>data-edit-desc</i>	R1003
R536	<i>data-i-do-object</i>	R535, C565, C567, C568, C572
R537	<i>data-i-do-variable</i>	R535, C572
R535	<i>data-implied-do</i>	R534, R536, C572
—	<i>data-pointer-component-name</i>	R736, C724
R736	<i>data-pointer-object</i>	R735, C717, C718, C719, C720, C721, C725
R612	<i>data-ref</i>	C612, C616, C617, R614, R617, R618, C723, C729, R1221, C1226, C1227
R532	<i>data-stmt</i>	R209, R212, C1206
R540	<i>data-stmt-constant</i>	C414, R538, C575
R534	<i>data-stmt-object</i>	R533, C565, C566, C567, C568
R539	<i>data-stmt-repeat</i>	R538, C573
R533	<i>data-stmt-set</i>	R532
R538	<i>data-stmt-value</i>	R533
R739	<i>data-target</i>	R461, C504, C505, C552, R735, C717, C718, C721, C727
R641	<i>dealloc-opt</i>	R640, C644
R640	<i>deallocate-stmt</i>	R214
R1019	<i>decimal-edit-desc</i>	R1012
R207	<i>declaration-construct</i>	R204
R403	<i>declaration-type-spec</i>	C404, C405, C424, R441, C445, C446, R501, R560, R1212, R1230, C1255
R726	<i>default-char-expr</i>	C707, R905, R906, R909, R913, R914
R607	<i>default-char-variable</i>	C607, R629, R907, R930
R605	<i>default-logical-variable</i>	C605, R930
R519	<i>deferred-co-shape-spec</i>	C448, C524, R511, C537
R516	<i>deferred-shape-spec</i>	R442, R444, C447, R510, C532, R550
R315	<i>define-macro-stmt</i>	R314
R723	<i>defined-binary-op</i>	R311, R722, C704, R1114, R1115
R311	<i>defined-operator</i>	C476, R1207, C1202
R703	<i>defined-unary-op</i>	R311, R702, C703, R1114, R1115
R430	<i>derived-type-def</i>	R207, C438, C442, C443
R457	<i>derived-type-spec</i>	R402, C403, R403, C405, C406, R459, C496, C503, R848, C833, C834, C835, R920, C937, C938
R431	<i>derived-type-stmt</i>	R430, C432, C439, C442, C443

R603	<i>designator</i>	R448, C467, R542, R601, C601, R615, C620, R616, C621, R701, C702, C745
—	<i>digit</i>	R302, R313, R410, R426, C412, R427, C413, R429, R426, C429, R427, C430, R429
R410	<i>digit-string</i>	R407, R408, R409, R413, R414
R544	<i>dimension-decl</i>	R543
R543	<i>dimension-stmt</i>	R212
R828	<i>do-block</i>	R822
R833	<i>do-body</i>	R832, R835, R837
R821	<i>do-construct</i>	R213, C822, C839
—	<i>do-construct-name</i>	R824, R825, R830, C815, R839, C822
R823	<i>do-stmt</i>	R822, C815, C816, C817
R834	<i>do-term-action-stmt</i>	C325, R832, C818, C819
R838	<i>do-term-shared-stmt</i>	R837, C820, C821
R827	<i>do-variable</i>	R474, R537, R826, C814, R919, C934
R1208	<i>dtio-generic-spec</i>	C478, R1207, C1202
R1235	<i>dummy-arg</i>	R1234, R1240, C1265
R1228	<i>dummy-arg-name</i>	R545, R546, R557, R1226, C1242, R1235, R1243, C1271, C1272, C1273
R1010	<i>e</i>	R1006, C1005, C1008
R842	<i>else-if-stmt</i>	R840, C828
R843	<i>else-stmt</i>	R840, C828
R750	<i>elsewhere-stmt</i>	R744, C735
R806	<i>end-associate-stmt</i>	R802, C804
R1122	<i>end-block-data-stmt</i>	R1120, C1119
R809	<i>end-block-stmt</i>	R807, C806
R820	<i>end-critical-stmt</i>	R818, C812
R829	<i>end-do</i>	R822, C815, C816, C817
R830	<i>end-do-stmt</i>	R829, C815, C816
R466	<i>end-enum-stmt</i>	R462
R758	<i>end-forall-stmt</i>	R752, C737
R1232	<i>end-function-stmt</i>	R214, C201, C324, C325, C818, C820, C829, R1205, R1225, C1254
R844	<i>end-if-stmt</i>	R840, C828
R1204	<i>end-interface-stmt</i>	R1201, C1202
R336	<i>end-macro-stmt</i>	R314
R1106	<i>end-module-stmt</i>	R1104, C1103
R1239	<i>end-mp-subprogram-stmt</i>	R1237, C1260
R1103	<i>end-program-stmt</i>	R214, C201, C324, C325, C818, C820, C829, R1101, C1102
R813	<i>end-select-stmt</i>	R810, C807
R849	<i>end-select-type-stmt</i>	R846, C838
R1119	<i>end-submodule-stmt</i>	R1116, C1118
R1236	<i>end-subroutine-stmt</i>	R214, C201, C324, C325, C818, C820, C829, R1205, R1233, C1258
R434	<i>end-type-stmt</i>	R430
R751	<i>end-where-stmt</i>	R744, C735
R924	<i>endfile-stmt</i>	R214, C1283
R503	<i>entity-decl</i>	R501, C502, C503, C505

—	<i>entity-name</i>	R530, R551
—	<i>entry-name</i>	C1237, R1240, C1261, C1264, C1266
R1240	<i>entry-stmt</i>	R206, R207, R209, C1104, C1115, C1206, C1262, C1263, C1264
R462	<i>enum-def</i>	R207
R463	<i>enum-def-stmt</i>	R462
R465	<i>enumerator</i>	R464, C506
R464	<i>enumerator-def-stmt</i>	R462
R721	<i>equiv-op</i>	R310, R717
R716	<i>equiv-operand</i>	R716, R717
R566	<i>equivalence-object</i>	R565, C585, C586, C587, C588, C589, C590, C591, C592, C593, C594, C595
R565	<i>equivalence-set</i>	R564
R564	<i>equivalence-stmt</i>	R212
R629	<i>errmsg-variable</i>	R627, R641, R858
R213	<i>executable-construct</i>	R208, R209, C1262
R208	<i>execution-part</i>	C201, R1101, C1101, R1225, R1233, R1237
R209	<i>execution-part-construct</i>	R208, R801, R833
R850	<i>exit-stmt</i>	R214, C325, C818, C820, C840
R338	<i>expand-stmt</i>	R322, C330
R520	<i>explicit-co-shape-spec</i>	R511, C538
R512	<i>explicit-shape-spec</i>	R444, C451, C452, C515, R510, C531, R517, C597
R416	<i>exponent</i>	R413
R415	<i>exponent-letter</i>	R413, C415
R722	<i>expr</i>	R461, R471, R601, C602, R630, R701, R722, R724, R725, R726, R727, R728, R730, R734, R739, C728, R742, C731, R805, C803, R916, R1223, R1243, C1269, C1270, C1274
R312	<i>extended-intrinsic-op</i>	R311
—	<i>external-name</i>	R1210
R1210	<i>external-stmt</i>	R212
R203	<i>external-subprogram</i>	R202, C1262
R906	<i>file-name-expr</i>	R905, R930
R902	<i>file-unit-number</i>	R901, R905, C904, C906, R909, C908, C920, C926, R922, C940, R923, R924, R925, R926, C943, R927, R928, C946, R930, C950, C1284
R456	<i>final-binding</i>	R452
—	<i>final-subroutine-name</i>	R456, C487, C488
R928	<i>flush-spec</i>	R927, C945, C946
R927	<i>flush-stmt</i>	R214, C1283
R757	<i>forall-assignment-stmt</i>	R756, C745, R759
R756	<i>forall-body-construct</i>	R752, C742, C743, C744
R752	<i>forall-construct</i>	R213, R756, C742
—	<i>forall-construct-name</i>	R753, R758, C737
R753	<i>forall-construct-stmt</i>	R752, C737
R754	<i>forall-header</i>	R753, R759, R826
R759	<i>forall-stmt</i>	R214, R756
R755	<i>forall-triplet-spec</i>	R754, C741
R914	<i>format</i>	R910, R912, R913, C917, C918, C921, C929, C930

R1003	<i>format-item</i>	R1002, C1002, R1003, R1004
R1002	<i>format-specification</i>	R1001
R1001	<i>format-stmt</i>	R206, R207, R209, C1001, C1104, C1115, C1206
—	<i>function-name</i>	R503, C508, C1203, R1226, C1237, C1238, R1232, C1254, C1266, R1243, C1272
R1219	<i>function-reference</i>	R506, C512, R701
R1226	<i>function-stmt</i>	R1205, C1203, R1225, C1239, C1246, C1247, C1254
R1225	<i>function-subprogram</i>	R203, R211, R1108
R454	<i>generic-binding</i>	R452, C473, C1110, C1111
—	<i>generic-name</i>	C475, R1207, C1202
R1207	<i>generic-spec</i>	R454, C473, C475, C476, C477, C478, R525, R1112, C1111, C1112, R1203, R1204, C1202, C1204
R851	<i>goto-stmt</i>	R214, C325, C818, C820, C841
R428	<i>hex-constant</i>	R425
R429	<i>hex-digit</i>	R428, R1021
R840	<i>if-construct</i>	R213, C828
—	<i>if-construct-name</i>	R841, R842, R843, R844, C828
R845	<i>if-stmt</i>	R214, C324, C829
R841	<i>if-then-stmt</i>	R840, C828
R419	<i>imag-part</i>	R417
R624	<i>image-selector</i>	R613, C615, C617
R862	<i>image-set</i>	C849, R863, R864
R860	<i>image-team</i>	C848, R905, R930
R205	<i>implicit-part</i>	R204
R206	<i>implicit-part-stmt</i>	R205
R560	<i>implicit-spec</i>	R559
R559	<i>implicit-stmt</i>	R205, R206
R518	<i>implied-shape-spec</i>	R510
—	<i>import-name</i>	R1209, C1211
R1209	<i>import-stmt</i>	R204
—	<i>index-name</i>	R755, C740, C741, C742
R448	<i>initial-data-target</i>	R447, C466, R505, C511, R540
R1217	<i>initial-proc-target</i>	R1216
R505	<i>initialization</i>	R503, C505, C506, C507, C510
R730	<i>initialization-expr</i>	R447, R505, R548, C712
R837	<i>inner-shared-do-construct</i>	R836, C821
R915	<i>input-item</i>	R910, C916, R918, C932, C933, C935
R930	<i>inquire-spec</i>	R929, C948, C949, C950, C951
R929	<i>inquire-stmt</i>	R214, C1283
R308	<i>int-constant</i>	C302, R539
—	<i>int-constant-name</i>	R408, C409
R541	<i>int-constant-subobject</i>	R539, C576
R727	<i>int-expr</i>	R401, R473, R611, R619, R622, R623, R625, R634, R635, C708, R729, C711, R732, R814, R826, R852, R862, C849, R902, R905, R913, R919, R922, R930, R1241, C1268
R732	<i>int-initialization-expr</i>	R405, C408, R420, C418, R437, R465, R535, C714, R817, R856, C846

