

Fortran Language Reference Manual,
Volume 3

Document Number 007-3694-004

St. Peter's Basilica image courtesy of ENEL SpA and InfoByte SpA. Disk Thrower image courtesy of Xavier Berenguer, Animatica.

Copyright © 1993, 1999 Silicon Graphics, Inc. All Rights Reserved. This manual or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Silicon Graphics, Inc.

The CF90 compiler includes United States software patents 5,247,696, 5,257,372, and 5,361,354.

LIMITED AND RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in the Rights in Data clause at FAR 52.227-14 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94043-1389.

Autotasking, CF77, CRAY, Cray Ada, CraySoft, CRAY Y-MP, CRAY-1, CRInform, CRI/*TurboKiva*, HSX, LibSci, MPP Apprentice, SSD, SUPERCLUSTER, UNICOS, and X-MP EA are federally registered trademarks and Because no workstation is an island, CCI, CCMT, CF90, CFT, CFT2, CFT77, ConCurrent Maintenance Tools, COS, Cray Animation Theater, CRAY APP, CRAY C90, CRAY C90D, Cray C++ Compiling System, CrayDoc, CRAY EL, CRAY J90, CRAY J90se, CrayLink, Cray NQS, Cray/REELibrarian, CRAY S-MP, CRAY SSD-T90, CRAY SV1, CRAY T90, CRAY T3D, CRAY T3E, CrayTutor, CRAY X-MP, CRAY XMS, CRAY-2, CSIM, CVT, Delivering the power . . ., DGauss, Docview, EMDS, GigaRing, HEXAR, IOS, ND Series Network Disk Array, Network Queuing Environment, Network Queuing Tools, OLNETH, RQS, SEGLDR, SMARTE, SUPERLINK, System Maintenance and Remote Testing Environment, Trusted UNICOS, UNICOS MAX, and UNICOS/mk are trademarks of Cray Research, Inc., a wholly owned subsidiary of Silicon Graphics, Inc.

IRIS, IRIX, and Silicon Graphics are registered trademarks and the Silicon Graphics logo is a trademark of Silicon Graphics, Inc.

MIPS is a registered trademark and MIPSpro is a trademark of MIPS Technologies, Inc. TotalView is a trademark of Bolt Baranek and Newman Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively to X/Open Limited. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited. X/Open is a registered trademark of X/Open Company Ltd. The X device is a trademark of The Open Group,

Adapted with permission of McGraw-Hill, Inc. from the FORTRAN 90 HANDBOOK, Copyright © 1992 by Walter S. Brainerd, Jeanne C. Adams, Jeanne T. Martin, Brian T. Smith, and Jerrold L. Wagener. All rights reserved. Cray Research, Inc. is solely responsible for the content of this work.

The UNICOS operating system is derived from UNIX® System V. The UNICOS operating system is also based in part on the Fourth Berkeley Software Distribution (BSD) under license from The Regents of the University of California.

New Features

This manual describes the Fortran 95 language as implemented by the CF90 compiler, release 3.2, and the MIPSpro 7 Fortran 90 compiler, revision 7.3. These compilers supported parts of the Fortran 95 standard in previous releases. The major changes to this manual reflect that the CF90 and the MIPSpro 7 Fortran 90 compilers now support the entire Fortran 95 standard.

This revision contains the following new features to support the Fortran 95 standard:

- The `FORALL` statement and construct.
- Pointer initialization.
- Default initialization in a derived type.
- A distinction between positive zero and negative real zero (-0.0) by the `SIGN(3I)` intrinsic function and input/output.
- The `ELSEWHERE` statement can accept a conditional expression as an argument.
- `MODULE PROCEDURE` statements can be intermixed with interface bodies. The `END INTERFACE` statement can repeat the generic identifier.
- User functions can be called from specification expressions.
- The `NULL` and `CPU_TIME` intrinsic procedures.
- The `CEILING` and `FLOOR` intrinsic procedures now accept `[KIND=]kind` arguments.

This revision also contains the following feature as an extension to the Fortran 95 standard:

- The `FORM=SYSTEM` specifier on the `OPEN` statement

The Fortran 95 standard has declared some Fortran 90 features to be obsolescent or has deleted them. The CF90 and MIPSpro 7 Fortran 90 compilers, however, have not removed their functionality from the system. Obsolescent features are described in the *Fortran Language Reference Manual, Volume 1* and in the *Fortran Language Reference Manual, Volume 2*; both manuals note when a particular feature is categorized as obsolescent. All features deleted from the current Fortran standard are described in the *Fortran Language Reference Manual, Volume 3*. When the compilers detect a deleted feature, they process the feature and issue a `NON-ANSI` message. The deleted features that generate `NON-ANSI` messages are as follows:

- Real and double precision `DO` variables
- Branching to an `END IF` from outside its `IF` construct
- `PAUSE` statement

- `ASSIGN`, assigned `GO TO` statements, and assigned format specifiers
- `H` edit descriptor

This revision describes the following feature, which is implemented on IRIX platforms only:

- The `VOLATILE` statement and attribute

Record of Revision

<i>Version</i>	<i>Description</i>
2.0	November 1995 Original Printing. The sections in this manual previously appeared in the <i>CF90 Fortran Language Reference Manual</i> , revision 1.0, publication SR-3902, and the <i>CF90 Commands and Directives Reference Manual</i> .
3.0	May 1997 This printing supports the Cray Research CF90 3.0 release running on UNICOS and UNICOS/mk operating systems, and the MIPSpro 7 Fortran 90 compiler release 7.2 running on the IRIX operating system. The implementation of features on IRIX operating system platforms is deferred.
3.0.1	August 1997 This online revision supports the Cray Research CF90 3.0.1 release, running on UNICOS and UNICOS/mk operating systems, and the MIPSpro 7 Fortran 90 compiler 7.2 release, running on the IRIX operating system. Includes minor updates and corrections to revision 3.0.
3.0.2	March 1998 This online revision supports the Cray Research CF90 3.0.2 release, running on UNICOS and UNICOS/mk operating systems, and the MIPSpro 7 Fortran 90 compiler 7.2.1 release, running on the IRIX operating system. Includes minor updates and corrections to revision 3.0.1.
3.1	August 1998 This online revision supports the Cray Research CF90 3.1 release, running on UNICOS and UNICOS/mk operating systems, and the MIPSpro 7 Fortran 90 compiler 7.2.1 release, running on the IRIX operating system. Includes minor updates and corrections to revision 3.0.2.
3.2	January 1999 This revision (007-3694-004) supports the CF90 3.2 release, running on the UNICOS and UNICOS/mk operating systems, and the MIPSpro 7 Fortran 90 7.3 release, running on the IRIX operating system. It includes major updates to revision 3.1.

Contents

	<i>Page</i>
About This Manual	xiii
Related CF90 and MIPSpro 7 Fortran 90 Compiler Publications	xiii
CF90 and MIPSpro 7 Fortran 90 Compiler Messages	xiv
CF90 and MIPSpro 7 Fortran 90 Compiler Man Pages	xiv
Related Fortran Publications	xv
Related Publications	xv
Obtaining Publications	xvi
Conventions	xvii
BNF Conventions	xix
Reader Comments	xx
Fortran Syntax [1]	1
Syntax Form	1
Syntax Rules Expressed in BNF	1
Definition Syntax Symbol: Is	2
Alternative Syntax Symbol: Or	2
Optional Symbol: []	3
Symbol for Repeated Items: [].	3
Syntax Rule Continuation	3
Assumed Syntax Rules	4
Example BNF Syntax	4
Constraints	5
Identifying Numbers	5
Syntax Rules and Constraints	5
Introduction	6
007-3694-004	iii

	<i>Page</i>
Fortran Terms and Concepts	6
Characters, Lexical Tokens, and Source Form	10
Intrinsic and Derived Data Types	13
Data Object Declarations and Specifications	18
Use of Data Objects	31
Expressions and Assignment	35
Execution Control	42
Input/Output (I/O) Statements	48
I/O Editing	54
Program Units	57
Procedures	59
Intrinsic Procedures	66
Scope, Association, and Definition	67
Decremental Features [2]	69
Deleted Features	69
PAUSE Statement	69
ASSIGN, Assigned GO TO Statements, and Assigned Format Specifiers	70
Form of the ASSIGN and Assigned GO TO Statements	71
Assigned Format Specifiers	72
H Edit Descriptor	72
Obsolescent Features	73
Character Set [3]	77
Extensions and Differences [4]	83
Fortran 95 Standard Differences and Incompatibilities With FORTRAN 77 Implementations	83
Fortran 95 and the G Edit Descriptor Output Differences	83
Fortran 95 and List-directed Output Differences	84
Delimited and Undelimited Character Strings in List-directed I/O	84

	<i>Page</i>
List-directed I/O and Floating-point Zero	84
CF90 extensions to Fortran 95	85
List-directed I/O and Hollerith Constants	85
Differences in the B, O, and Z Edit Descriptors	85
MIPSpro 7 Fortran 90 extensions to Fortran 95	86
CF90, MIPSpro 7 Fortran 90, and MIPSpro Fortran 77 Differences	86
MIPSpro 7 Fortran 90 and CF90 Compiler Differences	87
Numerical Model Differences	87
Fortran Statement Differences	87
Function and Procedure Differences	88
Modules Differences	88
I/O Library Differences	88
Library Function and Procedure Differences	89
Math Library Differences	89
MIPSpro FORTRAN 77 and MIPSpro 7 Fortran 90 Compiler Differences	89
Intrinsic Function and Subroutine Differences	89
DATA Statement Initialization Differences	90
I/O Record Length Differences	90
Special File Formats Differences	90
Miscellaneous Differences	91
Data Representation and Storage [5]	95
Data Representation for UNICOS Systems	95
Integer Type	95
Real Type	96
Normalized Floating-point Numbers	98
Double-precision Type	99
Single-precision Complex Type	99
Double-precision Complex Type	100

	<i>Page</i>
Character Type	101
Logical Type	102
Cray Character Pointers	102
Data Representation for IRIX systems	103
Integer Type	103
KIND=1	103
KIND=2	104
KIND=4	104
KIND=8	104
Real Type	105
KIND=4	105
KIND=8	107
KIND=16	108
Complex Type	109
KIND=4	109
KIND=8	110
KIND=16	110
Character Type	112
Logical Type	112
Cray Character Pointers (Deferred Implementation)	112
Data Representation for UNICOS/mk Systems	113
Integer Type	113
KIND=1, KIND=2, or KIND=4	113
KIND=8	114
Real Type	114
KIND=4	114
KIND=8	116
Complex Type	117
KIND=4	117

	<i>Page</i>
KIND=8	118
Character Type	119
Logical Type	119
Cray Character Pointers	120
Data Representation for CRAY T90 Systems That Support IEEE Floating-point Arithmetic	120
Integer Type	120
KIND=1, KIND=2, or KIND=4	120
KIND=8	121
Real Type	121
KIND=4 and KIND=8	122
KIND=16	122
Complex Type	123
KIND=4 and KIND=8	123
KIND=16	124
Character Type	126
Logical Type	126
Cray Character Pointers	127
Storage Issues	127
Storage Units and Sequences	128
Static and Stack Storage	130
Dynamic Memory Allocation (UNICOS Systems Only)	132
Changing Your Code: Standard Method	133
Changing Your Code: Nonstandard Method	134
Outmoded Features [6]	137
Hollerith Type	138
Hollerith Constants	138
Hollerith Values	140
Hollerith Relational Expressions	140

	<i>Page</i>
Formatted I/O and Internal Files	141
ENCODE Statement	141
DECODE Statement	143
Edit Descriptors	145
Asterisk Delimiters	145
Negative-valued X Descriptor	145
A and R Descriptors for Noncharacter Types	146
Type Declaration with Data Length	147
DATA Statement Features	150
IF Statements	150
Two-branch Arithmetic IF	150
Indirect Logical IF	151
TASK COMMON Statement (UNICOS Systems Only)	151
Nested Loop Termination	152
DOUBLE COMPLEX Statement (UNICOS Systems Only)	152
Bitwise Logical Expressions	153
CF90 Defined Externals [7]	157
Conformance Checks	157
Record Length	157
Glossary	159
Index	169
Figures	
Figure 1. Default 64-bit integers	96
Figure 2. Fast integer operations with INTEGER(KIND=8), UNICOS systems (except CRAY T90 systems that support IEEE floating-point arithmetic)	96
Figure 3. Real type	96

	<i>Page</i>
Figure 4. Binary version of 10.0	98
Figure 5. Double-precision type	99
Figure 6. Single-precision complex type	100
Figure 7. Double-precision complex type (real portion)	101
Figure 8. Double-precision complex type (imaginary portion)	101
Figure 9. Character type	101
Figure 10. 64-bit addressing for UNICOS systems (except CRAY T90 systems)	102
Figure 11. 64-bit addressing for CRAY T90 systems	103
Figure 12. INTEGER(KIND=1) on IRIX systems	103
Figure 13. INTEGER(KIND=2) on IRIX systems	104
Figure 14. INTEGER(KIND=4) on IRIX systems	104
Figure 15. INTEGER(KIND=8) on IRIX systems	105
Figure 16. REAL(KIND=4) on IRIX systems	105
Figure 17. Binary version of 10.0	107
Figure 18. REAL(KIND=8) on IRIX systems	107
Figure 19. REAL(KIND=16) on IRIX systems	108
Figure 20. COMPLEX(KIND=4) on IRIX systems (real portion)	109
Figure 21. COMPLEX(KIND=4) on IRIX systems (imaginary portion)	109
Figure 22. COMPLEX(KIND=8) on IRIX systems (real portion)	110
Figure 23. COMPLEX(KIND=8) on IRIX systems (imaginary portion)	110
Figure 24. COMPLEX(KIND=16) on IRIX systems (real portion)	111
Figure 25. COMPLEX(KIND=16) on IRIX systems (imaginary portion)	111
Figure 26. Character type	112
Figure 27. 32-bit addressing on IRIX systems	113
Figure 28. Integer KIND=1, 2, or 4 on UNICOS/mk systems	113
Figure 29. INTEGER(KIND=8) on UNICOS/mk systems	114
Figure 30. REAL(KIND=4) on UNICOS/mk systems	114

	<i>Page</i>
Figure 31. Binary version of 10.0	116
Figure 32. REAL(KIND=8) on UNICOS/mk systems	116
Figure 33. COMPLEX(KIND=4) on UNICOS/mk systems (real portion)	117
Figure 34. COMPLEX(KIND=4) on UNICOS/mk systems (imaginary portion)	117
Figure 35. COMPLEX(KIND=8) on UNICOS/mk systems (real portion)	118
Figure 36. COMPLEX(KIND=8) on UNICOS/mk systems (imaginary portion)	118
Figure 37. Character type	119
Figure 38. Cray character pointers on UNICOS/mk systems	120
Figure 39. Integer KIND=1, 2, or 4 on CRAY T90 systems that support IEEE floating-point arithmetic	120
Figure 40. Default INTEGER(KIND=8) on CRAY T90 systems that support IEEE floating-point arithmetic	121
Figure 41. Fast operations with INTEGER(KIND=8) on CRAY T90 systems that support IEEE floating-point arithmetic	121
Figure 42. Real KIND=4 or 8 on CRAY T90 systems that support IEEE floating-point arithmetic	122
Figure 43. REAL(KIND=16) on CRAY T90 systems that support IEEE floating-point arithmetic	123
Figure 44. Complex KIND=8 or 4 on CRAY T90 systems that support IEEE floating-point arithmetic (real portion)	124
Figure 45. Complex KIND=8 or 4 on CRAY T90 systems that support IEEE floating-point arithmetic (imaginary portion)	124
Figure 46. COMPLEX(KIND=16) on CRAY T90 systems that support IEEE floating-point arithmetic (real portion)	125
Figure 47. COMPLEX(KIND=16) on CRAY T90 systems that support IEEE floating-point arithmetic (imaginary portion)	125
Figure 48. Character type	126
Figure 49. Cray character pointer for CRAY T90 systems that support IEEE floating-point arithmetic	127
Figure 50. Memory use under UNICOS	133

Tables

Table 1. Syntax metalanguage abbreviations	2
Table 2. Summary of string edit descriptors	73

	<i>Page</i>
Table 3. Character set	77
Table 4. Outmoded features and preferred alternatives	137
Table 5. Data length (UNICOS systems)	148
Table 6. Data length (UNICOS/mk systems)	149
Table 7. Data length (IRIX systems)	149
Table 8. Standard alternatives to CF90 double-complex functions	153
Table 9. Standard alternatives to CF90 and MIPSpro 7 Fortran 90 bitwise functions	154
Table 10. Data types in bitwise logical operations	155

About This Manual

This manual describes the Fortran language as implemented by the CF90 compiler, revision 3.2, and by the MIPSpro 7 Fortran 90 compiler, revision 7.3. The CF90 and MIPSpro 7 Fortran 90 compilers implement the Fortran standard.

The CF90 and MIPSpro 7 Fortran 90 compilers run on UNICOS, UNICOS/mk, and IRIX operating systems. Specific hardware and operating system support information is as follows:

- The CF90 compiler runs under UNICOS 10.0, or later, on CRAY SV1, CRAY C90, CRAY J90, and CRAY T90 systems.
- The CF90 compiler runs under UNICOS/mk 2.0.3, or later, on CRAY T3E systems.
- The MIPSpro 7 Fortran 90 compiler runs under IRIX 6.2, or later, on Cray and Silicon Graphics systems.

The CF90 and MIPSpro 7 Fortran 90 compilers were developed to support the Fortran standard adopted by the American National Standards Institute (ANSI) and the International Standards Organization (ISO). This standard, commonly referred to in this manual as *the Fortran standard*, is ISO/IEC 1539-1:1997. Because the Fortran standard is, generally, a superset of previous standards, the CF90 and MIPSpro 7 Fortran 90 compilers will compile code written to previous standards.

Note: The Fortran 95 standard is a revision to the Fortran 90 standard. The standards organizations continue to interpret the Fortran standard for Silicon Graphics and for other vendors. To maintain conformance to the Fortran standard, Silicon Graphics may need to change the behavior of certain CF90 and MIPSpro 7 Fortran 90 compiler features in future releases based upon the outcomes of interpretations to the standard.

Related CF90 and MIPSpro 7 Fortran 90 Compiler Publications

This manual is one of a set of manuals that describes the CF90 and the MIPSpro 7 Fortran 90 compilers. The complete set of CF90 and MIPSpro 7 Fortran 90 compiler manuals is as follows:

- *Intrinsic Procedures Reference Manual*.

- *Fortran Language Reference Manual, Volume 1*. Chapters 1 through 8 correspond to sections 1 through 8 of the Fortran standard.
- *Fortran Language Reference Manual, Volume 2*. Chapters 1 through 6 of this manual correspond to sections 9 through 14 of the Fortran standard.
- *Fortran Language Reference Manual, Volume 3*. This manual contains CF90 and MIPSpro 7 Fortran 90 compiler information that supplements the Fortran standard. The standard is the complete, official description of the language. This manual also contains the complete Fortran syntax in Backus-Naur form (BNF).

The following publications contain information specific to the CF90 compiler:

- *CF90 Ready Reference*
- *CF90 Commands and Directives Reference Manual*
- *CF90 Co-array Programming Manual*

The following publication contains information specific to the MIPSpro 7 Fortran 90 compiler:

- *MIPSpro 7 Fortran 90 Commands and Directives Reference Manual*

CF90 and MIPSpro 7 Fortran 90 Compiler Messages

You can obtain CF90 and MIPSpro 7 Fortran 90 compiler message explanations by using the online `explain(1)` command.

CF90 and MIPSpro 7 Fortran 90 Compiler Man Pages

In addition to printed and online prose documentation, several online man pages describe aspects of the CF90 and MIPSpro 7 Fortran 90 compilers. Man pages exist for the library routines, the intrinsic procedures, and several programming environment tools.

You can print copies of online man pages by using the pipe symbol with the `man(1)`, `col(1)`, and `lpr(1)` commands. In the following example, these commands are used to print a copy of the `explain(1)` man page:

```
% man explain | col -b | lpr
```

Each man page includes a general description of one or more commands, routines, system calls, or other topics, and provides details of their usage

(command syntax, routine parameters, system call arguments, and so on). If more than one topic appears on a page, the entry in the printed manual is alphabetized under its primary name; online, secondary entry names are linked to these primary names. For example, `egrep` is a secondary entry on the page with a primary entry name of `grep`. To access `egrep` online, you can type `man grep` or `man egrep`. Both commands display the `grep` man page to your terminal.

Related Fortran Publications

The following commercially available reference books are among those that you can consult for more information on the history of Fortran and the Fortran language itself:

- Adams, J. C., W. S. Brainerd, J. T. Martin, B. T. Smith, and J. L. Wagener. *Fortran 95 Handbook : Complete ISO/ANSI Reference*. MIT Press, 1997. ISBN 0262510960.
- Chapman, S. *Fortran 90/95 for Scientists and Engineers*. McGraw Hill Text, 1998. ISBN 0070119384.
- Chapman, S. *Introduction to Fortran 90/95*. McGraw Hill Text, 1998. ISBN 0070119694.
- Counihan, M. *Fortran 95 : Including Fortran 90, Details of High Performance Fortran (HPF), and the Fortran Module for Variable-Length Character Strings*. UCL Press, 1997. ISBN 1857283678.
- Gehrke, W. *Fortran 95 Language Guide*. Springer Verlag, 1996. ISBN 3540760628.
- International Standards Organization. *ISO/IEC 1539-1:1997, Information technology — Programming languages — Fortran*. 1997.
- Metcalf, M. and J. Reid. *Fortran 90/95 Explained*. Oxford University Press, 1996. ISBN 0198518889.

Related Publications

Certain other publications from Silicon Graphics may also interest you.

On UNICOS and UNICOS/mk systems, the following documents contain information that may be useful when using the CF90 compiler:

- *Segment Loader (SEGLDR) and ld Reference Manual*
- *UNICOS User Commands Reference Manual*
- *UNICOS Performance Utilities Reference Manual*
- *Scientific Libraries Reference Manual*
- *Introducing the Program Browser*
- *Application Programmer's Library Reference Manual*
- *Guide to Parallel Vector Applications*
- *Introducing the Cray TotalView Debugger*
- *Introducing the MPP Apprentice Tool*
- *Application Programmer's I/O Guide*
- *Optimizing Code on Cray PVP Systems*
- *Compiler Information File (CIF) Reference Manual*

On IRIX systems, the following documents contain information that may be useful when using the MIPSpro 7 Fortran 90 compiler:

- *MIPSpro Compiling and Performance Tuning Guide*
- *MIPSpro Fortran 77 Programmer's Guide*
- *MIPSpro 64-Bit Porting and Transition Guide*
- *MIPSpro Assembly Language Programmer's Guide*

Obtaining Publications

The *User Publications Catalog* describes the availability and content of all Cray hardware and software documents that are available to customers. Customers who subscribe to the Cray Inform (CRInform) program can access this information on the CRInform system.

To order a document, call +1 651 683 5907. Silicon Graphics employees may send electronic mail to orderdsk@sgi.com (UNIX system users).

Customers who subscribe to the CRInform program can order software release packages electronically by using the Order Cray Software option.

Customers outside of the United States and Canada should contact their local service organization for ordering and documentation information.

Conventions

The following conventions are used throughout this document:

<u>Convention</u>	<u>Meaning</u>
command	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.
DEL or DELETED	The DEL or DELETED notation indicates that the feature being described has been deleted from the Fortran standard. The CF90 and MIPSpro 7 Fortran 90 compilers support these features, but the compilers issue a message when a deleted feature is encountered.
EXT or EXTENSION	The EXT or EXTENSION notation indicates that the feature being described is an extension to the Fortran standard. The CF90 and MIPSpro 7 Fortran 90 compilers issue a message when extensions are encountered.
OBS or OBSOLESCE	The OBS or OBSOLESCE notation indicates that the feature being described is considered to be obsolete in the Fortran standard. The CF90 and MIPSpro 7 Fortran 90 compilers support these features, but the compilers issue a message when an obsolescent feature is encountered.

<i>xyz_list</i>	When <i>_list</i> is part of a syntax description, it means that several items may be specified. For example, <i>xyz_list</i> can be expanded to mean <i>xyz [, xyz] . . .</i>
<i>scalar_</i>	When <i>scalar_</i> is the first item in a syntax description, it indicates that the item is a scalar, not an array, value.
<i>_name</i>	When <i>_name</i> is part of a syntax definition, it indicates that the item is a name with no qualification. For example, the item must not have a subscript list, so <code>ARRAY</code> is a name, but <code>ARRAY(I)</code> is not.
(Rnnnn)	Indicates that the Fortran 90 standard has rules regarding the characteristic of the language being discussed. All rules are numbered, and the numbered list appears in the <i>Fortran Language Reference Manual, Volume 3</i> . The numbering of the rules in the <i>Fortran Language Reference Manual, Volume 3</i> matches the numbering of the rules in the standard. The forms of the rules in the <i>Fortran Language Reference Manual, Volume 3</i> and the BNF syntax class terms that are used may differ from the rules and terms used in the standard.
POINTER	The term <code>POINTER</code> refers to the Fortran <code>POINTER</code> attribute.
Cray pointer	The term <i>Cray pointer</i> refers to the Cray pointer data type extension.
Fortran Fortran standard	These terms refer to the current Fortran standard, which is the Fortran 95 standard. For situations when it might otherwise be confusing, a specific

standard is mentioned along with its numeric identifier (FORTRAN 77, Fortran 90, Fortran 95).

BNF Conventions

This section describes some of the commonly used Backus-Naur Form (BNF) conventions.

Terms such as *goto_stmt* are called *variable entries*, *nonterminal symbols*, or simply, *nonterminals*. The metalanguage term *goto_stmt*, for example, represents the GO TO statement, as follows:

<i>goto_stmt</i>	is	GOTO <i>label</i>
------------------	----	-------------------

The syntax rule defines *goto_stmt* to be GO TO *label*, which describes the format of the GO TO statement. The description of the GO TO statement is incomplete until the definition of *label* is given. *label* is also a nonterminal symbol. A further search for *label* will result in a specification of *label* and thereby provide the complete statement definition. A *terminal* part of a syntax rule is one that does not need further definition. For example, GO TO is a terminal keyword and is a required part of the statement form. The complete BNF list appears in the *Fortran Language Reference Manual, Volume 3*.

The following abbreviations are commonly used in naming nonterminal keywords:

<u>Abbreviation</u>	<u>Term</u>
<i>arg</i>	argument
<i>attr</i>	attribute
<i>char</i>	character
<i>decl</i>	declaration
<i>def</i>	definition
<i>desc</i>	descriptor
<i>expr</i>	expression
<i>int</i>	integer
<i>op</i>	operator

spec specifier or specification
stmt statement

The term **is** separates the syntax class name from its definition. The term **or** indicates an alternative definition for the syntactic class being defined. The following example shows that *add_op*, the add operator, may be either a plus sign (+) or a minus sign (-):

<i>add_op</i>	is	+
	or	-

Indentation indicates syntax continuation. If a rule does not fit on one line, the second line is indented. This is shown in the following example:

<i>dimension_stmt</i>	is	DIMENSION [::] <i>array_name</i> (<i>array_spec</i>)
		[, <i>array_name</i> (<i>array_spec</i>)] . . .

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and part number of the document with your comments.

You can contact us in any of the following ways:

- Send electronic mail to the following address:

techpubs@sgi.com

- Send a facsimile to the attention of "Technical Publications" at fax number +1 650 932 0801.
- Use the Suggestion Box form on the Technical Publications Library World Wide Web page:

<http://techpubs.sgi.com/library/>

- Call the Technical Publications Group, through the Technical Assistance Center, using one of the following numbers:

For Silicon Graphics IRIX based operating systems: 1 800 800 4SGI

For UNICOS or UNICOS/mk based operating systems or CRAY Origin2000 systems: 1 800 950 2729 (toll free from the United States and Canada) or +1 651 683 5600

- Send mail to the following address:

Technical Publications
Silicon Graphics, Inc.
2011 North Shoreline Boulevard, M/S 535
Mountain View, California 94043-1389

We value your comments and will respond to them promptly.

Fortran Syntax [1]

This chapter contains a complete description of the Fortran syntax. Section 1.1 describes the format of the syntax. Section 1.2, page 5, contains the complete syntax and constraints as they appear in the Fortran standard. A high-level summary of the syntax appears in the *Fortran Language Reference Manual, Volume 1*.

1.1 Syntax Form

The syntax of Fortran programs is described using a variant of the Backus-Naur Form (BNF).

1.1.1 Syntax Rules Expressed in BNF

The BNF syntax rules are expressed as a definition. The metalanguage class being defined is first, followed by the symbol **is**, and finally the syntax definition, as in the following example:

```
goto_stmt                    is                    GO TO label
```

The term *goto_stmt* represents the GO TO statement; such terms are called *nonterminal symbols* or simply *nonterminals*. The syntax rule defines *goto_stmt* as GO TO *label*, which describes the form of the GO TO statement. The description of the GO TO statement is not complete until the definition of *label* is specified; *label* is also a nonterminal symbol. A further search for *label* in the BNF will result in a specification of *label* and thereby provide the complete statement definition. A *terminal* part of a syntax rule does not need further definition. For example, GO TO is a terminal and is a required part of the statement form.

In many cases, you can derive information about the metalanguage class from part of the descriptive term. The part can be a complete word, such as *_list*, or a common abbreviation. Some abbreviations used consistently in metalanguage classes are listed in Table 1, page 2.

Table 1. Syntax metalanguage abbreviations

Abbreviation	Term
<i>arg</i>	Argument
<i>attr</i>	Attribute
<i>char</i>	Character
<i>decl</i>	Declaration
<i>def</i>	Definition
<i>desc</i>	Descriptor
<i>expr</i>	Expression
<i>int</i>	Integer
<i>op</i>	Operator
<i>spec</i>	Specifier or specification
<i>stmt</i>	Statement

For example, all class definitions that end with *_stmt* might be used to generate a complete list of the statements in Fortran.

1.1.2 Definition Syntax Symbol: **Is**

As the following example shows, the symbol **is** separates the syntax class name from its definition:

```
goto_stmt           is           GO TO label  
power_op           is           **
```

1.1.3 Alternative Syntax Symbol: **Or**

The symbol **or** indicates an alternative definition for the syntactic class being defined. The following example shows that *add_op*, the add operator, can be either plus or minus.

1.1.7 Assumed Syntax Rules

In order to minimize the number of syntax rules and still convey an appropriate meaning, some portions of the BNF metaterms have assumed meanings. In the following example, *xyz* represents any BNF phrase:

<i>xyz_list</i>	means	<i>xyz</i> [, <i>xyz</i>] . . .
<i>xyz_name</i>	is	a name
<i>scalar_xyz</i>	is	an <i>xyz</i> that is a scalar

1.1.8 Example BNF Syntax

Consider the following example:

<i>read_stmt</i>	is	READ (<i>io_control_spec_list</i>) [<i>input_item_list</i>]
	or	READ <i>format</i> [, <i>input_item_list</i>]
<i>format</i>	is	<i>default_char_expr</i>
	or	<i>label</i>
	or	*
	or	<i>scalar_default_int_variable</i>

In this example, there are two alternatives to the READ statement. The first uses an input/output (I/O) control specification list; the second is a formatted READ statement where the unit is processor dependent. Both alternatives have an optional input item list, indicated by []. The syntax class *format* (a nonterminal) is further defined as either a default character expression containing the format specifications, or a statement label referring to a separate FORMAT statement that contains the format specifications, or an asterisk (*) indicating that the READ statement is list-directed, or a scalar default integer variable whose value specifies the label of a FORMAT statement. In the standard, the last alternative is printed in a smaller font because it is an obsolescent feature that may be removed in a later revision of the standard, including the next revision; this convention is not used in this manual.

There are other nonterminal symbols in the description of the `READ` statement and further BNF rules need to be examined to determine the complete description of the `READ` statement.

1.1.9 Constraints

The BNF forms do not provide a complete description of the syntax; additional constraints are described with text. The BNF rules and the constraints both describe the syntax of Fortran. Constraints are restrictions to the syntax rules that limit the form of the statement described. If present, constraints appear following a syntax rule.

1.1.10 Identifying Numbers

In the text of the standard, each BNF rule is given an identifying number, R201 for example. The numbering of the rules in the following subsections matches the numbering of the rules in the standard.

BNF rules are also used to describe extensions. In the following BNF description, for example, "EXT" in the leftmost column indicates that the CF90 and MIPSpro 7 Fortran 90 compilers also allow *unit_name* to be used as an *io_unit*:

R901	<i>io_unit</i>	is	<i>external_file_unit</i>
		or	*
		or	<i>internal_file_unit</i>
EXT		or	<i>unit_name</i>

1.2 Syntax Rules and Constraints

Each of the following sections contains the syntax rules and constraints from a section of the Fortran standard. The following sections use an underscore, rather than a hyphen, as a separator; this differs from the Fortran standard. The rules in the following sections have been amended to include BNF for the CF90 and MIPSpro 7 Fortran 90 compiler extensions to the Fortran standard, but the constraints have not been modified to reflect the extensions.

1.2.1 Introduction

There are no syntax rules described in section 1, "Introduction," of the Fortran 95 standard.