R407	<i>int-literal-constant</i>	R306, R406, R422, C419, R1005, R1007, R1008, R1009, R1010, R1015
R608	<i>int-variable</i>	C608, R628, R905, R909, R913, R922, R926, R928, R929, R930
—	<i>int-variable-name</i>	R827
R523	<i>intent-spec</i>	R502, R545, R1213
R545	<i>intent-stmt</i>	R212
R1201	<i>interface-block</i>	R207, C1201
R1205	<i>interface-body</i>	R1202, C1201, C1205, C1206, C1210
R1215	<i>interface-name</i>	R453, C471, C484, R1212, C1220
R1202	<i>interface-specification</i>	R1201
R1203	<i>interface-stmt</i>	R1201, C1202, C1203
R903	<i>internal-file-variable</i>	R901, C922
R211	<i>internal-subprogram</i>	R210
R210	<i>internal-subprogram-part</i>	R1101, R1225, C1253, R1233, C1257, R1237
R310	<i>intrinsic-operator</i>	R312, C703, C704
—	<i>intrinsic-procedure-name</i>	R1218, C1221
R1218	<i>intrinsic-stmt</i>	R212
R404	<i>intrinsic-type-spec</i>	R402, R403
R913	<i>io-control-spec</i>	R910, R911, C910, C911, C917, C918, C919, C920, C928
R917	<i>io-implied-do</i>	R915, R916
R919	<i>io-implied-do-control</i>	R917
R918	<i>io-implied-do-object</i>	R917, C935
R901	<i>io-unit</i>	R913, C911, C918, C919, C920, C922, C926, C1284
R907	<i>iomsg-variable</i>	R905, R909, R913, R922, R926, R928, R930
R1013	<i>k</i>	R1012, C1010
R215	<i>keyword</i>	R458, C493, C494, R460, C500, C501, R1222, C1228, C1229, C1230
R408	<i>kind-param</i>	R407, C410, C411, R413, C415, C416, C419, R423, C427, R424, C428, R425
R405	<i>kind-selector</i>	R404, R436
R313	<i>label</i>	C304, R824, C817, R851, C841, R852, C842, R853, C843, R905, C905, R909, C909, R913, C914, R914, C931, R922, C941, R926, C944, R928, C947, R930, C952, R1224, C1233
R824	<i>label-do-stmt</i>	R823, C817, R832, C819, R835, R837, C821
R508	<i>language-binding-spec</i>	R502, C502, C503, R529, C564, R1227
R469	<i>lbracket</i>	R343, R442, R467, R503, R509, R527, R544, R556, R624, R631
R421	<i>length-selector</i>	R420, C424, C425
—	<i>letter</i>	R302, R304, R561, C581, R703, R723
R561	<i>letter-spec</i>	R560
R702	<i>level-1-expr</i>	R704
R706	<i>level-2-expr</i>	R706, R710
R710	<i>level-3-expr</i>	R710, R712
R712	<i>level-4-expr</i>	R714
R717	<i>level-5-expr</i>	R717, R722
R306	<i>literal-constant</i>	R305
R1114	<i>local-defined-operator</i>	R1111

—	<i>local-name</i>	R1111
R724	<i>logical-expr</i>	C705, R733, R748, R814, R826, R841, R842, R845
R733	<i>logical-initialization-expr</i>	C715, R817
R424	<i>logical-literal-constant</i>	R306, C703, C704
R604	<i>logical-variable</i>	C604
R826	<i>loop-control</i>	R824, R825
R513	<i>lower-bound</i>	R512, R515, R517, R518
R634	<i>lower-bound-expr</i>	R633, R636, R637, R737, R738
R521	<i>lower-co-bound</i>	R520, C539
R1008	<i>m</i>	R1006, C1008
R339	<i>macro-actual-arg</i>	C327, C328, C330
R321	<i>macro-body-block</i>	R314, R323, R327
R322	<i>macro-body-construct</i>	R321, C309
R333	<i>macro-body-stmt</i>	R322, C315, C316
R332	<i>macro-condition</i>	R328, R329
R317	<i>macro-declaration-stmt</i>	R314
R314	<i>macro-definition</i>	R207, R322, C309
R323	<i>macro-do-construct</i>	R322
R325	<i>macro-do-limit</i>	C311
R324	<i>macro-do-stmt</i>	R323
—	<i>macro-do-variable-name</i>	C310
—	<i>macro-dummy-arg-name</i>	C305, C307
—	<i>macro-dummy-arg-name-list</i>	C305
—	<i>macro-dummy-name</i>	C329, C330
R329	<i>macro-else-if-stmt</i>	R327
R330	<i>macro-else-stmt</i>	R327
R326	<i>macro-end-do-stmt</i>	R323
R331	<i>macro-end-if-stmt</i>	R327
R337	<i>macro-expr</i>	R320, C308, R325, R332, C318
R327	<i>macro-if-construct</i>	R322
R328	<i>macro-if-then-stmt</i>	R327
—	<i>macro-local-variable-name</i>	C306
—	<i>macro-name</i>	R336, C317, C319
R319	<i>macro-optional-decl-stmt</i>	R317
R318	<i>macro-type-declaration-stmt</i>	R317
R1101	<i>main-program</i>	R202, C1101
R748	<i>mask-expr</i>	R743, R745, R749, R754, C738, C739
R749	<i>masked-elsewhere-stmt</i>	R744, C735
R1104	<i>module</i>	R202
—	<i>module-name</i>	R1105, R1106, C1103, R1109, C1107, C1108
R1110	<i>module-nature</i>	R1109, C1107, C1108
R1105	<i>module-stmt</i>	R1104, C1103
R1108	<i>module-subprogram</i>	R1107, C1262
R1107	<i>module-subprogram-part</i>	R1104, R1116
R1238	<i>mp-subprogram-stmt</i>	R1237
R708	<i>mult-op</i>	R310, R705

R704	<i>mult-operand</i>	R704, R705
R1015	<i>n</i>	R1014, C1011, C1012
R304	<i>name</i>	R102, R215, C301, R307, R504, R554, R602, R1215, C1213, C1214, R1228
R307	<i>named-constant</i>	R305, R418, R419, R465, R548
R548	<i>named-constant-def</i>	R547
—	<i>namelist-group-name</i>	R562, C582, C584, R913, C915, C916, C917, C919, C921, C929, C930
R563	<i>namelist-group-object</i>	R562, C583, C584
R562	<i>namelist-stmt</i>	R212
R343	<i>nested-token-sequence</i>	R341
R831	<i>nonblock-do-construct</i>	R821
R825	<i>nonlabel-do-stmt</i>	R823, C816
R718	<i>not-op</i>	R310, R714
R863	<i>notify-stmt</i>	R214
R506	<i>null-init</i>	R447, R505, R540, R1216
R638	<i>nullify-stmt</i>	R214
R728	<i>numeric-expr</i>	C709, R853, C844
R504	<i>object-name</i>	R503, C506, C509, C511, R526, R527, R528, R531, R550, R553, R555, R556, R558, R603
R427	<i>octal-constant</i>	R425
R1112	<i>only</i>	R1109
R1113	<i>only-use-name</i>	R1112
R904	<i>open-stmt</i>	R214, C1283
R546	<i>optional-stmt</i>	R212
R720	<i>or-op</i>	R310, R716
R715	<i>or-operand</i>	R715, R716
R835	<i>outer-shared-do-construct</i>	R831, R836, C821
R916	<i>output-item</i>	R911, R912, C916, R918, C935, C936, R929
R547	<i>parameter-stmt</i>	R206, R207
R1118	<i>parent-identifier</i>	R1117
R610	<i>parent-string</i>	R609, C609
—	<i>parent-submodule-name</i>	R1118, C1117
—	<i>parent-type-name</i>	R432, C433
—	<i>part-name</i>	R613, C610, C611, C612, C613, C614, C615, C618, C619, C624
R613	<i>part-ref</i>	C568, C571, C586, R612, C618, C619, C622, C623
R735	<i>pointer-assignment-stmt</i>	R214, C552, R757
R550	<i>pointer-decl</i>	R549
R639	<i>pointer-object</i>	R638, C642
R549	<i>pointer-stmt</i>	R212
R1014	<i>position-edit-desc</i>	R1012
R926	<i>position-spec</i>	R923, R924, R925, C942, C943
R707	<i>power-op</i>	R310, R704
R1229	<i>prefix</i>	R1226, C1243, C1244, C1245, C1246, C1248, C1249, C1250, C1251, R1234, C1255
R1230	<i>prefix-spec</i>	R1229, C1243
—	<i>primaries</i>	C1269

R701	<i>primary</i>	R702
R912	<i>print-stmt</i>	R214, C1283
R449	<i>private-components-stmt</i>	R433, C468
R433	<i>private-or-sequence</i>	R430, C438
R1213	<i>proc-attr-spec</i>	R1211
R452	<i>proc-binding-stmt</i>	R450
R446	<i>proc-component-attr-spec</i>	R445, C457, C458, C461
R445	<i>proc-component-def-stmt</i>	R440, C457
R741	<i>proc-component-ref</i>	R740, R742, R1221, R1223
R1214	<i>proc-decl</i>	R445, R1211, C1217, C1219
—	<i>proc-entity-name</i>	R550
R1212	<i>proc-interface</i>	R445, R1211, C1216, C1220
R1227	<i>proc-language-binding-spec</i>	C516, R1213, C1219, C1220, C1239, C1240, C1241, C1246, R1231, R1234
R1216	<i>proc-pointer-init</i>	R1214
R554	<i>proc-pointer-name</i>	R553, R568, C598, C601, R639, R740
R740	<i>proc-pointer-object</i>	R735
R742	<i>proc-target</i>	R461, C504, C552, R735, C733
—	<i>procedure-component-name</i>	C730
R1211	<i>procedure-declaration-stmt</i>	R207, C1213
R1221	<i>procedure-designator</i>	R1219, C1222, R1220, C1224
—	<i>procedure-entity-name</i>	R1214, C1216
—	<i>procedure-name</i>	R453, C470, C471, C472, R742, C732, R1206, C1207, C1208, C1209, R1217, C1218, R1221, C1225, R1223, C1232, R1238, R1239, C1259, C1260
R1206	<i>procedure-stmt</i>	R1202, C1204, C1208, C1209
—	<i>program-name</i>	R1102, R1103, C1102
R1102	<i>program-stmt</i>	R1101, C1102
R202	<i>program-unit</i>	R201
R551	<i>protected-stmt</i>	R212
R865	<i>query-spec</i>	R864, C850
R864	<i>query-stmt</i>	R214
R1005	<i>r</i>	R1003, C1003, C1004, R1012
R470	<i>rbracket</i>	R343, R442, R467, R503, R509, R527, R544, R556, R624, R631
R910	<i>read-stmt</i>	R214, C912, C1284
R413	<i>real-literal-constant</i>	R306, R412
R418	<i>real-part</i>	R417
R713	<i>rel-op</i>	R310, R712
R1111	<i>rename</i>	R1109, R1112
—	<i>rep-char</i>	R423
—	<i>result-name</i>	C1237, R1231, C1266
R334	<i>result-token</i>	R333, C313, C314
R1241	<i>return-stmt</i>	R214, C325, C818, C820, C1267
R925	<i>rewind-stmt</i>	R214, C1283
R1018	<i>round-edit-desc</i>	R1012
R552	<i>save-stmt</i>	R212