1.2.2 Fortran Terms and Concepts

The following syntax rules are described in section 2, "Fortran terms and concepts," of the Fortran 95 standard.

R201	<i>executable_program</i>	is	<i>program_unit</i> [<i>program_unit</i>] . . .
R202	<i>program_unit</i>	is	<i>main_program</i>
		or	<i>external_subprogram</i>
		or	<i>module</i>
		or	<i>block_data</i>
R1101	<i>main_program</i>	is	[<i>program_stmt</i>] [<i>specification_part</i>] [<i>execution_part</i>] [<i>internal_subprogram_part</i>] <i>end_program_stmt</i>

Constraint: An *execution_part* must not contain an *end_function_stmt*, *end_program_stmt*, or *end_subroutine_stmt*.

R203	<i>external_subprogram</i>	is	<i>function_subprogram</i>
		or	<i>subroutine_subprogram</i>
R1215	<i>function_subprogram</i>	is	<i>function_stmt</i> [<i>specification_part</i>] [<i>execution_part</i>] [<i>internal_subprogram_part</i>] <i>end_function_stmt</i>

R1219	<i>subroutine_subprogram</i>	is	<i>subroutine_stmt</i> [<i>specification_part</i>] [<i>execution_part</i>] [<i>internal_subprogram_part</i>] <i>end_subroutine_stmt</i>
R1104	<i>module</i>	is	<i>module_stmt</i> [<i>specification_part</i>] [<i>module_subprogram_part</i>] <i>end_module_stmt</i>
R1112	<i>block_data</i>	is	<i>block_data_stmt</i> [<i>specification_part</i>] <i>end_block_data_stmt</i>
R204	<i>specification_part</i>	is	[<i>use_stmt</i>] ... [<i>implicit_part</i>] [<i>declaration_construct</i>] ...
R205	<i>implicit_part</i>	is	[<i>implicit_part_stmt</i>] ... <i>implicit_stmt</i>
R206	<i>implicit_part_stmt</i>	is or or or	<i>implicit_stmt</i> <i>parameter_stmt</i> <i>format_stmt</i> <i>entry_stmt</i>
R207	<i>declaration_construct</i>	is or or or or or or or	<i>derived_type_def</i> <i>interface_block</i> <i>type_declaration_stmt</i> <i>specification_stmt</i> <i>parameter_stmt</i> <i>format_stmt</i> <i>entry_stmt</i> <i>stmt_function_stmt</i>
R208	<i>execution_part</i>	is	<i>executable_construct</i> [<i>execution_part_construct</i>] ...
R209	<i>execution_part_construct</i>	is or	<i>executable_construct</i> <i>format_stmt</i>

		or	<i>data_stmt</i>
		or	<i>entry_stmt</i>
R210	<i>internal_subprogram_part</i>	is	<i>contains_stmt</i> <i>internal_subprogram</i> [<i>internal_subprogram</i>] . . .
R211	<i>internal_subprogram</i>	is	<i>function_subprogram</i>
		or	<i>subroutine_subprogram</i>
R212	<i>module_subprogram_part</i>	is	<i>contains_stmt</i> <i>module_subprogram</i> [<i>module_subprogram</i>] . . .
R213	<i>module_subprogram</i>	is	<i>function_subprogram</i>
		or	<i>subroutine_subprogram</i>
R214	<i>specification_stmt</i>	is	<i>access_stmt</i>
		or	<i>allocatable_stmt</i>
EXT		or	<i>automatic_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>save_stmt</i>
		or	<i>target_stmt</i>
EXT		or	<i>volatile_stmt</i>
R215	<i>executable_construct</i>	is	<i>action_stmt</i>
		or	<i>case_construct</i>

		or	<i>do_construct</i>
		or	<i>if_construct</i>
		or	<i>forall_construct</i>
		or	<i>where_construct</i>
R216	<i>action_stmt</i>	is	<i>allocate_stmt</i>
OBS		or	<i>arithmetic_if_stmt</i>
EXT		or	<i>assign_stmt</i>
EXT		or	<i>assigned_goto_stmt</i>
		or	<i>assignment_stmt</i>
		or	<i>backspace_stmt</i>
EXT		or	<i>buffer_in_stmt</i>
EXT		or	<i>buffer_out_stmt</i>
		or	<i>call_stmt</i>
		or	<i>close_stmt</i>
OBS		or	<i>computed_goto_stmt</i>
		or	<i>continue_stmt</i>
		or	<i>cycle_stmt</i>
		or	<i>deallocate_stmt</i>
		or	<i>endfile_stmt</i>
		or	<i>end_function_stmt</i>
		or	<i>end_program_stmt</i>
		or	<i>end_subroutine_stmt</i>
		or	<i>exit_stmt</i>
		or	<i>forall_stmt</i>
		or	<i>goto_stmt</i>
		or	<i>if_stmt</i>
		or	<i>inquire_stmt</i>
		or	<i>nullify_stmt</i>
		or	<i>open_stmt</i>

EXT	or	<i>pause_stmt</i>
	or	<i>pointer_assignment_stmt</i>
	or	<i>print_stmt</i>
	or	<i>read_stmt</i>
	or	<i>return_stmt</i>
	or	<i>rewind_stmt</i>
	or	<i>stop_stmt</i>
	or	<i>where_stmt</i>
	or	<i>write_stmt</i>

1.2.3 Characters, Lexical Tokens, and Source Form

The following syntax rules are described in section 3, "Characters, lexical tokens, and source form," of the Fortran 95 standard.

R301	<i>character</i>	is	<i>alphanumeric_character</i>
		or	<i>special_character</i>
R302	<i>alphanumeric_character</i>	is	<i>letter</i>
		or	<i>digit</i>
		or	<i>underscore</i>
EXT		or	<i>currency_symbol</i>
EXT		or	<i>at_sign</i>

Note: The MIPSpro 7 Fortran 90 compiler does not support the *at_sign* (@).

R303	<i>underscore</i>	is	<i>_</i>
EXT	<i>currency_symbol</i>	is	<i>\$</i>
EXT	<i>at_sign</i>	is	<i>@</i>
R304	<i>name</i>	is	<i>letter</i> [<i>alphanumeric_character</i>] . . .

Constraint: The maximum length of a name is 31 characters.

R305	<i>constant</i>	is	<i>literal_constant</i>
		or	<i>named_constant</i>
R306	<i>literal_constant</i>	is	<i>int_literal_constant</i>
		or	<i>real_literal_constant</i>
		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>

Constraint: *int_constant* must be of type integer.

R309	<i>char_constant</i>	is	<i>constant</i>
------	----------------------	-----------	-----------------

Constraint: *char_constant* must 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>

R708	<i>power_op</i>	is	**
R709	<i>mult_op</i>	is	*
		or	/
R710	<i>add_op</i>	is	+
		or	-
R712	<i>concat_op</i>	is	//
R714	<i>rel_op</i>	is	.EQ.
		or	.NE.
		or	.LT.
		or	.LE.
		or	.GT.
		or	.GE.
EXT		or	.LG.
		or	==
		or	/=
		or	<
		or	<=
		or	>
		or	>=
EXT		or	<>
R719	<i>not_op</i>	is	.NOT.
EXT		or	.N.
R720	<i>and_op</i>	is	.AND.
EXT		or	.A.
R721	<i>or_op</i>	is	.OR.
EXT		or	.O.
R722	<i>equiv_op</i>	is	.EQV.
		or	.NEQV.
EXT	<i>exclusive_disjunct_op</i>	is	.XOR.

EXT		or	.X.
R311	<i>defined_operator</i>	is	<i>defined_unary_op</i>
		or	<i>defined_binary_op</i>
		or	<i>extended_intrinsic</i>
R704	<i>defined_unary_op</i>	is	. letter [letter]
R724	<i>defined_binary_op</i>	is	. letter [letter]
R312	<i>extended_intrinsic_op</i>	is	<i>intrinsic_operator</i>

Constraint: A *defined_unary_op* and a *defined_binary_op* must not contain more than 31 letters and must not be the same as any *intrinsic_operator* or *logical_literal_constant*.

R313	<i>label</i>	is	<i>digit</i> [<i>digit</i> [<i>digit</i> [<i>digit</i> [<i>digit</i>]]]]
------	--------------	-----------	---

Constraint: At least one digit in a *label* must be nonzero.

1.2.4 Intrinsic and Derived Data Types

The following syntax rules are described in section 4, "Intrinsic and derived data types," of the Fortran 95 standard.

R401	<i>signed_digit_string</i>	is	[<i>sign</i>] <i>digit_string</i>
R402	<i>digit_string</i>	is	<i>digit</i> [<i>digit</i>] . . .
R403	<i>signed_int_literal_constant</i>	is	[<i>sign</i>] <i>int_literal_constant</i>
R404	<i>int_literal_constant</i>	is	<i>digit_string</i> [<i>_kind_param</i>]
R405	<i>kind_param</i>	is	<i>digit_string</i>
		or	<i>scalar_int_constant_name</i>

Constraint: The value of *kind_param* must be nonnegative.

Constraint: The value of *kind_param* must specify a representation method that exists on the processor.

R406	<i>sign</i>	is	+
		or	-
R407	<i>boz_literal_constant</i>	is	<i>binary_constant</i>
		or	<i>octal_constant</i>
		or	<i>hex_constant</i>

Constraint: A *boz_literal_constant* may appear only in a DATA statement.

R408	<i>binary_constant</i>	is	B ' <i>digit</i> [<i>digit</i>] ... '
		or	B " <i>digit</i> [<i>digit</i>] ... "

Constraint: *digit* must have one of the values 0 or 1.

R409	<i>octal_constant</i>	is	O ' <i>digit</i> [<i>digit</i>] ... '
		or	O " <i>digit</i> [<i>digit</i>] ... "

Constraint: *digit* must have one of the values 0 through 7.

R410	<i>hex_constant</i>	is	Z ' <i>hex_digit</i> [<i>hex_digit</i>] ... '
		or	Z " <i>hex_digit</i> [<i>hex_digit</i>] ... "
R411	<i>hex_digit</i>	is	<i>digit</i>
		or	A
		or	B
		or	C
		or	D
		or	E
		or	F
R412	<i>signed_real_literal_constant</i>	is	[<i>sign</i>] <i>real_literal_constant</i>

R413	<i>real_literal_constant</i>	is	<i>significand</i> [<i>exponent_letter</i> <i>exponent</i>] [<i>_kind_param</i>]
		or	<i>digit_string</i> <i>exponent_letter</i> <i>exponent</i> [<i>_kind_param</i>]
R414	<i>significand</i>	is	<i>digit_string</i> . [<i>digit_string</i>]
		or	. <i>digit_string</i>
R415	<i>exponent_letter</i>	is	E
		or	D
EXT		or	Q
R416	<i>exponent</i>	is	<i>signed_digit_string</i>

Constraint: If both *kind_param* and *exponent_letter* are present, *exponent_letter* must be E.

Constraint: The value of *kind_param* must specify an approximation method that exists on the processor.

R417	<i>complex_literal_constant</i>	is	(<i>real_part</i> , <i>imag_part</i>)
R418	<i>real_part</i>	is	<i>signed_int_literal_constant</i>
		or	<i>signed_real_literal_constant</i>
R419	<i>imag_part</i>	is	<i>signed_int_literal_constant</i>
		or	<i>signed_real_literal_constant</i>
R420	<i>char_literal_constant</i>	is	[<i>kind_param</i> _] ' [<i>ASCII_char</i>] . . . '
		or	[<i>kind_param</i> _] " [<i>ASCII_char</i>] . . . "

Constraint: The value of *kind_param* must specify a representation method that exists on the processor.

R421	<i>logical_literal_constant</i>	is	.TRUE. [<i>_kind_param</i>]
		or	.FALSE. [<i>_kind_param</i>]

Constraint: The value of *kind_param* must specify a representation method that exists on the processor.

R422	<i>derived_type_def</i>	is	<i>derived_type_stmt</i> [<i>private_sequence_stmt</i>] . . . <i>component_def_stmt</i> [<i>component_def_stmt</i>] . . . <i>end_type_stmt</i>
R423	<i>derived_type_stmt</i>	is	TYPE [[, <i>access_spec</i>] ::] <i>type_name</i>
R424	<i>private_sequence_stmt</i>	is	PRIVATE or SEQUENCE

Constraint: The same *private_sequence_stmt* must not appear more than once in a given *derived_type_def*.

Constraint: If SEQUENCE is present, all derived types specified in component definitions must be sequence types.

Constraint: An *access_spec* or a PRIVATE statement within the definition is permitted only if the type definition is within the specification part of a module.

Constraint: A derived type *type_name* must not be the same as the name of any intrinsic type nor the same as any other accessible derived type *type_name*.

R425	<i>component_def_stmt</i>	is	<i>type_spec</i> [[, <i>component_attr_spec_list</i>] ::] <i>component_decl_list</i>
R426	<i>component_attr_spec</i>	is	POINTER or DIMENSION (<i>component_array_spec</i>)
R427	<i>component_array_spec</i>	is	<i>explicit_shape_spec_list</i> or <i>deferred_shape_spec_list</i>

Constraint: If a component of a derived type is of a type declared to be private, either the derived type definition must contain the PRIVATE statement or the derived type must be private.

Constraint: No *component_attr_spec* can appear more than once in a given *component_def_stmt*.

Constraint: If the POINTER attribute is not specified for a component, a *type_spec* in the *component_def_stmt* must specify an intrinsic type or a previously defined derived type.

Constraint: If the `POINTER` attribute is specified for a component, a *type_spec* in the *component_def_stmt* must specify an intrinsic type or any accessible derived type including the type being defined.

Constraint: If the `POINTER` attribute is not specified, each *component_array_spec* must be an *explicit_shape_spec_list*.

Constraint: If the `POINTER` attribute is specified, each *component_array_spec* must be a *deferred_shape_spec_list*.

R428	<i>component_decl</i>	is	<i>component_name</i> [(<i>component_array_spec</i>)] [* <i>char_length</i>] <i>component_initialization</i>
R429	<i>component_initialization</i>	is	= <i>initialization_expr</i>
		or	=> <code>NULL()</code>

Constraint: The * *char_length* option is permitted only if the type specified is character.

Constraint: The character length specified by a *char_length* in a *component_decl* or the *char_selector* in a *type_spec* must be a constant specification expression.

Constraint: Each bound in the *explicit_shape_spec* must be a constant specification expression.

Constraint: If *component_initialization* appears, a double colon separator (`::`) must appear before the *component_decl_list*.

Constraint: If `=>` appears in *component_initialization*, the `POINTER` attribute must appear in the *component_attr_spec_list*. If `=` appears in *component_initialization*, the `POINTER` attribute must not appear in the *component_attr_spec_list*.

R430	<i>end_type_stmt</i>	is	<code>END TYPE</code> [<i>type_name</i>]
------	----------------------	-----------	--

Constraint: If `END TYPE` is followed by a *type_name*, the *type_name* must be the same as that in the corresponding *derived_type_stmt*.

R431	<i>structure_constructor</i>	is	<i>type_name</i> (<i>expr_list</i>)
R432	<i>array_constructor</i>	is	(/ <i>ac_value_list</i> /)
R433	<i>ac_value</i>	is	<i>expr</i>
		or	<i>ac_implied_do</i>
R434	<i>ac_implied_do</i>	is	(<i>ac_value_list</i> , <i>ac_implied_do_control</i>)
R435	<i>ac_implied_do_control</i>	is	<i>ac_do_variable</i> = <i>scalar_int_expr</i> , <i>scalar_int_expr</i> [, <i>scalar_int_expr</i>]
R436	<i>ac_do_variable</i>	is	<i>scalar_int_variable</i>

Constraint: *ac_do_variable* must be a named variable.

Constraint: Each *ac_value* expression in the *array_constructor* must have the same type and kind type parameters.

1.2.5 Data Object Declarations and Specifications

The following syntax rules are described in section 5, "Data object declarations and specifications," of the Fortran 95 standard.