R553	<i>saved-entity</i>	R552
R103	<i>scalar-xyz</i>	C101
R620	<i>section-subscript</i>	R613, C614, C615, C623
R811	<i>select-case-stmt</i>	R810, C807
—	<i>select-construct-name</i>	R847, R848, R849, C838
R846	<i>select-type-construct</i>	R213, C836, C837, C838
R847	<i>select-type-stmt</i>	R846, C832, C838
R805	<i>selector</i>	R804, C801, R847, C830, C831, C832, C835
R1237	<i>separate-module-subprogram</i>	R1108, C1259
R435	<i>sequence-stmt</i>	R433
R836	<i>shared-term-do-construct</i>	R835
R411	<i>sign</i>	R406, R409, R412
R1016	<i>sign-edit-desc</i>	R1012
R409	<i>signed-digit-string</i>	R416
R406	<i>signed-int-literal-constant</i>	R418, R419, R540, R1011, R1013
R412	<i>signed-real-literal-constant</i>	R418, R419, R540
R414	<i>significand</i>	R413
R630	<i>source-expr</i>	R627, C628, C630, C636, C637, C638, C640
—	<i>special-character</i>	R301
R453	<i>specific-binding</i>	R452
R729	<i>specification-expr</i>	C404, C420, R513, R514, R521, R522, C1290
R204	<i>specification-part</i>	C473, C517, C562, R807, C805, R1101, R1104, C1104, C1105, C1106, R1116, C1115, C1116, R1120, C1120, C1121, R1205, R1225, R1233, R1237, C1275, C1276, C1277, C1278
R212	<i>specification-stmt</i>	R207
R628	<i>stat-variable</i>	R627, R641, R858
R1243	<i>stmt-function-stmt</i>	R207, C1104, C1115, C1206
R856	<i>stop-code</i>	R855
R855	<i>stop-stmt</i>	R214, C325, C818, C820, C1285
R622	<i>stride</i>	R621, R755, C741
R614	<i>structure-component</i>	R536, C570, C571, C572, R603, R610, R632, R639
R459	<i>structure-constructor</i>	R540, C575, R701
R1116	<i>submodule</i>	R202
—	<i>submodule-name</i>	R1117, R1119, C1118
R1117	<i>submodule-stmt</i>	R1116, C1118
—	<i>subroutine-name</i>	C1203, R1234, R1236, C1258
R1234	<i>subroutine-stmt</i>	R1205, C1203, C1239, C1246, C1247, R1233, C1255, C1258
R1233	<i>subroutine-subprogram</i>	R203, R211, R1108
R619	<i>subscript</i>	C571, C622, R620, R621, R755, C741
R621	<i>subscript-triplet</i>	R620, C626
R609	<i>substring</i>	R566, C596, R603
R611	<i>substring-range</i>	R609, R618, C624
R1231	<i>suffix</i>	R1226, R1240
R857	<i>sync-all-stmt</i>	R214
R861	<i>sync-images-stmt</i>	R214

R866	<i>sync-memory-stmt</i>	R214
R858	<i>sync-stat</i>	R857, C847, R863, R865, R866
R859	<i>sync-team-stmt</i>	R214
R556	<i>target-decl</i>	R555
R555	<i>target-stmt</i>	R212
R335	<i>token</i>	R334, C314
R432	<i>type-attr-spec</i>	R431, C432
R450	<i>type-bound-procedure-part</i>	R430, C441, C1504
R501	<i>type-declaration-stmt</i>	R207, C424, C425, C501
R848	<i>type-guard-stmt</i>	R846, C836, C837, C838
—	<i>type-name</i>	R431, C431, R434, C439, C478, R457, C490
R438	<i>type-param-attr-spec</i>	R436
R437	<i>type-param-decl</i>	R436
R436	<i>type-param-def-stmt</i>	R430, C442, C443
R616	<i>type-param-inquiry</i>	R701
—	<i>type-param-name</i>	R431, R437, C442, C443, R616, C621, R701, C701
R458	<i>type-param-spec</i>	R457, C491, C492, C493, C495
R401	<i>type-param-value</i>	C401, C402, C404, R420, R421, R422, C420, C421, C422, C423, C456, R458, C495, C631
R402	<i>type-spec</i>	R468, C507, C508, C509, R626, C628, C629, C630, C631, C632, C636, C639, R848, C833, C834, C835
R303	<i>underscore</i>	R302
R1004	<i>unlimited-format-item</i>	R1002
R514	<i>upper-bound</i>	R512
R635	<i>upper-bound-expr</i>	R633, R637, R738
R522	<i>upper-co-bound</i>	R520, C539
R1115	<i>use-defined-operator</i>	R1111, C1110, C1114
—	<i>use-name</i>	R525, C563, R1111, R1113, C1113
R1109	<i>use-stmt</i>	R204
R1011	<i>v</i>	R1006, C1008
R557	<i>value-stmt</i>	R212
R601	<i>variable</i>	R534, C566, C568, R604, R605, R606, R607, R608, R734, C716, R736, C723, C724, R739, C726, C729, C730, C745, R805, C801, C830, C831, R915, R1223
R602	<i>variable-name</i>	R563, R566, R568, C598, C601, C603, R610, R632, R639, R736, C722
R623	<i>vector-subscript</i>	R620, C625
R558	<i>volatile-stmt</i>	R212
R1007	<i>w</i>	R1006, C1006, C1007, C1008
R922	<i>wait-spec</i>	R921, C939, C940
R921	<i>wait-stmt</i>	R214, C1283
R747	<i>where-assignment-stmt</i>	R743, R746, C734
R746	<i>where-body-construct</i>	R744, C736
R744	<i>where-construct</i>	R213, R746, R756
—	<i>where-construct-name</i>	R745, R749, R750, R751, C735
R745	<i>where-construct-stmt</i>	R744, C735
R743	<i>where-stmt</i>	R214, R746, R756
R911	<i>write-stmt</i>	R214, C913, C1284

Annex E

(Informative)

Index

In this annex, entries in *italics* denote BNF terms, entries in **bold face** denote language keywords, and page numbers in **bold face** denote primary text or glossary definitions.

Symbols

<, 152
<=, 152
>, 152
>=, 152
(, 154
*, 31, 42, 44, 45, 50, 91, 93, 101, 125, 151, 231, 263, 276, 282, 309, 329
**, 151
+, 151
-, 151
.AND., 153, 154
.EQ., 152
.EQV., 153, 154
.GE., 152
.GT., 152
.LE., 152
.LT., 152
.NE., 152
.NEQV., 153, 154
.NOT., 153, 154
.OR., 153, 154
.XOR., 153, 154
/, 151
//, 51, 151
/=, 152
;;, 30
==, 152
&, 30, 280

A

abstract interface, **297**
abstract interface block, **300**
abstract type, **72**, **505**
ac-do-variable (R474), 79, **79**, 80, 142, 144, 486, 487
ac-implicit-do (R472), 79, **79**, 80, 146, 332, 487
ac-implicit-do-control (R473), 79, **79**, 142–144, 146
ac-spec (R468), 79, **79**
ac-value (R471), 79, **79**, 80
access methods, **206**
access-id (R525), 99, **99**

access-spec (R507), 33, 55, 56, 60, 62, 67–70, 83, 86, **86**, 99, 306
access-stmt (R524), 10, 99, **99**
ACCESS= specifier, 216, 247
accessibility attribute, **86**
accessibility statement, 99
action-stmt (R214), 11, **11**, 35, 183, 189, 192
action-term-do-construct (R832), 183, **183**
ACTION= specifier, 217, 247
actions, **206**
active, **183**
active combination of, **171**
actual argument, **295**, **505**
actual-arg (R1223), 309, **309**
actual-arg-spec (R1222), 75, 308, 309, **309**
add-op (R709), 27, 134, **134**
add-operand (R705), 134, **134**, 155
ADVANCE= specifier, 224
advancing input/output statement, **210**
affector, **225**
alloc-opt (R627), 124, **124**, 125
allocatable array, **91**
ALLOCATABLE attribute, **86**, 100
ALLOCATABLE statement, 100
allocatable variable, **505**
allocatable-decl (R527), 100, **100**
allocatable-stmt (R526), 10, **100**, 488
ALLOCATE statement, **124**
allocate-co-array-spec (R636), 125, **125**, 128
allocate-co-shape-spec (R637), 125, **125**
allocate-object (R632), 50, 51, 124, 125, **125**, 126, 127, 129, 503, 504
allocate-shape-spec (R633), 124, 125, **125**
allocate-stmt (R626), 11, **124**, 504
allocated, **127**
allocation (R631), 124, **124**, 126, 128
alphanumeric-character (R302), 25, **25**, 26
alt-return-spec (R1224), 192, 308, 309, **309**
ancestor, **291**
ancestor component, **73**
ancestor-module-name, 292
and-op (R719), 27, 136, **136**
and-operand (R714), 136, **136**

- approximation methods, **46**
 - arg-name*, 62, 64, 68, 69
 - arg-token* (R344), 36, **36**
 - argument, **505**
 - argument association, 310, 487, **505**
 - argument keyword, **21**, 310, 339, 486
 - arithmetic IF statement, 193
 - arithmetic-if-stmt* (R853), 11, 35, 183, 193, **193**
 - array, **19**, 88–92, **120**, 120–123, **505**
 - assumed-shape, 90
 - assumed-size, 91
 - deferred-shape, 90
 - explicit-shape, 90
 - array constructor, 79, **79**
 - array element, **19**, **120**, 121, **505**
 - array element order, **121**
 - array intrinsic assignment statement, **158**
 - array pointer, **91**, **505**
 - array section, **19**, **122**, **505**
 - array-constructor* (R467), 79, **79**, 80, 133
 - array-element* (R617), 101, 108, 115, 116, **120**
 - array-name*, 103, 488
 - array-section* (R618), 115, 120, **120**, 121, 122
 - array-spec* (R510), 7, 84, 86, 88, 89, **89**, 91, 100, 103, 105, 111, 127
 - ASCII character set, **49**
 - ASCII character type, **49**, 158, 212, 229, 262, 277, 409, 420
 - ASCII collating sequence, **52**, 350, 384, 387, 394–396, 409
 - assignment, 156–174
 - defined, 161
 - elemental array (FORALL), 168
 - intrinsic, 157
 - masked array (WHERE), 166
 - pointer, 162
 - assignment statement, 156, **505**
 - assignment-stmt* (R734), 11, 157, **157**, 166, 168, 169, 171, 335, 503
 - ASSOCIATE construct, 176, **176**, 487
 - associate name, **176**, **505**
 - associate-construct* (R802), 11, 176, **176**
 - associate-construct-name*, 176
 - associate-name*, 176, 189, 190, 192, 332, 486, 515
 - associate-stmt* (R803), 176, **176**, 192
 - associated, **20**
 - associating entity, **498**
 - association, **21**, **505**
 - argument, 310, 487
 - common, 112
 - host, 488
 - inheritance, 498
 - name, 487
 - pointer, 491
 - sequence, 318
 - storage, 494
 - use, 488
 - association* (R804), 176, **176**
 - association status, pointer, 491
 - assumed type parameter, **43**
 - assumed-shape array, **90**, **505**
 - assumed-shape-spec* (R515), 89, 90, **90**
 - assumed-size array, **91**, **505**
 - assumed-size-spec* (R517), 89, **91**
 - asynchronous, 217, 220, **224**, 226, 227, 230, 231, 238, 241–243, 245, 247, 250
 - ASYNCHRONOUS attribute, **86**, 100
 - ASYNCHRONOUS statement, 100
 - asynchronous-stmt* (R528), 10, **100**
 - ASYNCHRONOUS= specifier**, 217, 224, 247
 - attr-spec* (R502), 83, **83**, 84, 85, 104
 - attribute, 85–99, **505**
 - accessibility, 86
 - ALLOCATABLE, 86, 100
 - ASYNCHRONOUS, 86, 100
 - BIND, 55, 58, 100
 - CONTIGUOUS, **87**, 100
 - DIMENSION, 88, 103
 - EXTERNAL, 92
 - INTENT, 93, 103
 - INTRINSIC, 95
 - OPTIONAL, 95, 103
 - PARAMETER, 96, 104
 - POINTER, 96, 104
 - PRIVATE, 56, 86, 99
 - PROTECTED, 96, 104
 - PUBLIC, 56, 86, 99
 - SAVE, 101, 104
 - TARGET, 98, 105
 - VALUE, 98, 105
 - VOLATILE, 98, 105
 - attribute specification statements, 99–113
 - attributes, **83**
 - automatic data object, **84**, **505**
- ## B
- BACKSPACE statement, 243
 - backspace-stmt* (R923), 11, **243**, 335
 - base object, **117**
 - base type, **72**, **506**
 - basic-token* (R342), 36, **36**
 - basic-token-sequence* (R341), 34, 36, **36**
 - belong, **185**, **192**, **506**
 - belongs, **192**
 - binary-constant* (R426), 53, **53**, 54
 - BIND attribute, 55, 58, 100
 - BIND statement, 100
 - BIND(C)**, 78, 87, 296, 298, 475, 477, 481

bind-entity (R530), 100, **100**
bind-stmt (R529), 10, **100**
 binding, 69, 485
 binding label, **480, 481, 506**
 binding name, **69**
binding-attr (R455), 68, **68, 69**
binding-name, 68, 69, 309, 325, 486, 516
binding-private-stmt (R451), 68, **68, 70**
 bit model, 340
 bits compatible, **311, 506**
 bits conversion, 160
 bits editing, 263
 bits intrinsic assignment statement, **158**
 bits intrinsic operation, **138, 154**
 bits intrinsic operator, **138**
 bits relational intrinsic operation, **138**
 bits type, **53**
 blank common, **111**
blank-interp-edit-desc (R1017), 259, **260**
BLANK= specifier, 217, 225, 247
 block, **175, 506**
block (R801), **175, 176–178, 181, 182, 188, 189**
 block data program unit, **292, 506**
block-construct (R807), 11, **177, 178**
block-construct-name, 177, 178
block-data (R1120), 9, 292, **292**
block-data-name, 292, 293
block-data-stmt (R1121), 9, 292, **292**
block-do-construct (R822), 182, **182**
block-stmt (R808), 177, **177, 178, 192**
 blocking, **200**
 bounds, **506**
bounds-remapping (R738), 162, 163, **163, 164**
bounds-spec (R737), 162, 163, **163, 164**
boz-literal-constant (R425), 27, **53**
 branch target statement, **192**
 branching, **192**

C

C address, **468**
 C character kind, **468**
 C_(C type), 467–478
 C_LOC function, 471
 CALL statement, 308
call-stmt (R1220), 11, **308, 309, 310**
 CASE construct, **178**
 case index, **179**
case-construct (R810), 11, 178, **178, 179**
case-construct-name, 178
case-expr (R814), 178, **178, 179**
case-selector (R815), 178, **178**
case-stmt (R812), 178, **178**
case-value (R817), 178, **178**
case-value-range (R816), 178, **178, 179**

CHAR intrinsic, 52
char-constant (R309), 27, **27**
char-expr (R725), **139, 140, 144, 178**
char-initialization-expr (R731), 87, 144, **144, 178,**
 193, 222, 223
char-length (R422), 50, **50, 61, 83–85**
char-literal-constant (R423), 27, 32, **51, 237, 259,**
 260
char-selector (R420), 45, 50, **50**
char-string-edit-desc (R1020), 258, **260**
char-variable (R606), 115, **115, 212**
CHARACTER, 49
character (R301), **25**
 character context, **29**
 character intrinsic assignment statement, **158**
 character intrinsic operation, **138, 151**
 character intrinsic operator, **138**
 character length parameter, **506**
 character literal constant, **51**
 character relational intrinsic operation, **138**
 character sequence type, **57, 109, 497, 612**
 character set, 25
 character storage unit, **494, 506**
 character string, **49, 506**
 character string edit descriptor, **258**
 character type, **49, 49–53**
 CHARACTER_STORAGE_SIZE, **435**
 characteristics, **506**
 characteristics of a procedure, **296**
 child, **291**
 child data transfer statement, **234, 234–238, 255**
CLASS, 44
 class, **506**
 CLOSE statement, **220**
close-spec (R909), 221, **221**
close-stmt (R908), 11, **221, 335**
 co-array, **20, 506**
 explicit-co-shape, 93
co-array-spec (R511), 60–62, 84, 88, 89, **89, 92, 93,**
 100, 103, 105
 co-bound, **92, 93**
 co-bounds, **507**
 co-dimension, **20**
 co-indexed object, **20, 507**
co-name, 103
 co-rank, **20, 507**
 co-size, **20**
 co-subscript, **507**
co-subscript (R625), 20, 124, **124**
 collating sequence, 52, **52, 507**
 collective subroutine, **339, 507**
 comment, **30, 31**
 common association, **112**
 common block, **111, 485, 507, 552**

- common block storage sequence, **112**
 - COMMON statement, **111**, 111–113
 - common-block-name*, 100, 104, 111, 291
 - common-block-object* (R568), 111, **111**, 291, 488
 - common-stmt* (R567), 10, **111**, 488
 - companion processor, **23**, **507**
 - compatibility
 - FORTRAN 77, 4
 - Fortran 2003, 3
 - Fortran 90, 4
 - Fortran 95, 3
 - COMPILER_OPTIONS, **435**
 - COMPILER_VERSION, **435**
 - COMPLEX**, 48
 - complex part designator, **119**
 - complex type, 48, **48**
 - complex-literal-constant* (R417), 27, **48**
 - complex-part-designator* (R615), 115, 119, **119**, 120
 - component, **60**, 485, **507**
 - component keyword, **21**
 - component order, **66**, **507**
 - component value, **74**
 - component-array-spec* (R444), 60, 61, **61**, 62
 - component-attr-spec* (R442), 60, **60**, 61, 63, 64
 - component-data-source* (R461), 75, **75**, 76, 77
 - component-decl* (R443), 50, 60, **60**, 62, 64
 - component-def-stmt* (R440), 60, **60**, 61
 - component-initialization* (R447), 61, 64, **64**
 - component-name*, 60, 64
 - component-part* (R439), 55, **60**, 67, 70
 - component-spec* (R460), 75, **75**, 144
 - computed GO TO statement, 193
 - computed-goto-stmt* (R852), 11, 193, **193**
 - concat-op* (R711), 27, 135, **135**
 - concatenation, **51**
 - conform, 157
 - conformable, **19**, **507**
 - conformance, 141, **507**
 - connect team, **219**, **507**
 - connect-spec* (R905), 215, **215**, 216, 220
 - connected, **214**, **507**
 - constant, **18**, 27, 41, **507**
 - character, 51
 - integer, 46
 - named, 104
 - constant* (R305), 27, **27**, 101, 116, 133
 - constant subobject, **18**
 - constant-subobject* (R542), 101, 102, **102**
 - construct, **507**
 - construct association, 490, **507**
 - construct entity, **483**, **507**
 - construct-name*, 192
 - constructor
 - array, 79
 - derived-type, 75
 - structure, 75
 - CONTAINS statement, 333
 - contains-stmt* (R1242), 10, 68, 288, **333**
 - contiguous, **87**
 - CONTIGUOUS attribute, **87**, 100
 - CONTIGUOUS statement, 100
 - contiguous-stmt* (R531), **100**
 - continuation, 30, 31
 - CONTINUE statement, 193
 - continue-stmt* (R854), 11, 35, 182, 183, **193**
 - control character, **25**
 - control edit descriptor, **258**, 272
 - control information list, **222**
 - control mask, **507**
 - control-edit-desc* (R1012), 258, **259**
 - conversion
 - bits, 160
 - numeric, 159
 - correspond, **330**
 - critical-construct* (R818), 11, 181, **181**
 - critical-construct-name*, 181
 - critical-stmt* (R819), 181, **181**, 192
 - current record, **209**
 - CYCLE statement, 182, 185
 - cycle-stmt* (R839), 11, 35, 183, 185, **185**
- ## D
- d* (R1009), 259, **259**, 264–267, 270, 271, 279
 - data, **508**
 - data edit descriptor, **258**, 262
 - data edit descriptors, 272
 - data entity, **17**, **508**
 - data object, **18**, **508**
 - data object reference, **22**
 - data pointer, **20**
 - DATA statement, 100, 499
 - data transfer, 232
 - data transfer input statement, **205**
 - data transfer output statements, **205**
 - data transfer statements, 221
 - data type, *see* type, **508**
 - data-component-def-stmt* (R441), 60, **60**, 62
 - data-edit-desc* (R1006), 258, **258**
 - data-i-do-object* (R536), 101, **101**, 102
 - data-i-do-variable* (R537), 101, **101**, 102, 144, 486, 487
 - data-implied-do* (R535), 101, **101**, 102, 144, 487
 - data-pointer-component-name*, 162, 163
 - data-pointer-initialization compatible, **64**
 - data-pointer-object* (R736), 162, **162**, 163, 164, 171, 355, 503, 504
 - data-ref* (R612), 117, **117**, 118, 120, 163, 225, 309, 310, 317, 325

- data-stmt* (R532), 10, **100**, 299, 334, 488
- data-stmt-constant* (R540), 101, **101**, 102
- data-stmt-object* (R534), 100, 101, **101**, 102
- data-stmt-repeat* (R539), 101, **101**, 102
- data-stmt-set* (R533), 100, **101**
- data-stmt-value* (R538), 101, **101**, 102
- data-target* (R739), 75–77, 97, 162, 163, **163**, 164, 171, 319, 334, 355, 494
- datum, **508**
- dealloc-opt* (R641), 129, **129**, 131
- DEALLOCATE statement, **129**
- deallocate-stmt* (R640), 11, **129**, 504
- decimal symbol, **262**, **508**
- decimal-edit-desc* (R1019), 259, **260**
- DECIMAL= specifier**, 217, 225, 248
- declaration, **21**, 83–113
- declaration-construct* (R207), 10, **10**
- declaration-type-spec* (R403), 43, **43**, 44, 50, 60, 61, 83, 85, 105, 142, 306, 326, 329
- declared type, **44**, **508**
- default bits, **53**
- default character, **49**
- default complex, **48**
- default initialization, **64**, **508**
- default integer, **46**
- default logical, **53**
- default real, **47**
- default-char-expr* (R726), 140, **140**, 215–219, 221, 222, 224–227
- default-char-variable* (R607), 115, **115**, 124, 216, 246–252
- default-initialized, **65**, **508**
- default-logical-variable* (R605), 115, **115**, 246, 248, 249
- deferred binding, **69**, **508**
- deferred type parameter, **42**, **508**
- deferred-co-shape-spec* (R519), 61, 89, 92, **92**, 93
- deferred-shape array, **90**
- deferred-shape-spec* (R516), 60, 61, 89, 91, **91**, 104
- definable, **508**
- define-macro-stmt* (R315), 33, **33**, 488
- defined, **22**, **508**
- defined assignment, 303
- defined assignment statement, **161**, **508**
- defined binary operation, **139**
- defined elemental assignment statement, **161**
- defined elemental operation, **139**
- defined operation, **138**, 302, **508**
- defined unary operation, **138**
- defined-binary-op* (R723), 28, 136, **136**, 139, 155, 290
- defined-operator* (R311), **28**, 68, 290, 299
- defined-unary-op* (R703), 28, 134, **134**, 138, 139, 154, 290
- definition, **22**
- definition of variables, 498
- deleted feature, **508**
- deleted features, **7**
- DELIM= specifier**, 217, 226, 248
- delimiters, **29**
- derived type, **17**, 54–77, **508**
- derived type determination, 57
- derived-type intrinsic assignment statement, **158**
- derived-type type specifier, 44
- derived-type-def* (R430), 10, 44, **55**, 56, 59
- derived-type-spec* (R457), 43, 44, 50, **74**, 75, 189, 190, 235, 486
- derived-type-stmt* (R431), 55, **55**, 56, 59, 488
- descendant, **291**
- designator, **21**, **508**
- designator* (R603), 64, 102, 115, **115**, 119, 120, 133, 169
- designator*, 133
- digit*, 6, 25, **25**, 28, 46, 53, 54, 277
- digit-string* (R410), 46, **46**, 47, 263, 264, 269
- digit-string*, 46
- digits, **25**
- DIMENSION attribute, **88**, 103
- DIMENSION statement, 103
- dimension-decl* (R544), 103, **103**
- dimension-spec* (R509), **89**
- dimension-stmt* (R543), 10, **103**, 488
- direct access, 208
- direct access input/output statement, **226**
- direct component, **55**
- DIRECT= specifier**, 248
- disassociated, **20**, **509**
- distinguishable, **304**
- DO construct, 182, **182**
- DO statement, 182
- DO termination, **183**
- DO WHILE statement, 182
- do-block* (R828), 182, **182**, 183, 184
- do-body* (R833), 183, **183**
- do-construct* (R821), 11, **182**, 185, 192
- do-construct-name*, 182, 185
- do-stmt* (R823), 182, **182**, 192, 503
- do-term-action-stmt* (R834), 35, 183, **183**, 185, 192
- do-term-shared-stmt* (R838), 183, **183**, 184, 185, 192
- do-variable* (R827), 79, 101, 182, **182**, 184, 227, 228, 253–255, 278, 500, 502, 503, 544
- DOUBLE PRECISION**, 47
- double precision real, **47**
- dtio-generic-spec* (R1208), 68, 74, 234–236, 240, 299, **299**, 302, 305
- dtv-type-spec* (R920), **235**
- dummy argument, **295**, **509**