R501	<i>type_declaration_stmt</i>	is	<i>type_spec</i> [[, <i>attr_spec</i>] . . . ::] <i>entity_decl_list</i>
R502	<i>type_spec</i>	is	INTEGER [<i>kind_selector</i>]
EXT		or	INTEGER* <i>length_value</i>
		or	REAL [<i>kind_selector</i>]
EXT		or	REAL* <i>length_value</i>
		or	DOUBLE PRECISION
EXT		or	DOUBLE PRECISION* <i>length_value</i>
		or	COMPLEX [<i>kind_selector</i>]
EXT		or	COMPLEX* <i>length_value</i>
		or	CHARACTER [<i>char_selector</i>]
		or	LOGICAL [<i>kind_selector</i>]
EXT		or	LOGICAL* <i>length_value</i>
		or	TYPE (<i>type_name</i>)

EXT	or	POINTER (<i>pointer_name</i> , <i>pointee_name</i> [(<i>array_spec</i>)]) [, (<i>pointer_name</i> , <i>pointee_name</i> [(<i>array_spec</i>)])] ...
R503 <i>attr_spec</i>	is	PARAMETER
	or	<i>access_spec</i>
	or	ALLOCATABLE
EXT	or	AUTOMATIC
	or	DIMENSION (<i>array_spec</i>)
	or	EXTERNAL
	or	INTENT (<i>intent_spec</i>)
	or	INTRINSIC
	or	OPTIONAL
	or	POINTER
	or	SAVE
	or	TARGET
EXT	or	VOLATILE
R504 <i>entity_decl</i>	is	<i>object_name</i> [(<i>array_spec</i>)] [* <i>char_length</i>] [= <i>initialization_expr</i>]
	or	<i>function_name</i> [* <i>char_length</i>]
R505 <i>initialization</i>	is	= <i>initialization_expr</i>
	or	=> NULL()
R506 <i>kind_selector</i>	is	([KIND =] <i>scalar_int_initialization_expr</i>)

Constraint: The same *attr_spec* must not appear more than once in a given *type_declaration_stmt*.

Constraint: The *function_name* must be the name of an external function, an intrinsic function, a function dummy procedure, or a statement function.

Constraint: The *initialization_expr* must appear if the statement contains a PARAMETER attribute.

Constraint: If *initialization_expr* appears, a double colon separator (: :) must appear before the *entity_decl_list*.

Constraint: The *initialization_expr* must not appear if *object_name* is a dummy argument, a function result, or 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 array, an external name, an intrinsic name, or an automatic object.

Constraint: The **char_length* option is permitted only if the type specified is character.

Constraint: The `ALLOCATABLE` attribute may be used only when declaring an array that is not a dummy argument or a function result.

Constraint: An array declared with a `POINTER` or an `ALLOCATABLE` attribute must be specified with an *array_spec* that is a *deferred_shape_spec_list*.

Constraint: An *array_spec* for an *object_name* that is a function result that does not have the `POINTER` attribute must be an *explicit_shape_spec_list*.

Constraint: An *array_spec* for an *object_name* that is a function result that has the `POINTER` attribute must be a *deferred_shape_spec_list*.

Constraint: If the `POINTER` attribute is specified, the `TARGET`, `INTENT`, `EXTERNAL`, or `INTRINSIC` attribute must not be specified.

Constraint: If the `TARGET` attribute is specified, the `POINTER`, `EXTERNAL`, `INTRINSIC`, or `PARAMETER` attribute must not be specified.

Constraint: The `PARAMETER` attribute must not be specified for dummy arguments, pointers, allocatable arrays, functions, or objects in a common block.

Constraint: The `INTENT` and `OPTIONAL` attributes may be specified only for dummy arguments.

Constraint: An entity must not have the `PUBLIC` attribute if its type has the `PRIVATE` attribute.

Constraint: The `SAVE` attribute must not be specified for an object that is in a common block, a dummy argument, a procedure, a function result, an automatic data object, or an object with the `PARAMETER` attribute.

Constraint: An entity must not have the `EXTERNAL` attribute if it has the `INTRINSIC` attribute.

Constraint: An entity in an *entity_decl_list* must not have the `EXTERNAL` or `INTRINSIC` attribute specified unless it is a function.

Constraint: If `=>` appears in *initialization*, the object must have the `POINTER` attribute. If `=` appears in *initialization*, the object must not have the `POINTER` attribute.

Constraint: An array must not have both the `ALLOCATABLE` attribute and the `POINTER` attribute.

Constraint: An entity must not be given explicitly any attribute more than once in a scoping unit.

Constraint: The value of *scalar_int_initialization_expr* must be nonnegative and must specify a representation method that exists on the processor.

R507	<i>char_selector</i>	is	<i>length_selector</i>
		or	(<code>LEN = char_len_param_value</code> , <code>KIND = scalar_int_initialization_expr</code>)
		or	(<code>char_len_param_value</code> , [<code>KIND =</code>] <code>scalar_int_initialization_expr</code>)
		or	(<code>KIND = scalar_int_initialization_expr</code> [, <code>LEN = char_len_param_value</code>])
R508	<i>length_selector</i>	is	([<code>LEN =</code>] <code>char_len_param_value</code>)
OBS		or	* <code>char_length</code> [,]
R509	<i>char_length</i>	is	(<code>char_len_param_value</code>)
		or	<code>scalar_int_literal_constant</code>

Obsolescent Constraint: The optional comma in a *length_selector* is permitted only in a *type_spec* in a *type_declaration_stmt*.

Obsolescent Constraint: The optional comma in a *length_selector* is permitted only if no double colon separator appears in the *type_declaration_stmt*.

Constraint: The value of *scalar_int_initialization_expr* must be nonnegative and must specify a representation method that exists on the processor.

Constraint: The *scalar_int_literal_constant* must not include a *kind_param*.

R510	<i>char_len_param_value</i>	is	<i>specification_expr</i>
		or	*

Obsolescent Constraint: A function name must not be declared with a **char_len_param_value* unless it is the name of an external function or the name of a dummy function.

Constraint: A function name declared with a **char_len_param_value* must not be array-valued, pointer-valued, pure, or recursive.

R511	<i>access_spec</i>	is	PUBLIC
		or	PRIVATE

Constraint: An *access_spec* attribute may appear only in the specification part of a module.

R512	<i>intent_spec</i>	is	IN
		or	OUT
		or	INOUT

Constraint: The `INTENT` attribute must not be specified for a dummy argument that is a dummy procedure or a dummy pointer.

Constraint: A dummy argument with the `INTENT(IN)` attribute, or a subobject of such a dummy argument, must not appear as any of the following:

- The *variable* of an *assignment_stmt*
- The *pointer_object* of a *pointer_assignment_stmt*
- A DO-variable or implied DO-variable
- An *input_item* in a *read_stmt*
- A *variable_name* in a *namelist_stmt* if the *name_list_group_name* appears in an `NML=` specifier in a *read_stmt*
- An *internal_file_unit* in a *write_stmt*
- An `IOSTAT=` or `SIZE=` specifier in an I/O statement
- A definable variable in an `INQUIRE` statement
- A *stat_variable* or *allocate_object* in an *allocate_stmt* or a *deallocate_stmt*
- An actual argument in a reference to a procedure with an explicit interface when the associated dummy argument has the `INTENT(OUT)` or `INTENT(INOUT)` attribute

R513	<i>array_spec</i>	is	<i>explicit_shape_spec_list</i>
		or	<i>assumed_shape_spec_list</i>
		or	<i>deferred_shape_spec_list</i>
		or	<i>assumed_size_spec</i>

Constraint: The maximum rank is seven.

R514	<i>explicit_shape_spec</i>	is	[<i>lower_bound</i> :] <i>upper_bound</i>
R515	<i>lower_bound</i>	is	<i>specification_expr</i>
R516	<i>upper_bound</i>	is	<i>specification_expr</i>

Constraint: An explicit-shape array whose bounds depend on the values of nonconstant expressions must be a dummy argument, a function result, or an automatic array of a procedure.

R517	<i>assumed_shape_spec</i>	is	[<i>lower_bound</i>] :
R518	<i>deferred_shape_spec</i>	is	:
R519	<i>assumed_size_spec</i>	is	[<i>explicit_shape_spec_list</i> ,] [<i>lower_bound</i> :] *

Constraint: The function name of an array-valued function must not be declared as an assumed-size array.

Constraint: An assumed-size array with INTENT(OUT) must not be of a type for which default initialization is specified.

R520	<i>intent_stmt</i>	is	INTENT (<i>intent_spec</i>) [::] <i>dummy_arg_name_list</i>
------	--------------------	-----------	---

Constraint: An *intent_stmt* may appear only in the *specification_part* of a subprogram or an interface body.

Constraint: *dummy_arg_name* must not be the name of a dummy procedure or a dummy pointer.

R521	<i>optional_stmt</i>	is	OPTIONAL [[::] <i>dummy_arg_name_list</i>]
------	----------------------	-----------	--

Constraint: An *optional_stmt* can occur only in the *specification_part* of a subprogram or an interface body.

R522	<i>access_stmt</i>	is	<i>access_spec</i> [[::] <i>access_id_list</i>]
R523	<i>access_id</i>	is	<i>use_name</i>
		or	<i>generic_spec</i>

Constraint: An *access_stmt* can 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.

Constraint: Each *use_name* must be the name of a named variable, procedure, derived type, named constant, or namelist group.

Constraint: A module procedure that has a dummy argument or function result of a type that has PRIVATE accessibility must have PRIVATE accessibility and must not have a generic identifier that has PUBLIC accessibility.

R524	<i>save_stmt</i>	is	SAVE [[::] <i>saved_entity_list</i>]
R525	<i>saved_entity</i>	is	<i>object_name</i>
		or	/ <i>common_block_name</i> /

Constraint: An *object_name* must not be the name of an object in a common block, a dummy argument name, a procedure name, a function result name, an automatic data object name, or the name of an object with the PARAMETER attribute.

Constraint: If a SAVE statement with an omitted saved entity list occurs in a scoping unit, no other explicit occurrence of the SAVE attribute or SAVE statement is permitted in the same scoping unit.

R526	<i>dimension_stmt</i>	is	DIMENSION [::] <i>array_name</i> (<i>array_spec</i>) [, <i>array_name</i> (<i>array_spec</i>)] ...
R527	<i>allocatable_stmt</i>	is	ALLOCATABLE [::] <i>array_name</i> [(<i>deferred_shape_spec_list</i>)] [, <i>array_name</i> [(<i>deferred_shape_spec_list</i>)]] ...

Constraint: The *array_name* must not be a dummy argument or function result.

Constraint: If the DIMENSION attribute for an *array_name* is specified elsewhere in the scoping unit, the *array_spec* must be a *deferred_shape_spec_list*.

R528	<i>pointer_stmt</i>	is	POINTER [::] <i>object_name</i> [(<i>deferred_shape_spec_list</i>)] [, <i>object_name</i> [(<i>deferred_shape_spec_list</i>)]] ...
------	---------------------	-----------	---

Constraint: The INTENT attribute must not be specified for an *object_name*.

Constraint: If the DIMENSION attribute for an *object_name* is specified elsewhere in the scoping unit, the *array_spec* must be a *deferred_shape_spec_list*.

Constraint: The PARAMETER attribute must not be specified for an *object_name*.

R529	<i>target_stmt</i>	is	TARGET [::] <i>object_name</i> [(<i>array_spec</i>)] [, <i>object_name</i> [(<i>array_spec</i>)]] ...
------	--------------------	-----------	--

Constraint: The PARAMETER attribute must not be specified for an *object_name*.

R532	<i>data_stmt</i>	is	DATA <i>data_stmt_set</i> [[,] <i>data_stmt_set</i>] ...
R533	<i>data_stmt_set</i>	is	<i>data_stmt_object_list</i> / <i>data_stmt_value_list</i> / [[,] <i>data_stmt_object_list</i> / <i>data_stmt_value_list</i> /] ...
R534	<i>data_stmt_object</i>	is	<i>variable</i>
		or	<i>data_implied_do</i>

R535	<i>data_implied_do</i>	is	(<i>data_i_do_object_list</i> , <i>data_i_do_variable</i> = <i>scalar_int_expr</i> , <i>scalar_int_expr</i> [, <i>scalar_int_expr</i>])
R536	<i>data_i_do_object</i>	is	<i>array_element</i>
		or	<i>scalar_structure_component</i>
		or	<i>data_implied_do</i>
R538	<i>data_stmt_value</i>	is	[<i>data_stmt_repeat</i> *] <i>data_stmt_constant</i>
R539	<i>data_stmt_repeat</i>	is	<i>scalar_int_constant</i>
		or	<i>scalar_int_constant_subobject</i>
R540	<i>data_stmt_constant</i>	is	<i>scalar_constant</i>
		or	<i>scalar_constant_subobject</i>
		or	<i>signed_int_literal_constant</i>
		or	<i>signed_real_literal_constant</i>
		or	<i>structure_constructor</i>
		or	NULL()
		or	<i>boz_literal_constant</i>
EXT		or	<i>typeless_constant</i>

A *data_i_do_variable* must be a named variable.

Constraint: The *array_element* must not have a constant parent.

Constraint: The *scalar_structure_component* must not have a constant parent.

Constraint: In a *scalar_int_constant_subobject* that is a *data_stmt_repeat*, any subscript must be an initialization expression.

Constraint: In a *scalar_constant_subobject* that is a *data_stmt_constant*, any subscript, substring starting point, or substring ending point must be an initialization expression.

Constraint: If a DATA statement constant value is a named constant or a structure constructor, the named constant or derived type must have been declared previously in the scoping unit or made accessible by USE or host association.

R537	<i>data_i_do_variable</i>	is	<i>scalar_int_variable</i>
------	---------------------------	-----------	----------------------------

Constraint: *data_i_do_variable* must be a named variable.

Constraint: The DATA statement repeat factor must be positive or zero. If the DATA statement repeat factor is a named constant, it must have been declared previously in the scoping unit or made accessible by use association or host association.

Constraint: In a *scalar_int_constant_subobject* that is a *data_stmt_repeat*, any subscript must be an initialization expression.

Constraint: In a *scalar_constant_subobject* that is a *data_stmt_constant*, any subscript, substring starting point, or substring ending point must be an initialization expression.

Constraint: If a *data_stmt_constant* is a *structure_constructor*, each component must be an initialization expression.

Constraint: In a variable that is a *data_stmt_object*, any subscript, section subscript, substring starting point, and substring ending point must be an initialization expression.

Constraint: A variable whose name or designator is included in a *data_stmt_object_list* or a *data_i_do_object_list* must 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 array.

Constraint: In an *array_element* or a *scalar_structure_component* that is a *data_i_do_object*, any subscript must be an expression whose primaries are either constants, subobjects of constants, or DO variables of the containing *data_implied_do* elements, and each operation must be intrinsic.

Constraint: A *scalar_int_expr* of a *data_implied_do* must involve as primaries only constants, subobjects of constants, or DO variables of the containing *data_implied_dos*, and each operation must be intrinsic.

Constraint: The *scalar_structure_component* must contain at least one *part_ref* that contains a *subscript_list*.

EXT	<i>typeless_constant</i>	is	<i>octal_typeless_constant</i>
		or	<i>hexadecimal_typeless_constant</i>
		or	<i>binary_typeless_constant</i>
EXT	<i>octal_typeless_constant</i>	is	<i>digit</i> [<i>digit</i>] ... B
		or	O" <i>digit</i> [<i>digit</i>] ... "
		or	O' <i>digit</i> [<i>digit</i>] ... '
		or	" <i>digit</i> [<i>digit</i>] ... "O
		or	' <i>digit</i> [<i>digit</i>] ... 'O
EXT	<i>hexadecimal_typeless_constant</i>	is	X' <i>hex_digit</i> [<i>hex_digit</i>] ... '
		or	X" <i>hex_digit</i> [<i>hex_digit</i>] ... "
		or	' <i>hex_digit</i> [<i>hex_digit</i>] ... 'X
		or	" <i>hex_digit</i> [<i>hex_digit</i>] ... "X
		or	Z' <i>hex_digit</i> [<i>hex_digit</i>] ... '
		or	Z" <i>hex_digit</i> [<i>hex_digit</i>] ... "
EXT	<i>binary_typeless_constant</i>	is	B' <i>bin_digit</i> [<i>bin_digit</i>] ... '
		or	B" <i>bin_digit</i> [<i>bin_digit</i>] ... "

The following notes pertain to the definitions for *typeless_constant*, *octal_typeless_constant*, *hexadecimal_typeless_constant*, and *binary_typeless_constant*:

- *digit* must have one of the values 0 through 7 in *octal_typeless_constant*
- *digit* must have a value of 0 or 1 in *binary_typeless_constant*
- The B, O, X, and Z characters can be in uppercase or lowercase.

R530	<i>parameter_stmt</i>	is	PARAMETER (<i>named_constant_def_list</i>)
R531	<i>named_constant_def</i>	is	<i>named_constant</i> = <i>initialization_expr</i>
R541	<i>implicit_stmt</i>	is	IMPLICIT <i>implicit_spec_list</i>
		or	IMPLICIT NONE
EXT		or	IMPLICIT UNDEFINED

R542	<i>implicit_spec</i>	is	<i>type_spec</i> (<i>letter_spec_list</i>)
R543	<i>letter_spec</i>	is	<i>letter</i> [- <i>letter</i>]

Constraint: If `IMPLICIT NONE` is specified in a scoping unit, it must precede any `PARAMETER` statements that appear in the scoping unit and there must be no other `IMPLICIT` statements in the scoping unit.