restrictions, 319
 dummy array, **509**
 dummy data object, **509**
 dummy procedure, **296, 509**
dummy-arg (R1235), 329, **329**, 331
dummy-arg-name (R1228), 103, 105, 326, **326**, 329, 333, 488
 dynamic type, **44, 509**

E

e (R1010), 259, **259**, 265–267, 270–272, 279
 edit descriptor, *see* format descriptor, **258**
 effective argument, **311**
 effective item, **229, 509**
 effective position, **305**
 element sequence, **318**
 elemental, **19, 295, 509**
 elemental array assignment (FORALL), 168
 elemental intrinsic function, **339**
 elemental operation, 141
 elemental procedure, **336**
else-if-stmt (R842), 188, **188**
else-stmt (R843), 188, **188**
elsewhere-stmt (R750), 166, **166**
ENCODING= specifier, 217, 248
 END statement, 16, **16**
end-associate-stmt (R806), 176, **176**, 192
end-block-data-stmt (R1122), 10, 16, 292, **292**
end-block-stmt (R809), 177, **177**, 178, 192
end-critical-stmt (R820), 181, **181**, 192
end-do (R829), 182, **182**, 183, 185
end-do-stmt (R830), 182, **182**, 192
end-enum-stmt (R466), 77, **78**
end-forall-stmt (R758), 168, **168**, 169
end-function-stmt (R1232), 9, 11, 16, 35, 183, 189, 299, 326, 327, **327**, 333
end-if-stmt (R844), 188, **188**, 192
end-interface-stmt (R1204), 298, **298**, 299
end-macro-stmt (R336), 33, **34**
end-module-stmt (R1106), 9, 16, 288, **288**
end-mp-subprogram-stmt (R1239), 16, 330, **330**
 end-of-file condition, **252**
 end-of-record condition, **252**
end-program-stmt (R1103), 9, 11, 16, 35, 183, 189, 287, **287**
end-select-stmt (R813), 178, **178**, 179, 192
end-select-type-stmt (R849), 189, 190, **190**, 191, 192
end-submodule-stmt (R1119), 9, 16, 292, **292**
end-subroutine-stmt (R1236), 9, 11, 16, 35, 183, 189, 299, 329, **329**, 333
end-type-stmt (R434), 55, **56**
end-where-stmt (R751), 166, **166**
END= specifier, 253

endfile record, **206**
 ENDFILE statement, 206, 244
endfile-stmt (R924), 11, **243**, 335
 ending point, **116**
 entity, **509**
entity-decl (R503), 50, 83, 84, **84**, 85, 143, 144, 488
entity-name, 100, 104
 ENTRY statement, 331, **331**
entry-name, 326, 331, 485
entry-stmt (R1240), 10, 288, 292, 299, 331, **331**, 485, 488
enum-def (R462), 10, **77**, 78
enum-def-stmt (R463), 77, **78**
 enumeration, 77
 enumerator, 77
enumerator (R465), 78, **78**
enumerator-def-stmt (R464), 77, **78**
EOR= specifier, 254
equiv-op (R721), 27, 136, **136**
equiv-operand (R716), 136, **136**
 EQUIVALENCE statement, **108**, 108–110
equivalence-object (R566), 108, **108**, 109, 110, 291
equivalence-set (R565), 108, **108**, 109, 110
equivalence-stmt (R564), 10, **108**, 488
ERR= specifier, 253
errmsg-variable (R629), 124, **124**, 126, 129, 197, 503
ERRMSG= specifier, 129
 ERROR_UNIT, 213, **436**
 evaluation
 optional, 146
 operations, 145
 parentheses, 147
 executable construct, 175, **509**
 executable statement, **14, 509**
executable-construct (R213), 10, **11**, 14, 331
 execution control, 175
 execution cycle, **184**
execution-part (R208), 9, **10**, 11, 287, 326, 327, 329, 330
execution-part-construct (R209), 10, **10**, 175, 183
 exist, **207, 213**
EXIST= specifier, 248
exit-stmt (R850), 11, 35, 183, 192, **192**
expand-stmt (R338), 33, **35**, 36
 explicit, **297**
 explicit formatting, 257–276
 explicit initialization, **85**, 100, **509**
 explicit interface, 297, **509**
 explicit-co-shape co-array, **93**
explicit-co-shape-spec (R520), 89, 93, **93**
 explicit-shape array, **90, 509**
explicit-shape-spec (R512), 61, 62, 86, 89, 90, **90**, 91, 92, 111

exponent (R416), 47, **47**
exponent-letter (R415), 47, **47**, 48
expr (R722), 7, 71, 75, 76, 79, 115, 124, 133, 134, 136, **136**, 139, 140, 144, 157–163, 167, 168, 171, 176, 227, 309, 333, 335
 expression, **19**, **133**, 133–155, **509**
 extended type, **72**, **509**
extended-intrinsic-op (R312), 28, **28**
 extensible type, **72**, **509**
 extension operation, **139**, 154
 extension operator, **139**
 extension type, **72**, **509**
 extent, **19**, **510**
 EXTERNAL attribute, **92**
 external file, **206**, **510**
 external linkage, **510**
 external procedure, **13**, **295**, **510**
 EXTERNAL statement, **305**
 external subprogram, **12**, **510**
 external unit, **212**, **510**
external-name, 305
external-stmt (R1210), 10, **305**, 488
external-subprogram (R203), 9, **9**, 331

F

field, **260**
 field width, **260**
 file, **510**
 file access, 207
 file connection, 212
 file connection statements, **205**
 file inquiry statement, **205**, 245
 file position, 209
 file positioning statements, **205**, 243
 file storage unit, **211**, **510**
file-name-expr (R906), 216, **216**, 218, 246, 247, 249
file-unit-number (R902), 212, **212**, 215, 216, 221, 223, 236, 242–247, 335, 437
FILE= specifier, 218, 247
 FILE_STORAGE_SIZE, **436**
 files
 connected, 214
 external, 206
 internal, 212
 preconnected, 215
 final subroutine, **70**, **510**
final-binding (R456), 68, **70**
final-subroutine-name, 70
 finalizable, **70**, **510**
 finalization, 70, **510**
 finalized, **70**
 fixed source form, 31, **31**
 FLUSH statement, 244
flush-spec (R928), 244, **244**, 245

flush-stmt (R927), 11, **244**, 335
FMT= specifier, 224
 FORALL construct, 168
forall-assignment-stmt (R757), 168, **168**, 169, 171, 173, 335
forall-body-construct (R756), 168, **168**, 169, 171, 173
forall-construct (R752), 11, 168, **168**, 169, 170, 172
forall-construct-name, 168, 169
forall-construct-stmt (R753), 168, **168**, 169, 192
forall-header (R754), 168, **168**, 173, 174, 182
forall-stmt (R759), 11, 168, 172, 173, **173**
forall-triplet-spec (R755), 168, **168**, 169–172
FORM= specifier, 218, 248
format (R914), 222–224, **224**, 231, 257, 258
 format control, **260**
 format descriptor
 /, 274
 :, 274
 A, 270
 BN, 275
 BZ, 275
 control edit descriptor, 272
 D, 265
 data edit descriptor, 262–272
 E, 265
 EN, 266
 ES, 267
 F, 264
 G, 270
 I, 263
 L, 269
 P, 274
 S, 274
 SP, 274
 SS, 274
 TL, 273
 TR, 273
 X, 273
 format descriptors
 G, 270
 FORMAT statement, 224, 257, 626
format-item (R1003), 257, 258, **258**
format-specification (R1002), 257, **257**
format-stmt (R1001), 10, 257, **257**, 288, 292, 299
 formatted data transfer, 233
 formatted input/output statement, **223**
 formatted record, **205**
FORMATTED= specifier, 248
 formatting
 explicit, 257–276
 list-directed, 234, 276–280
 namelist, 234, 280–285
 forms, **206**

Fortran 2003 compatibility, 3
 FORTRAN 77 compatibility, 4
 Fortran 90 compatibility, 4
 Fortran 95 compatibility, 3
 Fortran character set, **25**
 free source form, 29, **29**
 function, **12, 510**
 function reference, **19**, 322
 function result, **510**
 FUNCTION statement, 326
 function subprogram, **326, 510**
function-name, 84, 299, 300, 326, 327, 331, 333, 485, 488
function-reference (R1219), 76, 84, 133, **308**, 310, 322
function-stmt (R1226), 9, 299, 300, 326, **326**, 327, 485, 488
function-subprogram (R1225), 9, 10, 288, **326**, 330

G

generic identifier, **302, 510**
 generic interface, 69, **302, 510**
 generic interface block, **300, 510**
 generic name, **302, 339**
 generic procedure reference, 304
generic-binding (R454), 68, **68**, 69, 290
generic-name, 68, 69, 299, 486, 488
generic-spec (R1207), 68, 69, 74, 99, 138, 139, 161, 162, 289, 290, 298, 299, **299**, 302, 486, 488
 global entity, **483, 510**
 global identifier, **483**
 GO TO statement, 192
goto-stmt (R851), 11, 35, 183, **192**, 193
 graphic character, **25**

H

hex-constant (R428), 53, 54, **54**
hex-digit (R429), 54, **54**, 269
hex-digit-string (R1021), 269, **269**
 host, **12, 13, 510**
 host association, **488, 510**
 host instance, **164**
 host scoping unit, **12, 510**

I

ICHAR intrinsic, 52
 ID= specifier, 226, 249
 IEEE_, 350, 439–465
 IF construct, 188, **188**
 IF statement, 188, **189**
if-construct (R840), 11, 188, **188**
if-construct-name, 188
if-stmt (R845), 11, 35, 189, **189**
if-then-stmt (R841), 188, **188**, 192

imag-part (R419), 48, **48**
 image, **14, 510**
 image control, **195**
 image index, **14, 511**
 image selector, **124**
image-selector (R624), 20, 117, **124**, 507
image-set (R862), 199, **199**, 200
image-team (R860), 197, **197**, 198, 200, 216, 219, 246, 252
 IMAGE_TEAM, **436**
 imaginary part, **48**
 implicit, **297**
 implicit interface, 308, **511**
 IMPLICIT statement, 105, **105**
 implicit team synchronization, **198**, 339
implicit-part (R205), 10, **10**
implicit-part-stmt (R206), 10, **10**
implicit-spec (R560), 105, **105**
implicit-stmt (R559), 10, **105**
 implied-shape array, **92**
implied-shape-spec (R518), 89, 92, **92**
import-name, 299
import-name-list, 300
import-stmt (R1209), 10, **299**
 inactive, **183**
INCLUDE, 32
 INCLUDE line, **32**
index-name, 168–174, 184, 486, 487, 501, 502
 inherit, **511**
 inheritance associated, **73**
 inheritance association, 498, **511**
 inherited, **72**
 initial point, **209**
initial-data-target (R448), 64, **64**, 84, 85, 101, 102
initial-proc-target (R1217), 65, 306, **306**, 307
 initialization, 65, 84, 85, 499, 607
initialization (R505), 84, **84**, 85
 initialization expression, **143**
 initialization target, **64, 307**
initialization-expr (R730), 43, 64, 65, 84, 85, 92, 96, 104, 144, **144**
inner-shared-do-construct (R837), 183, **183**
 input statement, **205**
input-item (R915), 222, 223, 227, **227**, 228, 241, 255, 503
 input/output editing, 257–285
 input/output list, 227
 input/output statements, 205–255
 INPUT_UNIT, 213, **436**
 inquire by file, **245**
 inquire by output list, **245**
 inquire by unit, **245**
 INQUIRE statement, 245
inquire-spec (R930), 246, **246**, 247, 255

- inquire-stmt* (R929), 11, **246**, 335
 - inquiry function, **339**, **511**
 - inquiry, type parameter, 119
 - instance, **329**, **511**
 - int-constant* (R308), 27, **27**, 101
 - int-constant-name*, 46
 - int-constant-subobject* (R541), 101, 102, **102**
 - int-expr* (R727), 16, 42, 79, 116, 120, 121, 124, 125, 140, **140**, 142–144, 146, 168, 178, 182, 184, 193, 199, 212, 213, 216, 222, 227, 230, 242, 246, 332
 - int-initialization-expr* (R732), 45, 50, 59, 78, 101, 144, **144**, 178, 193
 - int-literal-constant* (R407), 27, 46, **46**, 50, 258, 259
 - int-variable* (R608), 115, **115**, 124, 216, 221, 222, 242–244, 246, 249–254
 - int-variable-name*, 182
 - INTEGER**, 46
 - integer constant, 46
 - integer editing, 263
 - integer model, 341
 - integer type, **45**, 45
 - INTENT**, 340
 - intent, **511**
 - INTENT attribute, **93**, 103
 - INTENT statement, 103
 - intent-spec* (R523), 83, **93**, 103, 306
 - intent-stmt* (R545), 10, **103**
 - interface, **297**
 - (procedure), 297
 - abstract, **297**
 - explicit, 297
 - generic, 302
 - implicit, 308
 - interface block, **13**, 298–301, **511**
 - interface body, **13**, **511**
 - interface-block* (R1201), 10, **298**, 299
 - interface-body* (R1205), 298, 299, **299**, 488
 - interface-name* (R1215), 306, **306**
 - interface-name*, 68, 69
 - interface-specification* (R1202), 298, **298**
 - interface-stmt* (R1203), 298, **298**, 299, 302, 488
 - internal file, **212**, **511**
 - internal procedure, **13**, **295**, **511**
 - internal subprogram, **12**, **511**
 - internal unit, **212**
 - internal-file-variable* (R903), 212, **212**, 223, 503
 - internal-subprogram* (R211), 10, **10**
 - internal-subprogram-part* (R210), 9, **10**, 287, 326, 327, 329, 330
 - interoperable, 472, **472**, **475–477**, 478, **511**
 - interoperates, **472**
 - intrinsic, **22**, **511**
 - elemental, 339
 - inquiry function, 339
 - transformational, 339
 - intrinsic assignment statement, **157**
 - INTRINSIC attribute, **95**
 - intrinsic binary operation, **137**
 - intrinsic module, **288**
 - intrinsic operation, **137**
 - intrinsic operations, 150–154
 - intrinsic procedure, **295**, 339–434
 - INTRINSIC statement, **308**
 - intrinsic type, **17**, 45–53
 - intrinsic unary operation, **137**
 - intrinsic-operator* (R310), **27**, 28, 134, 136–139, 302
 - intrinsic-procedure-name*, 308, 488
 - intrinsic-stmt* (R1218), 10, **308**, 488
 - intrinsic-type-spec* (R404), 43, 44, **45**, 50
 - invoke, **511**
 - io-control-spec* (R913), 222, **222**, 223, 237, 255
 - io-implied-do* (R917), 227, **227**, 228, 229, 232, 255, 500, 502, 503, 544
 - io-implied-do-control* (R919), 227, **227**, 230
 - io-implied-do-object* (R918), 227, **227**
 - io-unit* (R901), 212, **212**, 222, 223, 335
 - iomsg-variable* (R907), 216, **216**, 221, 222, 242, 243, 245, 246, 253–255, 500
 - IOMSG= specifier**, 255
 - IOSTAT= specifier**, 254
 - IOSTAT_END, **437**
 - IOSTAT_INQUIRE_INTERNAL_UNIT, **437**
 - ISO 10646 character set, **49**
 - ISO 10646 character type, **49**, 158, 212, 217, 229, 262, 277, 409, 421
 - ISO_C_BINDING module, 467
 - ISO_FORTRAN_ENV module, 212, 213, 435
 - iteration count, **37**, **184**
- ## K
- k* (R1013), 259, **259**, 266, 271, 274
 - keyword, **21**, 310, **511**
 - keyword* (R215), 21, **21**, 74, 75, 309
 - kind, **45**, **46**, **48**, **49**, **53**
 - KIND intrinsic, 45, 46, 48, 49, 53
 - kind type parameter, **17**, 42, 45, 46, 48, 49, 53, **512**
 - kind-param* (R408), 46, **46**, 47, 48, 50, 51, 53, 54
 - kind-selector* (R405), 7, 45, **45**, 59
- ## L
- label, **512**
 - label* (R313), 28, **28**, 182, 192, 193, 216, 221, 222, 224, 242, 243, 245–247, 253, 254, 309
 - label-do-stmt* (R824), 182, **182**, 183
 - language-binding-spec* (R508), 83, **87**, 100, 326

- lbracket* (R469), 36, 60, 61, 79, **79**, 84, 89, 100, 103, 105, 124, 125
 - left tab limit, **273**
 - length, **49**
 - length of a character string, **512**
 - length type parameter, 42
 - length-selector* (R421), 7, 50, **50**
 - letter, **25**
 - letter*, 25, **25**, 26, 28, 105, 134, 136
 - letter-spec* (R561), 105, **105**, 106
 - level-1-expr* (R702), 134, **134**, 155
 - level-2-expr* (R706), 134, **134**, 135, 155
 - level-3-expr* (R710), 135, **135**
 - level-4-expr* (R712), 135, **135**, 136
 - level-5-expr* (R717), 136, **136**
 - lexical token, **26**, **512**
 - line, **29**, **512**
 - linkage association, **512**
 - list-directed formatting, 234, 276–280
 - list-directed input/output statement, **224**
 - literal constant, **18**, 116, **512**
 - literal-constant* (R306), 27, **27**
 - local entity, **512**
 - local identifier, **483**, 484
 - local variable, **18**, **512**
 - local-defined-operator* (R1114), 289, 290, **290**
 - local-name*, 289–291
 - LOGICAL**, 53
 - logical intrinsic assignment statement, **158**
 - logical intrinsic operation, **138**, 153
 - logical intrinsic operator, **138**
 - logical type, 53, **53**
 - logical-expr* (R724), 139, **139**, 144, 166, 178, 182, 184, 185, 188, 189
 - logical-initialization-expr* (R733), 144, **144**, 178
 - logical-literal-constant* (R424), 27, **53**, 134, 136
 - logical-variable* (R604), 115, **115**, 201
 - loop, **182**
 - loop-control* (R826), 182, **182**, 183, 184, 187
 - low-level syntax, **26**
 - lower-bound* (R513), 90, **90**, 91, 92
 - lower-bound-expr* (R634), 125, **125**, 128, 163
 - lower-co-bound* (R521), 93, **93**
- ## M
- m* (R1008), 258, 259, **259**, 263, 264, 269, 279
 - macro-actual-arg* (R339), 35, 36, **36**
 - macro-actual-arg-value* (R340), 36, **36**
 - macro-attribute* (R316), 33, **33**
 - macro-body-block* (R321), 33, **33**, 34
 - macro-body-construct* (R322), 33, **33**
 - macro-body-stmt* (R333), 33, 34, **34**
 - macro-condition* (R332), 34, **34**
 - macro-declaration-stmt* (R317), 33, **33**
 - macro-definition* (R314), 10, 33, **33**
 - macro-do-construct* (R323), 33, **33**
 - macro-do-limit* (R325), 33, 34, **34**
 - macro-do-stmt* (R324), 33, **33**, 37
 - macro-do-variable-name*, 33, 34, 37
 - macro-dummy-arg-name*, 33
 - macro-dummy-arg-name-list*, 33
 - macro-dummy-name*, 36
 - macro-else-if-stmt* (R329), 34, **34**
 - macro-else-stmt* (R330), 34, **34**
 - macro-end-do-stmt* (R326), 33, **34**
 - macro-end-if-stmt* (R331), 34, **34**
 - macro-expr* (R337), 33, 34, **34**, 37
 - macro-if-construct* (R327), 33, **34**
 - macro-if-then-stmt* (R328), 34, **34**
 - macro-local-variable-name*, 33
 - macro-local-variable-name-list*, 33
 - macro-name*, 33–35
 - macro-optional-decl-stmt* (R319), 33, **33**
 - macro-type-declaration-stmt* (R318), 33, **33**
 - macro-type-spec* (R320), 33, **33**, 36
 - main program, 12, **287**, **512**
 - main-program* (R1101), 9, 287, **287**
 - many-one array section, **123**, **512**
 - mask-expr* (R748), 166, **166**, 167–173
 - masked array assignment (WHERE), 166
 - masked-elsewhere-stmt* (R749), 166, **166**, 167, 172
 - model
 - bit, 340
 - integer, 341
 - real, 341
 - module, **13**, **288**, **512**
 - module* (R1104), 9, **288**
 - module procedure, **13**, **296**, **512**
 - module procedure interface, **300**, **512**
 - module procedure interface body, **300**
 - module reference, **22**, **289**
 - module subprogram, **12**, **512**
 - module-name*, 288–290, 488
 - module-nature* (R1110), 289, **289**, 290
 - module-stmt* (R1105), 9, 288, **288**
 - module-subprogram* (R1108), 10, 288, **288**, 331
 - module-subprogram-part* (R1107), 9, 69, 74, 288, **288**, 292, 559
 - mp-subprogram-stmt* (R1238), 330, **330**
 - mult-op* (R708), 27, 134, **134**
 - mult-operand* (R704), 134, **134**, 155
- ## N
- n* (R1015), 259, **259**
 - name, **21**, **512**
 - name* (R304), 6, 21, 26, **26**, 27, 84, 104, 115, 190, 232, 306, 326
 - name association, **487**, **512**

name-value subsequences, **280**
NAME= specifier, 249
 named, **512**
 named common block, **111**
 named constant, **18**, 96, 104, 116, **513**
 named file, **206**
named-constant (R307), 27, **27**, 32, 48, 78, 104, 488
named-constant-def (R548), 104, **104**, 488
NAMED= specifier, 249
 namelist formatting, 234, 280–285
 namelist input/output statement, **224**
 NAMELIST statement, 107, **107**
namelist-group-name, 107, 108, 222–224, 231, 233, 257, 280, 284, 291, 488, 503
namelist-group-object (R563), 107, **107**, 108, 232, 234, 241, 255, 280, 281, 291
namelist-stmt (R562), 10, **107**, 488, 503
 Names, **26**
 NaN, 350, 444, **513**
nested-token-sequence (R343), 36, **36**
 next record, **209**
NEXTREC= specifier, 249
NML= specifier, 224
 nonadvancing input/output statement, **210**
nonblock-do-construct (R831), 182, **183**
 nonexecutable statement, **14**, **513**
 nonintrinsic module, **288**
nonlabel-do-stmt (R825), 182, **182**
 normal, **444**
not-op (R718), 27, 136, **136**
notify-stmt (R863), 11, **200**
null-init (R506), 64, 84, **84**, 85, 101, 102, 306, 307
 NULL_IMAGE_TEAM, **437**
 NULLIFY statement, **129**
nullify-stmt (R638), 11, **129**, 503, 504
NUMBER= specifier, 249
 numeric conversion, 159
 numeric editing, 263
 numeric intrinsic assignment statement, **158**
 numeric intrinsic operation, **137**, 151
 numeric intrinsic operator, **137**
 numeric relational intrinsic operation, **138**
 numeric sequence type, **57**, 109, 497, 612
 numeric storage unit, **494**, **513**
 numeric type, **45**, **513**
numeric-expr (R728), 47, 140, **140**, 193
 NUMERIC_STORAGE_SIZE, **437**