Constraint: If the minus and second letter appear, the second letter must follow the first letter alphabetically.

R544	<i>namelist_stmt</i>	is	<code>NAMELIST</code> / <i>namelist_group_name</i> / <i>namelist_group_object_list</i> [[,] / <i>namelist_group_name</i> / <i>namelist_group_object_list</i>] ...
R545	<i>namelist_group_object</i>	is	<i>variable_name</i>

Constraint: A *namelist_group_object* must not be an array dummy argument with a nonconstant bound, a variable with nonconstant character length, an automatic object, a pointer, a variable of a type that has an ultimate component that is a pointer, or an allocatable array.

Constraint: If a *namelist_group_name* has the `PUBLIC` attribute, no item in the *namelist_group_object_list* can have the `PRIVATE` attribute or have private components.

Constraint: The *namelist_group_name* must not be a name made accessible by `USE` association.

R546	<i>equivalence_stmt</i>	is	<code>EQUIVALENCE</code> <i>equivalence_set_list</i>
R547	<i>equivalence_set</i>	is	(<i>equivalence_object</i> , <i>equivalence_object_list</i>)
R548	<i>equivalence_object</i>	is	<i>variable_name</i>
		or	<i>array_element</i>
		or	<i>substring</i>

Constraint: An *equivalence_object* must not be a dummy argument, a pointer, an allocatable array, an object of a nonsequence derived type or of a sequence derived type containing a pointer at any level of component selection, an

automatic object, a function name, an entry name, a result name, a named constant, a structure component, or a subobject of any of the preceding objects.

Constraint: An *equivalence_object* must not have the TARGET attribute.

Constraint: Each subscript or substring range expression in an *equivalence_object* must be an integer initialization expression.

Constraint: If an *equivalence_object* is of type default integer, default real, double-precision real, default complex, default logical, or numeric sequence type, all of the objects in the equivalence set must be of these types.

Constraint: If an *equivalence_object* is of type default character or character sequence type, all of the objects in the equivalence set must be of these types.

Constraint: If an *equivalence_object* is of a derived type that is not a numeric sequence or character sequence type, all of the objects in the equivalence set must be of the same type.

Constraint: 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 must be of the same type with the same kind type parameter value.

Constraint: The name of an *equivalence_object* must not be a name made accessible by USE association.

Constraint: A *substring* must not have length zero.

R549	<i>common_stmt</i>	is	COMMON [/ [<i>common_block_name</i>] /] <i>common_block_object_list</i> [[,] / [<i>common_block_name</i>] / <i>common_block_object_list</i>] . . .
R550	<i>common_block_object</i>	is	<i>variable_name</i> [(<i>explicit_shape_spec_list</i>)]

Constraint: Only one appearance of a given *variable_name* is permitted in all *common_block_object_lists* within a scoping unit.

Constraint: A *common_block_object* must not be a dummy argument, an allocatable array, an automatic object, a function name, an entry name, or a result name.

Constraint: Each bound in the *explicit_shape_spec* must be a constant specification expression.

Constraint: If a *common_block_object* is of a derived type, it must be a sequence type with no default initialization.

Constraint: If a *variable_name* appears with an *explicit_shape_spec_list*, it must not have the `POINTER` attribute.

Constraint: A variable name must not be a name made accessible by `USE` association.

1.2.6 Use of Data Objects

The following syntax rules are described in section 6, "Use of data objects," of the Fortran 95 standard.

R601	<i>variable</i>	is	<i>scalar_variable_name</i>
		or	<i>array_variable_name</i>
		or	<i>subobject</i>

Constraint: *array_variable_name* must be the name of a data object that is an array.

Constraint: *array_variable_name* must not have the `PARAMETER` attribute.

Constraint: *scalar_variable_name* must not have the `PARAMETER` attribute.

Constraint: *subobject* must not be a subobject designator (for example, a substring) whose parent is a constant.

R602	<i>subobject</i>	is	<i>array_element</i>
		or	<i>array_section</i>
		or	<i>structure_component</i>
		or	<i>substring</i>
R603	<i>logical_variable</i>	is	<i>variable</i>

Constraint: *logical_variable* must be of type logical.

R604	<i>default_logical_variable</i>	is	<i>variable</i>
------	---------------------------------	-----------	-----------------

Constraint: *default_logical_variable* must be of type default logical.

R605	<i>char_variable</i>	is	<i>variable</i>
------	----------------------	-----------	-----------------

Constraint: *char_variable* must be of type character.

R606	<i>default_char_variable</i>	is	<i>variable</i>
------	------------------------------	-----------	-----------------

Constraint: *default_char_variable* must be of type default character.

R607	<i>int_variable</i>	is	<i>variable</i>
------	---------------------	-----------	-----------------

Constraint: *int_variable* must be of type integer.

R608	<i>default_int_variable</i>	is	<i>variable</i>
------	-----------------------------	-----------	-----------------

Constraint: *default_int_variable* must be of type default integer.

R609	<i>substring</i>	is	<i>parent_string</i> (<i>substring_range</i>)
R610	<i>parent_string</i>	is	<i>scalar_variable_name</i>
		or	<i>array_element</i>
		or	<i>scalar_structure_component</i>
		or	<i>scalar_constant</i>
R611	<i>substring_range</i>	is	[<i>scalar_int_expr</i>] : [<i>scalar_int_expr</i>]

Constraint: *parent_string* must be of type character.

R612	<i>data_ref</i>	is	<i>part_ref</i> [% <i>part_ref</i>] . . .
------	-----------------	-----------	---

R613	<i>part_ref</i>	is	<i>part_name</i> [(<i>section_subscript_list</i>)]
------	-----------------	-----------	--

Constraint: In a *data_ref*, each *part_name* except the rightmost must be of derived type.

Constraint: In a *data_ref*, each *part_name* except the leftmost must be the name of a component of the derived type definition of the type of the preceding *part_name*.

Constraint: In a *part_ref* containing a *section_subscript_list*, the number of *section_subscripts* must equal the rank of *part_name*.

Constraint: In a *data_ref*, there must not be more than one *part_ref* with nonzero rank. A *part_name* to the right of a *part_ref* with nonzero rank must not have the `POINTER` attribute.

R614	<i>structure_component</i>	is	<i>data_ref</i>
------	----------------------------	-----------	-----------------

Constraint: In a *structure_component*, there must be more than one *part_ref* and the rightmost *part_ref* must be of the form *part_name*.

R615	<i>array_element</i>	is	<i>data_ref</i>
------	----------------------	-----------	-----------------

Constraint: In an *array_element*, every *part_ref* must have rank zero and the last *part_ref* must contain a *subscript_list*.

R616	<i>array_section</i>	is	<i>data_ref</i> [(<i>substring_range</i>)]
------	----------------------	-----------	--

Constraint: In an *array_section*, exactly one *part_ref* must have nonzero rank, and either the final *part_ref* has a *section_subscript_list* with nonzero rank or another *part_ref* must have nonzero rank.

Constraint: In an *array_section* with a *substring_range*, the rightmost *part_name* must be of type character.

R617	<i>subscript</i>	is	<i>scalar_int_expr</i>
R618	<i>section_subscript</i>	is	<i>subscript</i>
		or	<i>subscript_triplet</i>
		or	<i>vector_subscript</i>
R619	<i>subscript_triplet</i>	is	[<i>subscript</i>] : [<i>subscript</i>] [: <i>stride</i>]
R620	<i>stride</i>	is	<i>scalar_int_expr</i>
R621	<i>vector_subscript</i>	is	<i>int_expr</i>

Constraint: A *vector_subscript* must be an integer array expression of rank one.

Constraint: The second subscript must not be omitted from a *subscript_triplet* in the last dimension of an assumed-size array.

R622	<i>allocate_stmt</i>	is	ALLOCATE (<i>allocation_list</i> [, STAT = <i>stat_variable</i>])
R623	<i>stat_variable</i>	is	<i>scalar_int_variable</i>
R624	<i>allocation</i>	is	<i>allocate_object</i> [(<i>allocate_shape_spec_list</i>)]
R625	<i>allocate_object</i>	is	<i>variable_name</i>
		or	<i>structure_component</i>
R626	<i>allocate_shape_spec</i>	is	[<i>allocate_lower_bound</i> :] <i>allocate_upper_bound</i>
R627	<i>allocate_lower_bound</i>	is	<i>scalar_int_expr</i>
R628	<i>allocate_upper_bound</i>	is	<i>scalar_int_expr</i>

Constraint: Each *allocate_object* must be a pointer or an allocatable array.

Constraint: The number of *allocate_shape_specs* in an *allocate_shape_spec_list* must be the same as the rank of the pointer or allocatable array.

R629	<i>nullify_stmt</i>	is	NULLIFY (<i>pointer_object_list</i>)
R630	<i>pointer_object</i>	is	<i>variable_name</i>
		or	<i>structure_component</i>

Constraint: Each *pointer_object* must have the POINTER attribute.

R631	<i>deallocate_stmt</i>	is	DEALLOCATE (<i>allocate_object_list</i> [, STAT = <i>stat_variable</i>])
------	------------------------	-----------	---

Constraint: Each *allocate_object* must be a pointer or an allocatable array.

1.2.7 Expressions and Assignment

The following syntax rules are described in section 7, "Expressions and assignment," of the Fortran 95 standard.

Note: The language of the Fortran 95 standard is presented in this subsection in its original form. Chapter 7 of the *Fortran Language Reference Manual, Volume 1*, however, sometimes uses terms that are different from those found in the standard. The terminology was changed to improve clarity. The following list shows the terms used in this compiler manual set and the equivalent term used in the Fortran 95 standard.

<u>Standard</u>	<u>Silicon Graphics term</u>
<i>level_1_expr</i>	<i>defined_unary_expr</i>
<i>defined_unary_op</i>	<i>defined_operator</i>
<i>mult_operand</i>	<i>exponentiation_expr</i>
<i>power_op</i>	**
<i>add_operand</i>	<i>multiplication_expr</i>
<i>mult_op</i>	* or /
<i>level_2_expr</i>	<i>summation_expr</i>
<i>add_op</i>	+ or -
<i>level_3_expr</i>	<i>concatenation_expr</i>
<i>concat_op</i>	//
<i>level_4_expr</i>	<i>comparison_expr</i>
<i>rel_op</i>	<i>rel_op</i>
<i>and_operand</i>	<i>not_expr</i>
<i>not_op</i>	.NOT.

<i>or_operand</i>	<i>conjunct_expr</i>
<i>and_op</i>	.AND.
<i>or_op</i>	.OR.
<i>equiv_operand</i>	<i>inclusive_disjunct_expr</i>
<i>level_5_expr</i>	<i>equivalence_expr</i>
<i>mask_expr</i>	<i>logical_expr</i>

R701	<i>primary</i>	is	<i>constant</i>
		or	<i>constant_subobject</i>
		or	<i>variable</i>
		or	<i>array_constructor</i>
		or	<i>structure_constructor</i>
		or	<i>function_reference</i>
		or	(<i>expr</i>)
R702	<i>constant_subobject</i>	is	<i>subobject</i>

Constraint: *subobject* must be a subobject designator whose parent is a constant. A *variable* that is a primary must not be an assumed-size array.

R703	<i>level_1_expr</i>	is	[<i>defined_unary_op</i>] <i>primary</i>
R704	<i>defined_unary_op</i>	is	. <i>letter</i> [<i>letter</i>]

Constraint: A *defined_unary_op* must not contain more than 31 letters and must not be the same as any *intrinsic_operator* or *logical_literal_constant*.

R705	<i>mult_operand</i>	is	<i>level_1_expr</i> [<i>power_op mult_operand</i>]
R706	<i>add_operand</i>	is	[<i>add_operand mult_op</i>] <i>mult_operand</i>
R707	<i>level_2_expr</i>	is	[[<i>level_2_expr</i>] <i>add_op</i>] <i>add_operand</i>

R708	<i>power_op</i>	is	**	
R709	<i>mult_op</i>	is	*	
		or	/	
R710	<i>add_op</i>	is	+	
		or	-	
R711	<i>level_3_expr</i>	is	[<i>level_3_expr concat_op</i>] <i>level_2_expr</i>	
R712	<i>concat_op</i>	is	//	
R713	<i>level_4_expr</i>	is	[<i>level_3_expr rel_op</i>] <i>level_3_expr</i>	
R714	<i>rel_op</i>	is	.EQ.	
		or	.NE.	
		or	.LT.	
		or	.LE.	
		or	.GT.	
		or	.GE.	
		EXT	or	.LG.
		or	==	
		or	/=	
		or	<	
		or	<=	
		or	>	
		or	>=	
		EXT	or	<>
R715	<i>and_operand</i>	is	[<i>not_op</i>] <i>level_4_expr</i>	
R716	<i>or_operand</i>	is	[<i>or_operand and_op</i>] <i>and_operand</i>	
R717	<i>equiv_operand</i>	is	[<i>equiv_operand or_op</i>] <i>or_operand</i>	
R718	<i>level_5_expr</i>	is	[<i>level_5_expr equiv_op</i>] <i>equiv_operand</i>	
R719	<i>not_op</i>	is	.NOT.	
R720	<i>and_op</i>	is	.AND.	
R721	<i>or_op</i>	is	.OR.	

R722	<i>equiv_op</i>	is	<i>.EQV.</i>
		or	<i>.NEQV.</i>
R723	<i>expr</i>	is	[<i>expr defined_binary_op</i>] <i>level_5_expr</i>
R724	<i>defined_binary_op</i>	is	<i>. letter [letter]</i>

Constraint: A *defined_binary_op* must not contain more than 31 letters and must not be the same as any *intrinsic_operator* or *logical_literal_constant*.

R725	<i>logical_expr</i>	is	<i>expr</i>
------	---------------------	-----------	-------------

Constraint: *logical_expr* must be type logical.

R726	<i>char_expr</i>	is	<i>expr</i>
------	------------------	-----------	-------------

Constraint: *char_expr* must be type character.

R727	<i>default_char_expr</i>	is	<i>expr</i>
------	--------------------------	-----------	-------------

Constraint: *default_char_expr* must be of type default character.

R728	<i>int_expr</i>	is	<i>expr</i>
------	-----------------	-----------	-------------

Constraint: *int_expr* must be type integer.

R729	<i>numeric_expr</i>	is	<i>expr</i>
------	---------------------	-----------	-------------

Constraint: *numeric_expr* must be of type integer, real, or complex.

R730	<i>initialization_expr</i>	is	<i>expr</i>
------	----------------------------	-----------	-------------

Constraint: An *initialization_expr* must be an initialization expression.

R731	<i>char_initialization_expr</i>	is	<i>char_expr</i>
------	---------------------------------	-----------	------------------

Constraint: A *char_initialization_expr* must be an initialization expression.

R732	<i>int_initialization_expr</i>	is	<i>int_expr</i>
------	--------------------------------	-----------	-----------------

Constraint: An *int_initialization_expr* must be an initialization expression.

R733	<i>logical_initialization_expr</i>	is	<i>logical_expr</i>
------	------------------------------------	-----------	---------------------

Constraint: A *logical_initialization_expr* must be an initialization expression.

R734	<i>specification_expr</i>	is	<i>scalar_int_expr</i>
------	---------------------------	-----------	------------------------

Constraint: The *scalar_int_expr* must be a restricted expression.

R735	<i>assignment_stmt</i>	is	<i>variable = expr</i>
------	------------------------	-----------	------------------------

Constraint: A variable in an *assignment_stmt* must not be an assumed-size array.

R736	<i>pointer_assignment_stmt</i>	is	<i>pointer_object => target</i>
------	--------------------------------	-----------	------------------------------------

R737	<i>target</i>	is	<i>variable</i>
------	---------------	-----------	-----------------

		or	<i>expr</i>
--	--	-----------	-------------

Constraint: The *pointer_object* must have the POINTER attribute.

Constraint: The variable must have the `TARGET` attribute or be a subobject of an object with the `TARGET` attribute, or it must have the `POINTER` attribute.

Constraint: The *target* must be of the same type, kind type parameters, and rank as the pointer.

Constraint: The *target* must not be an array section with a vector subscript.

Constraint: The *expr* must deliver a pointer result.