O

object, *see* data object, **18**, **513**
 object designator, **21**, **513**
object-name (R504), 84, **84**, 100, 104, 105, 115, 488
 obsolescent feature, **7**, **513**
octal-constant (R427), 53, 54, **54**

only (R1112), 289, **289**, 290
only-use-name (R1113), 289, **289**, 290
 OPEN statement, 215, **215**
open-stmt (R904), 11, **215**, 335
OPENED= specifier, 249
 operand, **22**, **513**
 operation, **513**
 character intrinsic, 151
 defined, **138**
 extension, 139
 intrinsic, **137**
 intrinsic bits, 154
 logical intrinsic, 153
 numeric intrinsic, 151
 relational intrinsic, 152
 operations, 41
 operator, **22**, 27, **513**
 operator precedence, 154
 OPTIONAL attribute, **95**, 103
 optional dummy argument, 319
 OPTIONAL statement, 103
optional-stmt (R546), 10, **103**
or-op (R720), 27, 136, **136**
or-operand (R715), 136, **136**
outer-shared-do-construct (R835), 183, **183**
 output statement, **205**
output-item (R916), 222, 223, 227, **227**, 241, 246
 OUTPUT_UNIT, 213, **438**
 override, 73, **513**
 overrides, **65**

P

PAD= specifier, 218, 226, 250
 PARAMETER attribute, **96**, 104
 PARAMETER statement, 104, **104**
parameter-stmt (R547), 10, **104**, 488
 parent, **291**
 parent component, **73**, **513**
 parent data transfer statement, **234**, 234–238, 255
 parent type, **72**, **513**
parent-identifier (R1118), 291, 292, **292**
parent-string (R610), 88, 116, **116**
parent-submodule-name, 292
parent-type-name, 55, 56
 parentheses, 147
part-name, 117, 118, 120
part-ref (R613), 88, 101, 108, 117, **117**, 118, 120, 122, 316, 317
 partially [storage] associated, **496**
PASS attribute, 310
 passed-object dummy argument, **63**, **513**
PENDING= specifier, 250
 pointer, **20**, **513**
 pointer assignment, 162, **513**

- pointer assignment statement, **513**
 - pointer associated, **513**
 - pointer association, 491, **513**
 - pointer association context, 504
 - pointer association status, **491**
 - POINTER attribute, **96**, 104
 - POINTER statement, 104
 - pointer-assignment-stmt* (R735), 11, 97, **162**, 168, 171, 334, 503, 504
 - pointer-decl* (R550), 104, **104**
 - pointer-object* (R639), 129, **129**, 503, 504
 - pointer-stmt* (R549), 10, **104**, 488
 - polymorphic, **44**, **513**
 - POS= specifier**, 226, 250
 - position, **207**
 - position-edit-desc* (R1014), 259, **259**
 - position-spec* (R926), 243, **243**
 - POSITION= specifier**, 218, 250
 - positional arguments, 339
 - power-op* (R707), 27, 134, **134**
 - pre-existing, **498**
 - precedence of operators, 154
 - preceding record, **209**
 - preconnected, **513**
 - preconnected files, 215
 - Preconnection, **215**
 - prefix* (R1229), 326, **326**, 327, 329
 - prefix-spec* (R1230), 326, **326**, 327–329, 334, 336
 - present, **319**
 - present (dummy argument), 319
 - PRESENT intrinsic, 95
 - primaries*, 333
 - primary, 133
 - primary* (R701), 133, **133**, 134
 - PRINT statement, **221**
 - print-stmt* (R912), 11, **222**, 335
 - PRIVATE attribute, 56, 86
 - PRIVATE statement, 99
 - private-components-stmt* (R449), 56, 67, **67**
 - private-or-sequence* (R433), 55, 56, **56**
 - proc-attr-spec* (R1213), 306, **306**, 307
 - proc-binding-stmt* (R452), 68, **68**, 69
 - proc-component-attr-spec* (R446), 62, **62**
 - proc-component-def-stmt* (R445), 60, 62, **62**
 - proc-component-ref* (R741), 163, **163**, 309
 - proc-decl* (R1214), 62, 65, 306, **306**, 307
 - proc-entity-name*, 104, 307
 - proc-interface* (R1212), 62, 306, **306**, 307
 - proc-language-binding-spec* (R1227), 86, 306, 307, 326, **326**, 327, 329, 333, 477
 - proc-pointer-init* (R1216), 306, **306**
 - proc-pointer-name* (R554), 104, **104**, 111, 129, 163, 488
 - proc-pointer-object* (R740), 162, **163**, 164, 165, 171, 355, 503, 504
 - proc-target* (R742), 75, 77, 97, 162, **163**, 164, 165, 171, 319, 355, 494
 - procedure, **12**, **514**
 - characteristics of, 296
 - dummy, 296
 - elemental, 336
 - external, 295
 - internal, 295
 - intrinsic, 339–434
 - non-Fortran, 333
 - pure, 334
 - type-bound, **69**, 325
 - procedure designator, **21**, **514**
 - procedure interface, 297, **514**
 - procedure pointer, **20**, 306
 - procedure reference, **22**, 308
 - generic, 304
 - resolving, 323
 - type-bound, 325
 - procedure-component-name*, 163
 - procedure-declaration-stmt* (R1211), 10, 306, **306**, 307, 488
 - procedure-designator* (R1221), 308, 309, **309**, 325
 - procedure-entity-name*, 306
 - procedure-name*, 68, 69, 163, 165, 299, 306, 309, 330
 - procedure-stmt* (R1206), 298, 299, **299**
 - processor, **1**, **514**
 - processor dependent, **3**, **514**
 - program, **12**, **514**
 - program* (R201), 9, **9**
 - program unit, **12**, **514**
 - program-name*, 287
 - program-stmt* (R1102), 9, 287, **287**
 - program-unit* (R202), 7, 9, **9**
 - PROTECTED attribute, **96**, 104
 - PROTECTED statement, 104, **104**
 - protected-stmt* (R551), 10, **104**
 - prototype, **514**
 - PUBLIC attribute, 56, 86
 - PUBLIC statement, 99
 - pure procedure, **334**, **514**
- ## Q
- query-spec* (R865), 200, **200**
 - query-stmt* (R864), 11, **200**
- ## R
- r* (R1005), 258, **258**, 259, 260
 - range, **183**
 - RANGE intrinsic, 45
 - rank, **19**, **514**

- rbracket* (R470), 36, 60, 61, 79, **79**, 84, 89, 100, 103, 105, 124, 125
 - READ statement, **221**
 - read-stmt* (R910), 11, 222, **222**, 335, 503
 - READ= specifier**, 251
 - reading, **205**
 - READWRITE= specifier**, 251
 - REAL**, 47
 - real and complex editing, 264
 - real model, 341
 - real part, **48**
 - real type, **46**, 46–48
 - real-literal-constant* (R413), 27, 47, **47**
 - real-part* (R418), 48, **48**
 - REC= specifier**, 226
 - RECL= specifier**, 218, 251
 - record, **205**, **514**
 - record file, **205**
 - record length, **206**
 - record number, **208**
 - RECURSIVE**, 327, 329
 - recursive input/output statement, **255**
 - reference, **514**
 - rel-op* (R713), 27, 135, **135**, 152
 - relational intrinsic operation, **138**, 152
 - relational intrinsic operator, **138**
 - rename* (R1111), 289, **289**, 290, 484
 - rep-char*, 51, **51**, 260, 277, 282
 - repeat specification., **258**
 - representable character, **51**
 - representation method, **45**, **46**, **49**, **53**
 - resolving procedure reference, 323
 - resolving procedure references
 - derived-type input/output, 240
 - restricted expression, **142**
 - result variable, **12**, **514**
 - result-name*, 326, 327, 331, 488
 - result-token* (R334), 34, **34**
 - RETURN statement, 332
 - return-stmt* (R1241), 11, 16, 35, 183, 332, **332**
 - REWIND statement, 244
 - rewind-stmt* (R925), 11, **243**, 335
 - round-edit-desc* (R1018), 259, **260**
 - ROUND= specifier**, 219, 226, 251
 - rounding mode, **514**
- S**
- satisfied, **200**
 - SAVE attribute, 101, 104
 - SAVE statement, 104
 - save-stmt* (R552), 10, **104**, 488
 - saved, **97**
 - saved-entity* (R553), 104, **104**
 - scalar, **19**, **116**, **514**
 - scalar-xyz* (R103), 6, **6**
 - scale factor, **259**
 - scope, **483**, **514**
 - scoping unit, **12**, **514**
 - section subscript, **515**
 - section-subscript* (R620), 117, 120, **121**, 122, 316
 - segment, **195**
 - SELECT CASE statement, 178
 - SELECT TYPE, **189**
 - SELECT TYPE construct, 189, 487
 - SELECT TYPE construct**, 490
 - select-case-stmt* (R811), 178, **178**, 192
 - select-construct-name*, 189, 190
 - select-type-construct* (R846), 11, **189**, 190
 - select-type-stmt* (R847), 189, **189**, 190, 192
 - SELECTED_INT_KIND intrinsic, 45
 - selector, **515**
 - selector* (R805), 176, **176**, 177, 189–191, 319, 491, 503
 - separate module procedure, **330**
 - separate-module-subprogram* (R1237), 10, 288, 330, **330**
 - sequence, **22**
 - sequence association, 318
 - SEQUENCE property, 58
 - SEQUENCE statement, **57**
 - sequence structure, **55**
 - sequence type, **57**
 - sequence-stmt* (R435), 56, **57**
 - sequential access, 207
 - sequential access input/output statement, **226**
 - SEQUENTIAL= specifier**, 251
 - shape, **19**, **515**
 - shape conformance, **141**
 - shared-term-do-construct* (R836), 183, **183**
 - sign* (R411), 46, **46**, 47, 264
 - sign-edit-desc* (R1016), 259, **259**
 - SIGN= specifier**, 219, 227, 251
 - signed-digit-string* (R409), **46**, 47, 263, 264
 - signed-int-literal-constant* (R406), 46, **46**, 48, 101, 259
 - signed-real-literal-constant* (R412), **47**, 48, 101
 - significand* (R414), 47, **47**
 - size, **19**, **515**
 - size of a common block, **112**
 - size of a storage sequence, **494**
 - SIZE= specifier**, 227, 251
 - source forms, **29**
 - source-expr* (R630), 124, **124**, 125–127
 - special characters, **26**
 - special-character*, 25, **26**
 - specific interface, **299**
 - specific interface block, **300**
 - specific name, **339**