R738	<i>where_stmt</i>	is	<code>WHERE (<i>mask_expr</i>) <i>where_assignment_stmt</i></code>
R739	<i>where_construct</i>	is	<code><i>where_construct_stmt</i> [<i>where_body_construct</i>] ... [<i>masked_elsewhere_stmt</i> [<i>where_body_construct</i>] ...] ... [<i>elsewhere_stmt</i> [<i>where_body_construct</i>] ...] <i>end_where_stmt</i></code>
R740	<i>where_construct_stmt</i>	is	<code>[<i>where_construct_name</i>:]WHERE (<i>mask_expr</i>)</code>
R741	<i>where_construct_stmt</i>	is	<code><i>where_assignment_stmt</i></code>
		or	<code><i>where_stmt</i></code>
		or	<code><i>where_construct</i></code>
R742	<i>where_assignment_stmt</i>	is	<code><i>assignment_stmt</i></code>
R743	<i>mask_expr</i>	is	<code><i>logical_expr</i></code>
R742	<i>masked_elsewhere_stmt</i>	is	<code>ELSEWHERE (<i>mask_expr</i>) [<i>where_construct_name</i>]</code>
R745	<i>elsewhere_stmt</i>	is	<code>ELSEWHERE [<i>where_construct_name</i>]</code>
R746	<i>end_where_stmt</i>	is	<code>END WHERE [<i>where_construct_name</i>]</code>

Constraint: A *where_assignment_stmt* that is a defined assignment must be elemental.

Constraint: If the *where_construct_stmt* is identified by a *where_construct_name*, the corresponding *end_where_stmt* must specify the same *where_construct_name*. If the *where_construct_stmt* is not identified by a *where_construct_name*, the corresponding *end_where_stmt* must 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* must specify the same *where_construct_name*.

R747	<i>forall_construct</i>	is	<i>forall_construct_stmt</i> [<i>forall_body_construct</i>] ... <i>end_forall_stmt</i>
R748	<i>forall_construct_stmt</i>	is	[<i>forall_construct_name</i> :]FORALL <i>forall_header</i>
R749	<i>forall_header</i>	is	(<i>forall_triplet_spec_list</i> [, <i>scalar_mask_expr</i>])
R750	<i>forall_triplet_spec</i>	is	<i>index_name</i> = <i>subscript</i> : <i>subscript</i> [: <i>stride</i>]
R617	<i>subscript</i>	is	<i>scalar_int_expr</i>
R603	<i>stride</i>	is	<i>scalar_int_expr</i>
R751	<i>forall_body_construct</i>	is	<i>forall_assignment_stmt</i>
R751	<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>
R752	<i>forall_assignment_stmt</i>	is	<i>assignment_stmt</i> or <i>pointer_assignment_stmt</i>
R753	<i>end_forall_stmt</i>	is	END FORALL [<i>forall_construct_name</i>]

Constraint: If the *forall_construct_stmt* has a *forall_construct_name*, the *end_forall_stmt* must have the same *forall_construct_name*. If the *end_forall_stmt* has a *forall_construct_name*, the *forall_construct_stmt* must have the same *forall_construct_name*.

Constraint: The *scalar_mask_expr* must be scalar and of type logical.

Constraint: Any procedure referenced in the *scalar_mask_expr*, including one referenced by a defined operation, must be a pure procedure.

Constraint: The *index_name* must be a named scalar variable of type integer.

Constraint: A *subscript* or *stride* in a *forall_triplet_spec* must not contain a reference to any *index_name* in the *forall_triplet_spec_list* in which it appears.

Constraint: A statement in a *forall_body_construct* must not define an *index_name* of the *forall_construct*.

Constraint: Any procedure referenced in a *forall_body_construct*, including one referenced by a defined operation or assignment, must be a pure procedure.

Constraint: A *forall_body_construct* must not be a branch target.

R754	<i>forall_stmt</i>	is	FORALL <i>forall_header forall_assignment_stmt</i>
------	--------------------	-----------	--

1.2.8 Execution Control

The following syntax rules are described in section 8, "Execution control," of the Fortran 95 standard.

R801	<i>block</i>	is	[<i>execution_part_construct</i>] . . .
R802	<i>if_construct</i>	is	<i>if_then_stmt</i> <i>block</i> [<i>else_if_stmt</i> <i>block</i>] . . . [<i>else_stmt</i> <i>block</i>] <i>end_if_stmt</i>
R803	<i>if_then_stmt</i>	is	[<i>if_construct_name</i> :] IF (<i>scalar_logical_expr</i>) THEN
R804	<i>else_if_stmt</i>	is	ELSE IF (<i>scalar_logical_expr</i>) THEN [<i>if_construct_name</i>]
R805	<i>else_stmt</i>	is	ELSE [<i>if_construct_name</i>]
R806	<i>end_if_stmt</i>	is	END IF [<i>if_construct_name</i>]

Constraint: If the *if_then_stmt* of an *if_construct* is identified by an *if_construct_name*, the corresponding *end_if_stmt* must specify the same *if_construct_name*. If the *if_then_stmt* of an *if_construct* is not identified by an *if_construct_name*, the corresponding *end_if_stmt* must not specify an *if_construct_name*. If an *else_if_stmt* or *else_stmt* is identified by an *if_construct_name*, the corresponding *if_then_stmt* must specify the same *if_construct_name*.

R807	<i>if_stmt</i>	is	IF (<i>scalar_logical_expr</i>) <i>action_stmt</i>
------	----------------	-----------	--

Constraint: The *action_stmt* in the *if_stmt* must not be an *if_stmt*, *end_program_stmt*, *end_function_stmt*, or *end_subroutine_stmt*.

R808	<i>case_construct</i>	is	<i>select_case_stmt</i> [<i>case_stmt</i> <i>block</i>] . . . <i>end_select_stmt</i>
------	-----------------------	-----------	---

R809	<i>select_case_stmt</i>	is	[<i>case_construct_name</i>] : SELECT CASE (<i>case_expr</i>)
------	-------------------------	-----------	---

R810	<i>case_stmt</i>	is	CASE <i>case_selector</i> [<i>case_construct_name</i>]
------	------------------	-----------	--

R811	<i>end_select_stmt</i>	is	END SELECT [<i>case_construct_name</i>]
------	------------------------	-----------	---

Constraint: If the *select_case_stmt* of a *case_construct* is identified by a *case_construct_name*, the corresponding *end_select_stmt* must specify the same *case_construct_name*. If the *select_case_stmt* of a *case_construct* is not identified by a *case_construct_name*, the corresponding *end_select_stmt* must not specify a *case_construct_name*. If a *case_stmt* is identified by a *case_construct_name*, the corresponding *select_case_stmt* must specify the same *case_construct_name*.

R812	<i>case_expr</i>	is	<i>scalar_int_expr</i>
------	------------------	-----------	------------------------

or	<i>scalar_char_expr</i>
-----------	-------------------------

or	<i>scalar_logical_expr</i>
-----------	----------------------------

R813	<i>case_selector</i>	is	(<i>case_value_range_list</i>)
------	----------------------	-----------	----------------------------------

or	DEFAULT
-----------	---------

Constraint: No more than one of the selectors of one of the CASE statements may be DEFAULT.

R814	<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>
R815	<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>

Constraint: For a given *case_construct*, each *case_value* must be of the same type as *case_expr*. For character type, length differences are allowed, but the kind type parameters must be the same.

Constraint: A *case_value_range* using a colon must not be used if *case_expr* is of type logical.

Constraint: For a given *case_construct*, the *case_value_ranges* must not overlap; that is, there must be no possible value of the *case_expr* that matches more than one *case_value_range*.

R816	<i>do_construct</i>	is	<i>block_do_construct</i>
		or	<i>nonblock_do_construct</i>
R817	<i>block_do_construct</i>	is	<i>do_stmt</i> <i>do_block</i> <i>end_do</i>
R818	<i>do_stmt</i>	is	<i>label_do_stmt</i>
		or	<i>nonlabel_do_stmt</i>
R819	<i>label_do_stmt</i>	is	[<i>do_construct_name</i> :] DO <i>label</i> [<i>loop_control</i>]
R820	<i>nonlabel_do_stmt</i>	is	[<i>do_construct_name</i> :] DO [<i>loop_control</i>]
R821	<i>loop_control</i>	is	[,] <i>do_variable</i> = <i>scalar_int_expr</i> ,
		or	<i>scalar_int_expr</i> [, <i>scalar_int_expr</i>]
			[,] WHILE (<i>scalar_logical_expr</i>)
R822	<i>do_variable</i>	is	<i>scalar_int_variable</i>

Constraint: The *do_variable* must be a named scalar variable of type integer, default real, or double-precision real. The Fortran standard does not allow for *do_variables* of type default real or double-precision real.

Constraint: Each *scalar_numeric_expr* in *loop_control* must be of type integer, default real, or double-precision real. The Fortran standard does not allow for *do_variables* of type default real or double-precision real.

R823	<i>do_block</i>	is	<i>block</i>
R824	<i>end_do</i>	is	<i>end_do_stmt</i>
		or	<i>continue_stmt</i>
R825	<i>end_do_stmt</i>	is	END DO [<i>do_construct_name</i>]

Constraint: If the *do_stmt* of a *block_do_construct* is identified by a *do_construct_name*, the corresponding *end_do* must be an *end_do_stmt* specifying the same *do_construct_name*. If the *do_stmt* of a *block_do_construct* is not identified by a *do_construct_name*, the corresponding *end_do* must not specify a *do_construct_name*.

Constraint: If the *do_stmt* is a *nonlabel_do_stmt*, the corresponding *end_do* must be an *end_do_stmt*.

Constraint: If the *do_stmt* is a *label_do_stmt*, the corresponding *end_do* must be identified with the same label.

OBS	<i>nonblock_do_construct</i>	is	<i>action_term_do_construct</i>
		or	<i>outer_shared_do_construct</i>
OBS	<i>action_term_do_construct</i>	is	<i>label_do_stmt</i>
			<i>do_body</i>
			<i>do_term_action_stmt</i>
OBS	<i>do_body</i>	is	[<i>execution_part_construct</i>] . . .
OBS	<i>do_term_action_stmt</i>	is	<i>action_stmt</i>

Obsolescent Constraint: A *do_term_action_stmt* must 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*.

Obsolescent Constraint: The *do_term_action_stmt* must be identified with a label and the corresponding *label_do_stmt* must refer to the same label.

OBS	<i>outer_shared_do_construct</i>	is	<i>label_do_stmt</i> <i>do_body</i> <i>shared_term_do_construct</i>
OBS	<i>shared_term_do_construct</i>	is or	<i>outer_shared_do_construct</i> <i>inner_shared_do_construct</i>
OBS	<i>inner_shared_do_construct</i>	is	<i>label_do_stmt</i> <i>do_body</i> <i>do_term_shared_stmt</i>
OBS	<i>do_term_shared_stmt</i>	is	<i>action_stmt</i>

Obsolescent Constraint: A *do_term_shared_stmt* must 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*.

Obsolescent Constraint: The *do_term_shared_stmt* must be identified with a label, and all of the *label_do_stmts* of the *shared_term_do_construct* must refer to the same label.

R834	<i>cycle_stmt</i>	is	CYCLE [<i>do_construct_name</i>]
------	-------------------	-----------	------------------------------------

Constraint: If a *cycle_stmt* refers to a *do_construct_name*, it must be within the range of that *do_construct*; otherwise, it must be within the range of at least one *do_construct*.

R835	<i>exit_stmt</i>	is	EXIT [<i>do_construct_name</i>]
------	------------------	-----------	-----------------------------------

Constraint: If an *exit_stmt* refers to a *do_construct_name*, it must be within the range of that *do_construct*; otherwise, it must be within the range of at least one *do_construct*.

R836	<i>goto_stmt</i>	is	GO TO <i>label</i>
------	------------------	-----------	--------------------

Constraint: The *label* must be the statement label of a branch target statement that appears in the same scoping unit as the *goto_stmt*.

OBS	<i>computed_goto_stmt</i>	is	GO TO (<i>label_list</i>) [,] <i>scalar_int_expr</i>
-----	---------------------------	-----------	--

Obsolescent Constraint: Each *label* in *label_list* must be the statement label of a branch target statement that appears in the same scoping unit as the *computed_goto_stmt*.

EXT	<i>assign_stmt</i>	is	ASSIGN <i>label</i> TO <i>scalar_int_variable</i>
-----	--------------------	-----------	---

Extension Constraint: The *label* must be the statement label of a branch target statement or *format_stmt* that appears in the same scoping unit as the *assign_stmt*.

Extension Constraint: *scalar_int_variable* must be named and of type default integer.

EXT	<i>assigned_goto_stmt</i>	is	GO TO <i>scalar_int_variable</i> [[,] (<i>label_list</i>)]
-----	---------------------------	-----------	--

Extension Constraint: Each *label* in *label_list* must be the statement label of a branch target statement that appears in the same scoping unit as the *assigned_goto_stmt*.

Extension Constraint: *scalar_int_variable* must be named and of type default integer.

OBS	<i>arithmetic_if_stmt</i>	is	IF (<i>scalar_numeric_expr</i>) <i>label</i> , <i>label</i> , <i>label</i>
-----	---------------------------	-----------	--

Obsolescent Constraint: Each *label* must be the label of a branch target statement that appears in the same scoping unit as the *arithmetic_if_stmt*.

Obsolescent Constraint: The *scalar_numeric_expr* must not be of type complex.

R839	<i>continue_stmt</i>	is	CONTINUE
R840	<i>stop_stmt</i>	is	STOP [<i>stop_code</i>]
R841	<i>stop_code</i>	is	<i>scalar_char_constant</i>
EXT		or	<i>digit</i> [<i>digit</i> [<i>digit</i> [<i>digit</i> [<i>digit</i>]]]

Constraint: *scalar_char_constant* must be of type default character.

EXT	<i>pause_stmt</i>	is	PAUSE [<i>stop_code</i>]
-----	-------------------	-----------	----------------------------

1.2.9 Input/Output (I/O) Statements

The following syntax rules are described in section 9, "Input/Output statements," of the Fortran 95 standard.

R901	<i>io_unit</i>	is	<i>external_file_unit</i>
		or	*
		or	<i>internal_file_unit</i>
EXT		or	<i>unit_name</i>
R902	<i>external_file_unit</i>	is	<i>scalar_int_expr</i>
R903	<i>internal_file_unit</i>	is	<i>default_char_variable</i>

Constraint: The *default_char_variable* must not be an array section with a vector subscript.

R904	<i>open_stmt</i>	is	OPEN (<i>connect_spec_list</i>)
R905	<i>connect_spec</i>	is	[UNIT =] <i>external_file_unit</i>
		or	IOSTAT = <i>scalar_default_int_variable</i>
		or	ERR = <i>label</i>
		or	FILE = <i>file_name_expr</i>
		or	STATUS = <i>scalar_char_expr</i>
		or	ACCESS = <i>scalar_char_expr</i>
		or	FORM = <i>scalar_char_expr</i>
		or	RECL = <i>scalar_int_expr</i>
		or	BLANK = <i>scalar_char_expr</i>
		or	POSITION = <i>scalar_char_expr</i>
		or	ACTION = <i>scalar_char_expr</i>
		or	DELIM = <i>scalar_char_expr</i>
		or	PAD = <i>scalar_char_expr</i>
R906	<i>file_name_expr</i>	is	<i>scalar_char_expr</i>

Constraint: If the optional characters UNIT= are omitted from the unit specifier, the unit specifier must be the first item in the *connect_spec_list*.

Constraint: Each specifier must not appear more than once in a given *open_stmt*; an *external_file_unit* must be specified.

Constraint: The *label* used in the ERR= specifier must be the statement label of a branch target statement that appears in the same scoping unit as the OPEN statement.

R907	<i>close_stmt</i>	is	CLOSE (<i>close_spec_list</i>)
R908	<i>close_spec</i>	is	[UNIT =] <i>external_file_unit</i>
		or	IOSTAT = <i>scalar_default_int_variable</i>
		or	ERR = <i>label</i>
		or	STATUS = <i>scalar_char_expr</i>

Constraint: If the optional characters UNIT= are omitted from the unit specifier, the unit specifier must be the first item in the *close_spec_list*.

Constraint: Each specifier must not appear more than once in a given *close_stmt*; an *external_file_unit* must be specified.

Constraint: The label used in the ERR= specifier must be the statement label of a branch target statement that appears in the same scoping unit as the CLOSE statement.

R909	<i>read_stmt</i>	is	READ (<i>io_control_spec_list</i>) [<i>input_item_list</i>]
EXT		or	READ <i>format</i> [, <i>input_item_list</i>]
R910	<i>write_stmt</i>	is	WRITE (<i>io_control_spec_list</i>) [<i>output_item_list</i>]
EXT		or	WRITE <i>format</i> [, <i>output_item_list</i>]
R911	<i>print_stmt</i>	is	PRINT <i>format</i> [, <i>output_item_list</i>]
R912	<i>io_control_spec</i>	is	[UNIT =] <i>io_unit</i>
		or	[FMT =] <i>format</i>
		or	[NML =] <i>namelist_group_name</i>
		or	REC = <i>scalar_int_expr</i>
		or	IOSTAT = <i>scalar_default_int_variable</i>
		or	ERR = <i>label</i>
		or	END = <i>label</i>
		or	ADVANCE = <i>scalar_default_char_expr</i>
		or	SIZE = <i>scalar_default_int_variable</i>
		or	EOR = <i>label</i>

Constraint: An *io_control_spec_list* must contain exactly one *io_unit* and may contain at most one of each of the other specifiers.

Constraint: An END=, EOR=, or SIZE= specifier must not appear in a *write_stmt*.

Constraint: The label in the ERR=, EOR=, or END= specifier must be the statement label of a branch target statement that appears in the same scoping unit as the data transfer statement.

Constraint: A *namelist_group_name* must not be present if an *input_item_list* or an *output_item_list* is present in the data transfer statement.

Constraint: An *io_control_spec_list* must not contain both a format and a *namelist_group_name*.

Constraint: If the optional characters `UNIT=` are omitted from the unit specifier, the unit specifier must be the first item in the control information list.

Constraint: If the optional characters `FMT=` are omitted from the format specifier, the format specifier must be the second item in the control information list and the first item must be the unit specifier without the optional characters `UNIT=`.

Constraint: If the optional characters `NML=` are omitted from the namelist specifier, the namelist specifier must be the second item in the control information list and the first item must be the unit specifier without the optional characters `UNIT=`.

Constraint: If the unit specifier specifies an internal file, the *io_control_spec_list* must not contain a `REC=` specifier or a *namelist_group_name*.

Constraint: If the `REC=` specifier is present, an `END=` specifier must not appear, a *namelist_group_name* must not appear, and the format, if any, must not be an asterisk specifying *list_directed* I/O.

Constraint: An `ADVANCE=` specifier may be present only in a formatted sequential I/O statement with explicit format specification whose control information list does not contain an internal file unit specifier.

Constraint: If an `EOR=` specifier is present, an `ADVANCE=` specifier also must appear.

Constraint: If a `SIZE=` specifier is present, an `ADVANCE=` specifier must also appear.

R913	<i>format</i>	is	<i>default_char_expr</i>
		or	<i>label</i>
		or	*
EXT		or	<i>scalar_default_int_variable</i>

Constraint: The *label* must be the label of a `FORMAT` statement that appears in the same scoping unit as the statement containing the format specifier.

R914	<i>input_item</i>	is	<i>variable</i> or <i>io_implied_do</i>
R915	<i>output_item</i>	is	<i>expr</i> or <i>io_implied_do</i>
R916	<i>io_implied_do</i>	is	(<i>io_implied_do_object_list</i> , <i>io_implied_do_control</i>)
R917	<i>io_implied_do_object</i>	is	<i>input_item</i> or <i>output_item</i>
R918	<i>io_implied_do_control</i>	is	<i>do_variable</i> = <i>scalar_int_expr</i> , <i>scalar_int_expr</i> [, <i>scalar_int_expr</i>]

Constraint: A variable that is an *input_item* must not be an assumed-size array.

Constraint: The DO variable must be a named scalar of type integer.

Constraint: In an *input_item_list*, an *io_implied_do_object* must be an *input_item*.
In an *output_item_list*, an *io_implied_do_object* must be an *output_item*.

EXT	<i>buffer_in_stmt</i>	is	BUFFER IN (<i>io_unit</i> , <i>mode</i>) (<i>start_loc</i> , <i>end_loc</i>)
EXT	<i>buffer_out_stmt</i>	is	BUFFER OUT (<i>io_unit</i> , <i>mode</i>) (<i>start_loc</i> , <i>end_loc</i>)
EXT	<i>io_unit</i>	is	<i>external_file_unit</i> or <i>file_name_expr</i>
EXT	<i>mode</i>	is	<i>scalar_integer_expr</i>
EXT	<i>start_loc</i>	is	<i>variable</i>
EXT	<i>end_loc</i>	is	<i>variable</i>

In the preceding definition, the *variable* specified for *start_loc* and *end_loc* cannot be of a derived type if you are performing implicit data conversion. The data items between *start_loc* and *end_loc* must be of the same type and same kind type.

R919	<i>backspace_stmt</i>	is	BACKSPACE <i>external_file_unit</i>
		or	BACKSPACE (<i>position_spec_list</i>)
R920	<i>endfile_stmt</i>	is	ENDFILE <i>external_file_unit</i>
		or	ENDFILE (<i>position_spec_list</i>)
R921	<i>rewind_stmt</i>	is	REWIND <i>external_file_unit</i>
		or	REWIND (<i>position_spec_list</i>)
R922	<i>position_spec</i>	is	[UNIT =] <i>external_file_unit</i>
		or	IOSTAT = <i>scalar_default_int_variable</i>
		or	ERR = <i>label</i>

Constraint: The *label* in the ERR= specifier must be the statement label of a branch target statement that appears in the same scoping unit as the file positioning statement.

Constraint: If the optional characters UNIT= are omitted from the unit specifier; the unit specifier must be the first item in the *position_spec_list*.

Constraint: A *position_spec_list* must contain exactly one *external_file_unit* and may contain at most one of each of the other specifiers.

R923	<i>inquire_stmt</i>	is	INQUIRE (<i>inquire_spec_list</i>)
		or	INQUIRE (IOLENGTH = <i>scalar_default_int_variable</i>) <i>output_item_list</i>
R924	<i>inquire_spec</i>	is	[UNIT =] <i>external_file_unit</i>
		or	FILE = <i>file_name_expr</i>
		or	IOSTAT = <i>scalar_default_int_variable</i>
		or	ERR = <i>label</i>
		or	EXIST = <i>scalar_default_logical_variable</i>
		or	OPENED = <i>scalar_default_logical_variable</i>
		or	NUMBER = <i>scalar_default_int_variable</i>
		or	NAMED = <i>scalar_default_logical_variable</i>
		or	NAME = <i>scalar_default_char_variable</i>
		or	ACCESS = <i>scalar_default_char_variable</i>

```

or    SEQUENTIAL = scalar_default_char_variable
or    DIRECT = scalar_default_char_variable
or    FORM = scalar_default_char_variable
or    FORMATTED = scalar_default_char_variable
or    UNFORMATTED = scalar_default_char_variable
or    RECL = scalar_default_int_variable
or    NEXTREC = scalar_default_int_variable
or    BLANK = scalar_default_char_variable
or    POSITION = scalar_default_char_variable
or    ACTION = scalar_default_char_variable
or    READ = scalar_default_char_variable
or    WRITE = scalar_default_char_variable
or    READWRITE = scalar_default_char_variable
or    DELIM = scalar_default_char_variable
or    PAD = scalar_default_char_variable
    
```

Constraint: An *inquire_spec_list* must contain one FILE= specifier or one UNIT= specifier, but not both, and at most one of each of the other specifiers.

Constraint: In the inquire by unit form of the INQUIRE statement, if the optional characters UNIT= are omitted from the unit specifier, the unit specifier must be the first item in the *inquire_spec_list*.

1.2.10 I/O Editing

The following syntax rules are described in section 10, "Input/Output editing," of the Fortran 95 standard.

R1001	<i>format_stmt</i>	is	FORMAT <i>format_specification</i>
R1002	<i>format_specification</i>	is	([<i>format_item_list</i>])

Constraint: The *format_stmt* must be labeled.

Constraint: The comma used to separate *format_items* in a *format_item_list* may be omitted as follows:

- Between a P edit descriptor and an immediately following F, E, EN, ES, D, or G edit descriptor
- Before a slash edit descriptor when the optional repeat specification is not present
- After a slash edit descriptor
- Before or after a colon edit descriptor

R1003	<i>format_item</i>	is	[<i>r</i>] <i>data_edit_desc</i>
		or	<i>control_edit_desc</i>
		or	<i>char_string_edit_desc</i>
		or	[<i>r</i>] (<i>format_item_list</i>)
R1004	<i>r</i>	is	<i>int_literal_constant</i>

Constraint: *r* must be positive.

Constraint: *r* must not have kind parameter specified for it.

R1005	<i>data_edit_desc</i>	is	I <i>w</i> [. <i>m</i>]
		or	B <i>w</i> [. <i>m</i>]
		or	O <i>w</i> [. <i>m</i>]
		or	Z <i>w</i> [. <i>m</i>]
		or	F <i>w</i> . <i>d</i>
		or	E <i>w</i> . <i>d</i> [E <i>e</i>]
		or	EN <i>w</i> . <i>d</i> [E <i>e</i>]
		or	ES <i>w</i> . <i>d</i> [E <i>e</i>]
		or	G <i>w</i> . <i>d</i> [E <i>e</i>]
		or	L <i>w</i>

		or	A [<i>w</i>]
		or	D <i>w</i> . <i>d</i>
EXT		or	D <i>w</i> . <i>d</i> E <i>e</i>
EXT		or	R <i>w</i>
EXT		or	Q
R1006	<i>w</i>	is	<i>int_literal_constant</i>
R1007	<i>m</i>	is	<i>int_literal_constant</i>
R1008	<i>d</i>	is	<i>int_literal_constant</i>
R1009	<i>e</i>	is	<i>int_literal_constant</i>

Constraint: *e* must be positive.

Constraint: *w* must be zero or positive for the I, B, O, Z, and F edit descriptors. *w* must be positive for all other edit descriptors.

Constraint: *w*, *m*, *d*, and *e* must not have kind parameters specified for them.

R1010	<i>control_edit_desc</i>	is	<i>position_edit_desc</i>
		or	[<i>r</i>] /
		or	:
		or	<i>sign_edit_desc</i>
		or	<i>k</i> P
		or	<i>blank_interp_edit_desc</i>
R1011	<i>k</i>	is	<i>signed_int_literal_constant</i>

Constraint: *k* must not have a kind parameter specified for it.

R1012	<i>position_edit_desc</i>	is	T <i>n</i>
		or	TL <i>n</i>
		or	TR <i>n</i>
		or	<i>n</i> X

EXT		or	\
EXT		or	\$
R1013	<i>n</i>	is	<i>int_literal_constant</i>

Constraint: *n* must be positive.

Constraint: *n* must not have a kind parameter specified for it.

R1014	<i>sign_edit_desc</i>	is	S
		or	SP
		or	SS
R1015	<i>blank_interp_edit_desc</i>	is	BN
		or	BZ
R1016	<i>char_string_edit_desc</i>	is	<i>char_literal_constant</i>
EXT		or	<i>c</i> H <i>rep_char</i> [<i>rep_char</i>] . . .
EXT	<i>c</i>	is	<i>int_literal_constant</i>

Constraint: The *char_literal_constant* must not have a kind parameter specified for it.

1.2.11 Program Units

The following syntax rules are described in section 11, "Program units," of the Fortran 95 standard.

R1101	<i>main_program</i>	is	[<i>program_stmt</i>] [<i>specification_part</i>] [<i>execution_part</i>] [<i>internal_subprogram_part</i>] <i>end_program_stmt</i>
R1102	<i>program_stmt</i>	is	PROGRAM <i>program_name</i> [(<i>args</i>)]
EXT	<i>args</i>	is	Any character in the CF90 character

set. The CF90 compiler ignores any
args specified after *program_name*.

R1103	<i>end_program_stmt</i>	is	END [PROGRAM [<i>program_name</i>]]
-------	-------------------------	-----------	---

Constraint: In a *main_program*, the *execution_part* must not contain a RETURN statement or an ENTRY statement.

Constraint: The *program_name* may be included in the *end_program_stmt* only if the optional *program_stmt* is used and, if included, must be identical to the *program_name* specified in the *program_stmt*.

Constraint: An automatic object must not appear in the *specification_part* of a main program.

R1104	<i>module</i>	is	<i>module_stmt</i> [<i>specification_part</i>] [<i>module_subprogram_part</i>] <i>end_module_stmt</i>
-------	---------------	-----------	--

R1105	<i>module_stmt</i>	is	MODULE <i>module_name</i>
-------	--------------------	-----------	---------------------------

R1106	<i>end_module_stmt</i>	is	END [MODULE [<i>module_name</i>]]
-------	------------------------	-----------	---------------------------------------

Constraint: If the *module_name* is specified in the *end_module_stmt*, it must be identical to the *module_name* specified in the *module_stmt*.

Constraint: A module *specification_part* must not contain a *stmt_function_stmt*, an *entry_stmt*, or a *format_stmt*.

Constraint: An automatic object must not appear in the *specification_part* of a module.

Constraint: If an object of a type for which *component_initialization* is specified appears in the *specification_part* of a module and does not have the ALLOCATABLE or POINTER attribute, the object must have the SAVE attribute.

R1107	<i>use_stmt</i>	is	USE <i>module_name</i> [, <i>rename_list</i>]
		or	USE <i>module_name</i> , ONLY : [<i>only_list</i>]

R1108	<i>rename</i>	is	<i>local_name</i> => <i>use_name</i>
-------	---------------	-----------	--------------------------------------

R1109	<i>only</i>	is	<i>generic_spec</i>
		or	<i>only_use_name</i>
		or	<i>only_rename</i>
R1110	<i>only_use_name</i>	is	<i>use_name</i>
R1111	<i>only_rename</i>	is	<i>local_name</i> => <i>use_name</i>

Constraint: Each *generic_spec* must be a public entity in the module.

Constraint: Each *use_name* must be the name of a public entity in the module.

R1112	<i>block_data</i>	is	<i>block_data_stmt</i> [<i>specification_part</i>] <i>end_block_data_stmt</i>
R1113	<i>block_data_stmt</i>	is	BLOCK DATA [<i>block_data_name</i>]
R1114	<i>end_block_data_stmt</i>	is	END [BLOCK DATA [<i>block_data_name</i>]]

Constraint: The *block_data_name* may be included in the *end_block_data_stmt* only if it was provided in the *block_data_stmt* and, if included, must be identical to the *block_data_name* in the *block_data_stmt*.

Constraint: A *block_data_specification_part* may contain only USE statements, type declaration statements, IMPLICIT statements, PARAMETER statements, derived-type definitions, and the following specification statements: COMMON, DATA, DIMENSION, EQUIVALENCE, INTRINSIC, POINTER, SAVE, and TARGET.

Constraint: A type declaration statement in a *block_data_specification_part* must not contain ALLOCATABLE, EXTERNAL, INTENT, OPTIONAL, PRIVATE, or PUBLIC attribute specifiers.

1.2.12 Procedures

The following syntax rules are described in section 12, "Procedures," of the Fortran 95 standard.

R1201	<i>interface_block</i>	is	<i>interface_stmt</i> [<i>interface_specification</i>] . . . <i>end_interface_stmt</i>
R1202	<i>interface_specification</i>	is	<i>interface_body</i>
		or	<i>module_procedure_stmt</i>
R1203	<i>interface_stmt</i>	is	INTERFACE [<i>generic_spec</i>]
R1204	<i>end_interface_stmt</i>	is	END INTERFACE [<i>generic_spec</i>]
R1204	<i>interface_body</i>	is	<i>function_stmt</i> [<i>specification_part</i>] <i>end_function_stmt</i>
		or	<i>subroutine_stmt</i> [<i>specification_part</i>] <i>end_subroutine_stmt</i>
R1206	<i>module_procedure_stmt</i>	is	MODULE PROCEDURE <i>procedure_name_list</i>
R1207	<i>generic_spec</i>	is	<i>generic_name</i>
		or	OPERATOR (<i>defined_operator</i>)
		or	ASSIGNMENT (=)

Constraint: An interface body of a pure procedure must specify the intents of all dummy procedures except pointer, alternate return, and procedure arguments.

Constraint: An *interface_body* must not contain an *entry_stmt*, *data_stmt*, *format_stmt*, or *stmt_function_stmt*.

Constraint: The MODULE PROCEDURE statement is allowed only if the *interface_block* has a *generic_spec* and is in a scoping unit where each *procedure_name* is accessible as a module procedure.

Constraint: An *interface_block* in a subprogram must not contain an *interface_body* for a procedure defined by that subprogram.

Constraint: A *procedure_name* in a *module_procedure_stmt* must not be one that previously had been established to be associated with the *generic_spec* of the *interface_block* in which it appears, either by a previous appearance in an *interface_block* or by use or host association.

Constraint: The *generic_spec* can be included in the *end_interface_stmt* only if it was provided in the *interface_stmt*. If included, it must be identical to the *generic_spec* in the *interface_stmt*.

R1208	<i>external_stmt</i>	is	EXTERNAL [::] <i>external_name_list</i>
R1209	<i>intrinsic_stmt</i>	is	INTRINSIC [::] <i>intrinsic_procedure_name_list</i>

Constraint: Each *intrinsic_procedure_name* must be the name of an intrinsic procedure.

R1210	<i>function_reference</i>	is	<i>function_name</i> ([<i>actual_arg_spec_list</i>])
-------	---------------------------	-----------	--

Constraint: The *actual_arg_spec_list* for a function reference must not contain an *alt_return_spec*.