- specific-binding* (R453), 68, **68**
- specification, 83–113
- specification expression, **142, 515**
- specification function, **143, 515**
- specification inquiry, **142**
- specification-expr* (R729), 43, 50, 84, 90, 93, 142, **142, 336**
- specification-part* (R204), 9, **10**, 16, 68, 86, 99, 143, 144, 177, 178, 287, 288, 292, 299, 326, 329, 330, 334
- specification-stmt* (R212), 10, **10**
- standard-conforming program, **2, 515**
- starting point, **116**
- stat-variable* (R628), 124, **124**, 126, 127, 129, 197, 503
- statement, **29, 515**
 - accessibility, 99
 - ALLOCATABLE, 100
 - ALLOCATE, 124
 - arithmetic IF, 193
 - assignment, 156
 - ASSOCIATE, 176
 - ASYNCHRONOUS, 100
 - attribute specification, 99–113
 - BACKSPACE, 243
 - BIND, 100
 - CALL, 308
 - CASE, 178
 - CLASS IS, 189
 - CLOSE, 220
 - COMMON, 111–113
 - computed GO TO, 193
 - CONTAINS, 333
 - CONTIGUOUS, 100
 - CONTINUE, 193
 - CYCLE, 185
 - DATA, 100
 - data transfer, 221
 - DEALLOCATE, 129
 - DIMENSION, 103
 - DO, 182
 - DO WHILE, 182
 - END, 16
 - ENDFILE, 244
 - ENTRY, 331
 - EQUIVALENCE, 108–110
 - executable, 14
 - EXTERNAL, 305
 - file inquiry, 245
 - file positioning, 243
 - FLUSH, 244
 - FORALL, 168, 173
 - FORMAT, 257
 - formatted input/output, 223
 - FUNCTION, 326
 - GO TO, 192
 - IF, 189
 - IMPLICIT, 105
 - input/output, 205–255
 - INTENT, 103
 - INTRINSIC, 308
 - list-directed input/output, 224
 - NAMELIST, 107
 - namelist input/output, 224
 - NULLIFY, 129
 - OPEN, 215
 - OPTIONAL, 103
 - PARAMETER, 104
 - POINTER, 104
 - pointer assignment, 162
 - PRINT, 221
 - PRIVATE, 99
 - PROTECTED, 104
 - PUBLIC, 99
 - READ, 221
 - RETURN, 332
 - REWIND, 244
 - SAVE, 104
 - SELECT CASE, 178
 - SELECT TYPE, 189
 - STOP, 193
 - TARGET, 105
 - TYPE, 55
 - type declaration, 83–85
 - TYPE IS, 189
 - unformatted input/output, 223
 - VALUE, 105
 - VOLATILE, 105
 - WHERE, 166
 - WRITE, 221
- statement entity, **483, 515**
- statement function, **296, 333, 515**
- statement label, 28, **28, 515**
- statement order, 14
- statements
 - INQUIRE, 245
- STATUS= specifier**, 219, 221
- stmt-function-stmt* (R1243), 10, 288, 292, 299, **333, 488**
- STOP statement, 193
- stop-code* (R856), 193, **193, 194**
- stop-stmt* (R855), 11, 35, 183, **193, 335**
- storage associated, **496**
- storage association, 108–113, 494, **494, 515**
- storage sequence, 112, **494, 515**
- storage unit, **494, 515**
- stream access, 208
- stream access input/output statement, **226**

- stream file, **205**
- STREAM= specifier**, 252
- stride, **122, 515**
- stride* (R622), 121, **121**, 168, 169, 171, 172, 230, 316
- string, *see* character string
- struct, **515**
- structure, **17, 55, 515**
- structure component, **117, 515**
- structure constructor, **75, 515**
- structure-component* (R614), 101, 115, 116, **118**, 125, 129
- structure-constructor* (R459), **75**, 76, 101, 102, 133, 335
- subcomponent, **118, 515**
- submodule, **13, 291, 516**
- submodule* (R1116), 9, **292**
- submodule identifier, **291**
- submodule-name*, 292
- submodule-stmt* (R1117), 9, 291, 292, **292**
- subobject, **18, 516**
- subprogram, **12, 516**
- subroutine, **12, 516**
- subroutine reference, 322
- subroutine subprogram, **329, 516**
- subroutine-name*, 299, 300, 329, 485
- subroutine-stmt* (R1234), 9, 299, 300, 326, 329, **329**, 485, 488
- subroutine-subprogram* (R1233), 9, 10, 288, **329**, 330
- subscript, **516**
- subscript* (R619), 7, 101, 120, **120**, 121, 168, 169, 172, 230, 316
- subscript triplet, 122, **516**
- subscript-triplet* (R621), 121, **121**, 122, 316
- substring, **116, 516**
- substring* (R609), 108, 109, 115, **116**
- substring-range* (R611), 88, 116, **116**, 118, 120, 121, 230, 316
- suffix* (R1231), 326, **327**, 331
- sync-all-stmt* (R857), 11, **197**
- sync-images-stmt* (R861), 11, **199**
- sync-memory-stmt* (R866), 11, **202**
- sync-stat* (R858), 197, **197**, 199, 200, 202
- sync-team-stmt* (R859), 11, **197**
- synchronization
 - implicit team, 339
- synchronous, **225**, 227, 229
- T**
- target, **20, 516**
- TARGET attribute, **98**, 105
- TARGET statement, 105
- target-decl* (R556), 105, **105**
- target-stmt* (R555), 10, **105**, 488
- team, **14, 516**
- team synchronization, **197, 516**
 - implicit, 198
- TEAM= specifier**, 252
- terminal point, **209**
- TKR compatible, **304**
- token* (R335), 34, **34**
- totally [storage] associated, **496**
- transfer of control, 175, 192, 253, 254
- transformational function, **339, 516**
- type, **17**, 41–80, **516**
 - abstract, **72**
 - base, 72
 - character, 49–53
 - complex, 48
 - concept, **41**
 - declared, 44
 - derived, 54–77
 - dynamic, 44
 - expression, 139
 - extended, 72
 - extensible, 72
 - extension, 72
 - integer, 45
 - intrinsic, 45–53
 - logical, 53
 - operation, 140
 - parent, 72
 - primary, 140
 - real, 46–48
- type compatible, **44, 516**
- type conformance, 157
- type declaration statement, 83–85, **516**
- type equality, 57
- type incompatible, **44**
- type parameter, **17**, 42, 45, 46, 485, **516**
- type parameter inquiry, **119**
- type parameter keyword, **21**
- type parameter order, **60, 516**
- type specifier, **43**
 - derived type, 44
 - TYPE, 44
- TYPE statement, 55
- TYPE type specifier, 44
- type-attr-spec* (R432), 55, **55**
- type-bound procedure, **69, 325, 516**
- type-bound-procedure-part* (R450), 55, 57, **68**, 70, 475
- type-declaration-stmt* (R501), 10, 50, 83, **83**, 334, 488
- type-guard-stmt* (R848), 189, **189**, 190
- type-name*, 55, 56, 59, 68, 74, 516
- type-param-attr-spec* (R438), 59, **59**

type-param-decl (R437), 59, **59**
type-param-def-stmt (R436), 55, 59, **59**
type-param-inquiry (R616), **119**, 120, 133, 486
type-param-name, 55, 59, 62, 119, 133, 134, 486, 488
type-param-spec (R458), 74, **74**, 75
type-param-value (R401), 42, **42**, 43, 50, 61, 62, 74, 75, 125, 126
type-spec (R402), 43, **43**, 44, 50, 51, 79, 80, 124–126, 189, 190

U

ultimate component, **55**, **517**
 unallocated, **127**
 undefined, **22**, **517**
 undefinition of variables, 498
underscore (R303), 25, **25**
 unformatted data transfer, 232
 unformatted input/output statement, **223**
 unformatted record, **206**
UNFORMATTED= specifier, 252
 Unicode, **217**
 unit, **212**
 unlimited polymorphic, **44**
unlimited-format-item (R1004), 257, **258**, 261
 unsaved, **97**
 unsigned, **517**
 unspecified storage unit, **494**, **517**
upper-bound (R514), 90, **90**
upper-bound-expr (R635), 125, **125**, 128, 163
upper-co-bound (R522), 93, **93**
 use associated, **289**
 use association, 488, **517**
 USE statement, **289**
use-defined-operator (R1115), 289, 290, **290**
use-name, 99, 289, 290, 484
use-stmt (R1109), 10, **289**, 290, 488

V

v (R1011), 237, 259, **259**, 272
 value, **171**
 VALUE attribute, **98**, 105
 value separator, **276**
 VALUE statement, 105
value-stmt (R557), 11, **105**
 variable, **18**, **517**
variable (R601), xv, 76, 101, 115, **115**, 157–160, 162, 163, 167, 169, 171, 176, 189, 190, 227, 309, 503
variable-name (R602), 107, 108, 111, 115, **115**, 116, 125, 129, 162, 163, 488, 503
 variables
 definition & undefinition, 498
 vector subscript, **123**, **517**

vector-subscript (R623), 121, **121**, 122
 void, **517**
 VOLATILE attribute, **98**, 105
 VOLATILE statement, 105
volatile-stmt (R558), 11, **105**

W

w (R1007), 258, 259, **259**, 263–267, 269–272, 277, 279, 282
 wait operation, 220, 229, 241, **241**
WAIT statement, 242
wait-spec (R922), 242, **242**
wait-stmt (R921), 11, **242**, 335
 WHERE construct, 166
 WHERE statement, 166
where-assignment-stmt (R747), 166, **166**, 167, 168, 172, 508
where-body-construct (R746), 166, **166**, 167
where-construct (R744), 11, 166, **166**, 168, 172
where-construct-name, 166
where-construct-stmt (R745), 166, **166**, 172, 192
where-stmt (R743), 11, 166, **166**, 168, 172
 whole array, **120**, **517**
 WRITE statement, **221**
write-stmt (R911), 11, 222, **222**, 335, 503
WRITE= specifier, 252
 writing, **205**

X

xyz-list (R101), **6**
xyz-name (R102), **6**

Z

zero-size array, 19, 90, 102
 ZZZUTI005, 194
 ZZZUTI006, 194
 ZZZUTI007, 20
 ZZZUTI010, 288
 ZZZUTI011, 88
 ZZZUTI016, 311
 ZZZUTI020, 117
 ZZZUTI023, 157
 ZZZUTI033, 195
 ZZZUTI039, 214
 ZZZUTI042, 220
 ZZZUTI043, 220
 ZZZUTI044, 249
 ZZZUTI049, 298
 ZZZUTI050, 312
 ZZZUTI051, 312
 ZZZUTI054, 317
 ZZZUTI055, 314
 ZZZUTI061, 360
 ZZZUTI063, 380

ZZZUTI073, 436
ZZZUTI075, 494
ZZZUTI076, 492
ZZZUTI077, 495
ZZZUTI079, 317
ZZZUTI080, xv
ZZZUTI081, 127
ZZZUTI082, 567
ZZZUTI083, 332
ZZZUTI084, 178
ZZZUTI085, 429

ZZZUTI086, 198
ZZZUTI087, 437
ZZZUTI088, 198
ZZZUTI089, 119
ZZZUTI090, 436
ZZZUTI091, 230
ZZZUTI092, 419
ZZZUTI093, 422
ZZZUTI094, 61
ZZZUTI095, 14
ZZZUTI5003, 512