R1211	<i>call_stmt</i>	is	CALL <i>subroutine_name</i> [([<i>actual_arg_spec_list</i>])]
R1212	<i>actual_arg_spec</i>	is	[<i>keyword</i> =] <i>actual_arg</i>
R1213	<i>keyword</i>	is	<i>dummy_arg_name</i>
R1214	<i>actual_arg</i>	is	<i>expr</i>
		or	<i>variable</i>
		or	<i>procedure_name</i>
OBS		or	<i>alt_return_spec</i>
R1215	<i>alt_return_spec</i>	is	* <i>label</i>

Constraint: The *keyword* = must not appear if the interface of the procedure is implicit in the scoping unit.

Constraint: The *keyword* = may be omitted from an *actual_arg_spec* only if the *keyword* = has been omitted from each preceding *actual_arg_spec* in the argument list.

Constraint: Each *keyword* must be the name of a dummy argument in the explicit interface of the procedure.

Constraint: A *procedure_name actual_arg* must not be the name of an internal procedure or of a statement function and must not be the generic name of a procedure.

Constraint: The *label* used in the *alt_return_spec* must be the statement label of a branch target statement that appears in the same scoping unit as the *call_stmt*.

Constraint: A nonintrinsic elemental procedure must not be used as an actual argument.

Constraint: In a reference to a pure procedure, a *procedure_name actual_arg* must be the name of a pure procedure.

R1216	<i>function_subprogram</i>	is	<i>function_stmt</i> [<i>specification_part</i>] [<i>execution_part</i>] [<i>internal_subprogram_part</i>] <i>end_function_stmt</i>
R1217	<i>function_stmt</i>	is	[<i>prefix</i>] FUNCTION <i>function_name</i> ([<i>dummy_arg_name_list</i>]) [RESULT (<i>result_name</i>)]

Constraint: If RESULT is specified, the *function_name* must not appear in any specification statement in the scoping unit of the function subprogram.

Constraint: A *prefix* must contain at most one of each *prefix_spec*.

Constraint: If ELEMENTAL is present, RECURSIVE must not be present.

R1218	<i>prefix</i>	is	<i>prefix_spec</i> [<i>prefix_spec</i>] ...
R1219	<i>prefix_spec</i>	is	<i>type_spec</i> or RECURSIVE or PURE or ELEMENTAL
R1220	<i>end_function_stmt</i>	is	END [FUNCTION [<i>function_name</i>]]

Constraint: If `RESULT` is specified, `result_name` must not be the same as `function_name`.

Constraint: `FUNCTION` must be present on the `end_function_stmt` of an internal or module function.

Constraint: An internal function subprogram must not contain an `ENTRY` statement.

Constraint: An internal function subprogram must not contain an `internal_subprogram_part`.

Constraint: If a `function_name` is present on the `end_function_stmt`, it must be identical to the `function_name` specified in the `function_stmt`.

R1221	<code>subroutine_subprogram</code>	is	<code>subroutine_stmt</code> [<code>specification_part</code>] [<code>execution_part</code>] [<code>internal_subprogram_part</code>] <code>end_subroutine_stmt</code>
R1222	<code>subroutine_stmt</code>	is	[<code>prefix</code>] <code>SUBROUTINE</code> <code>subroutine_name</code> [([<code>dummy_arg_list</code>])]
R1223	<code>dummy_arg</code>	is	<code>dummy_arg_name</code>
		or	*
R1224	<code>end_subroutine_stmt</code>	is	<code>END</code> [<code>SUBROUTINE</code> [<code>subroutine_name</code>]]

Constraint: The prefix of a `subroutine_stmt` must not contain a `type_spec`.

Constraint: `SUBROUTINE` must be present on the `end_subroutine_stmt` of an internal or module subroutine.

Constraint: An internal subroutine must not contain an `ENTRY` statement.

Constraint: An internal subroutine must not contain an `internal_subprogram_part`.

Constraint: If a `subroutine_name` is present on the `end_subroutine_stmt`, it must be identical to the `subroutine_name` specified in the `subroutine_stmt`.

R1225	<code>entry_stmt</code>	is	<code>ENTRY</code> <code>entry_name</code> [([<code>dummy_arg_list</code>]) [<code>RESULT</code> (<code>result_name</code>)]]
-------	-------------------------	-----------	--

Constraint: If *RESULT* is specified, the *entry_name* must not appear in any specification statement in the scoping unit of the function program.

Constraint: An *entry_stmt* may appear only in an *external_subprogram* or *module_subprogram*. An *entry_stmt* must not appear within an *executable_construct*.

Constraint: *RESULT* may be present only if the *entry_stmt* is contained in a function subprogram.

Constraint: Within the subprogram containing the *entry_stmt*, the *entry_name* must not appear as a dummy argument in the *FUNCTION* or *SUBROUTINE* statement or in another *ENTRY* statement and it must not appear in an *EXTERNAL* or *INTRINSIC* statement.

Constraint: A *dummy_arg* can be an alternate return indicator only if the *ENTRY* statement is in a subroutine subprogram.

Constraint: If *RESULT* is specified, *result_name* must not be the same as *entry_name*.

R1226	<i>return_stmt</i>	is	RETURN
OBS		or	RETURN [<i>scalar_int_expr</i>]

Constraint: The *return_stmt* must be in the scoping unit of a function or subroutine subprogram.

Obsolescent Constraint: The *scalar_int_expr* is allowed only in the scoping unit of a subroutine subprogram.

R1227	<i>contains_stmt</i>	is	CONTAINS
R1228	<i>stmt_function_stmt</i>	is	<i>function_name</i> ([<i>dummy_arg_name_list</i>]) = <i>scalar_expr</i>

Obsolescent Constraint: The *primaries* of the *scalar_expr* must 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 function dummy procedure, the reference must not require an explicit interface; the function must not require an explicit interface unless it is an intrinsic; the function must not be a transformational intrinsic; and the result

must be scalar. If an argument to a function or a function dummy procedure is array valued, it must be an array name. If a reference to a statement function appears in *scalar_expr*, its definition must have been provided earlier in the scoping unit and must not be the name of the statement function being defined.

Obsolescent Constraint: Named constants in *scalar_expr* must have been declared earlier in the scoping unit or made accessible by `USE` or host association. If array elements appear in *scalar_expr*, the parent array must not have been declared as an array earlier in the scoping unit or made accessible by `USE` or host association.

Obsolescent Constraint: 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 must confirm this implied type and the values of any implied type parameters.

Constraint: The *function_name* and each *dummy_arg_name* must be specified, explicitly or implicitly, to be scalar data objects.

Constraint: A given *dummy_arg_name* may appear only once in any *dummy_arg_name_list*.

Constraint: Each scalar variable reference in *scalar_expr* may be either a reference to a dummy argument of the statement function or a reference to a variable local to the same scoping unit as the statement function statement.

Constraint: The *specification_part* of a pure function subprogram must specify that all dummy arguments have `INTENT(IN)` except procedure arguments and arguments with the `POINTER` attribute.

Constraint: The *specification_part* of a pure subroutine subprogram must specify the intents of all dummy arguments except procedure arguments, alternate return indicators, and arguments with the `POINTER` attribute. Note that alternate return indicators are obsolete.

Constraint: A local variable declared in the *specification_part* or *internal_subprogram_part* of a pure subprogram must not have the `SAVE` attribute.

Constraint: The *specification_part* of a pure subprogram must specify that all dummy arguments that are procedure arguments are pure.

Constraint: If a procedure that is neither an intrinsic procedure nor a statement function is used in a context that requires it to be pure, its interface must be explicit in the scope of that use. The interface must specify that the procedure is pure. Note that statement functions are obsolete.

Constraint: All internal subprograms in a pure subprogram must be pure.

Constraint: In a pure subprogram, any variable that is in common or is accessed by host or USE association, is a dummy argument to a pure function, is a dummy argument with `INTENT(IN)` to a pure subroutine, or an object that is storage associated with any such variable, must not be used in the following contexts:

- As the *variable* of an *assignment_stmt*
- As a DO variable or implied DO variable
- As an *input_item* in a *read_stmt* from an internal file
- As an *internal_file_unit* in a *write_stmt*
- As an `IOSTAT=` specifier in an input or output statement with an internal file
- As the *pointer_object* of a *pointer_assignment_stmt*
- As the *target* of a *pointer_assignment_stmt*
- As the *expr* of an *assignment_stmt* in which the *variable* is of a derived type if the derived type has a pointer component at any level of component selection
- As an *allocate_object* or *stat_variable* in an *allocate_stmt* or *deallocate_stmt*, or as a *pointer_object* in a *nullify_stmt*
- As an actual argument associated with a dummy argument with `INTENT(OUT)` or `INTENT(INOUT)` or with the `POINTER` attribute

Constraint: Any procedure referenced in a pure subprogram, including one referenced via a defined operation or assignment, must be pure.

Constraint: A pure subprogram must not contain a *print_stmt*, *open_stmt*, *close_stmt*, *backspace_stmt*, *endfile_stmt*, *rewind_stmt*, or *inquire_stmt*.

Constraint: A pure subprogram must not contain a *read_stmt* or *write_stmt* with an *io_unit* that is an *external_file_unit* or an asterisk (*).

Constraint: A pure subprogram must not contain a *stop_stmt*.

1.2.13 Intrinsic Procedures

There are no syntax rules described in section 13, "Intrinsic procedures," of the Fortran 95 standard.

1.2.14 Scope, Association, and Definition

There are no syntax rules described in section 14, "Scope, association, and definition," of the Fortran 95 standard.

Decremental Features [2]

This chapter describes Fortran features that have been deleted or declared obsolescent in the current (Fortran 95) standard. The CF90 and MIPSpro 7 Fortran 90 compilers continue to support these features as extensions to the standard. The CF90 and MIPSpro 7 Fortran 90 compilers can generate messages when the deleted or obsolescent features are detected; to enable these messages, specify one of the following on the `f90(1)` command line: `-ansi` and `-fullwarn` (on IRIX systems) or `-e n` (on UNICOS or UNICOS/mk systems). For more information on these options, see the `f90(1)` man page.

2.1 Deleted Features

This section describes features that the Fortran 95 standard declares to be deleted from the Fortran language. These features had been included in previous revisions of the Fortran standard. The CF90 and MIPSpro 7 Fortran 90 compilers include these features as extensions. The deleted features are as follows:

- Real and double-precision `DO` variables. The preferred alternative is integer.
- The ability to branch to an `END IF` statement from outside its block. The preferred alternative is to branch to the statement following the `END IF`.
- The `PAUSE` statement. For more information on this, see Section 2.1.1, page 69.
- The `ASSIGN` statement, assigned `GO TO` statements, and assigned format specifiers. For more information on this, see Section 2.1.2, page 70.
- The `H` edit descriptor. For more information on this, see Section 2.1.3, page 72.

2.1.1 PAUSE Statement

Execution of a `PAUSE` statement requires operator or system-specific intervention to resume execution. In most cases, the same functionality can be achieved as effectively and in a more portable way with the use of an appropriate `READ` statement that awaits some input data.

The execution of the `PAUSE` statement suspends the execution of a program. This is now redundant, because a `WRITE` statement can be used to send a

message to any device, and a READ statement can be used to wait for and receive a message from the same device.

The PAUSE statement is defined as follows:

EXT	<i>pause_stmt</i>	is	PAUSE [<i>stop_code</i>]
-----	-------------------	----	----------------------------

The character constant or list of digits identifying the PAUSE statement is called the *stop_code* because it follows the same rules as those for the STOP statement's stop code. The stop code is accessible following program suspension. The CF90 and MIPSpro 7 Fortran 90 compilers send the *stop_code* to the standard error file (`stderr`). The following are examples of PAUSE statements:

```
PAUSE
PAUSE 'Wait #823'
PAUSE 100
```

2.1.2 ASSIGN, Assigned GO TO Statements, and Assigned Format Specifiers

The ASSIGN statement assigns a statement label to an integer variable. During program execution, the variable can be assigned labels of branch target statements, providing a dynamic branching capability in a program. The unsatisfactory property of these statements is that the integer variable name can be used to hold both a label and an ordinary integer value, leading to errors that can be hard to discover and programs that can be difficult to read.

A frequent use of the ASSIGN statement and assigned GO TO statement is to simulate internal procedures, using the ASSIGN statement to record the return point after a reusable block of code has completed. The internal procedure mechanism of Fortran now provides this capability.

A second use of the ASSIGN statement is to simulate dynamic format specifications by assigning labels corresponding to different format statements to an integer variable and using this variable in I/O statements as a format specifier. This use can be accomplished in a clearer way by using character strings as format specifications. Thus, it is no longer necessary to use either the ASSIGN statement or the assigned GO TO statement.

Execution of an ASSIGN statement causes the variable in the statement to become defined with a statement label value.

When a numeric storage unit becomes defined, all associated numeric storage units of the same type become defined. Variables associated with the variable in an ASSIGN statement, however, become undefined as integers when the ASSIGN statement is executed. When an entity of double-precision real type becomes defined, all totally associated entities of double-precision real type become defined.

Execution of an ASSIGN statement causes the variable in the statement to become undefined as an integer. Variables that are associated with the variable also become undefined.

2.1.2.1 Form of the ASSIGN and Assigned GO TO Statements

Execution of an ASSIGN statement assigns a label to an integer variable. Subsequently, this value can be used by an assigned GO TO statement or by an I/O statement to reference a FORMAT statement. The ASSIGN statement is defined as follows:

EXT	<i>assign_stmt</i>	is	ASSIGN <i>label</i> TO <i>scalar_int_variable</i>
-----	--------------------	----	---

The term *default integer type* in this section means that the integer variable must occupy a full word in order to be able to hold the address of the statement label. On CRAY T3E systems, programs that contain an ASSIGN statement and are compiled with `-i 32` or `-s default32` must ensure that the *scalar_int_variable* is declared as `INTEGER(KIND=8)`. This ensures that it occupies a full word.

The variable must be a named variable of default integer type. It must not be an array element, an integer component of a structure, or an object of nondefault integer type.

The label must be the label of a branch target statement or the label of a FORMAT statement in the same scoping unit as the ASSIGN statement.

When defined with an integer value, the integer variable cannot be used as a label.

When assigned a label, the integer variable cannot be used as anything other than a label.

When the integer variable is used in an assigned GO TO statement, it must be assigned a label.

As the following example shows, the variable can be redefined during program execution with either another label or an integer value:

```
ASSIGN 100 TO K
```

Execution of the assigned GO TO statement causes a transfer of control to the branch target statement with the label that had previously been assigned to the integer variable.

The assigned GO TO statement is defined as follows:

EXT *assigned_goto_stmt* **is** GO TO *scalar_int_variable* [[,] (*label_list*)]

The variable must be a named variable of default integer type. That is, it must not be an array element, a component of a structure, or an object of nondefault integer type.

The variable must be assigned the label of a branch target statement in the same scoping unit as the assigned GO TO statement.

If a label list appears, such as in the following examples, the variable must have been assigned a label value that is in the list:

```
GO TO K
GO TO K (10, 20, 100)
```

The ASSIGN statement also allows the label of a FORMAT statement to be dynamically assigned to an integer variable, which can later be used as a format specifier in READ, WRITE, or PRINT statements. This hinders readability, permits inconsistent usage of the integer variable, and can be an obscure source of error.

This functionality is available through character variables, arrays, and constants.

2.1.2.2 Assigned Format Specifiers

When an I/O statement containing the integer variable as a format specifier is executed, the integer variable can be defined with the label of a FORMAT specifier.

2.1.3 H Edit Descriptor

This edit descriptor can be a source of error because the number of characters following the descriptor can be miscounted easily. The same functionality is

available using the character constant edit descriptor, for which no count is required.

The following information pertains to the H edit descriptor:

Table 2. Summary of string edit descriptors

Descriptor	Description
H	Transfer of text character to output record
'text'	Transfer of a character literal constant to output record
"text"	Transfer of a character literal constant to output record

For more information on edit descriptors, see the *Fortran Language Reference Manual, Volume 2*.

2.2 Obsolescent Features

The obsolescent features are those features of previous Fortran standards that are considered by the Fortran 95 standard to be redundant. The Fortran 95 standard states that these features are obsolescent and provides preferred alternatives.

The obsolescent features and their preferred alternatives are as follows:

- Arithmetic IF. The preferred alternative is the IF statement or IF construct. For more information on the arithmetic IF, see the *Fortran Language Reference Manual, Volume 1*.
- Shared DO termination and termination on a statement other than END DO or CONTINUE statements. The preferred alternative is an END DO or a CONTINUE statement for each DO statement. For more information on this DO termination, see the *Fortran Language Reference Manual, Volume 1*.
- Alternate return. An alternate return introduces labels into an argument list to allow the called procedure to direct the execution of the caller upon return. The preferred alternative is to use a return code that is used in a CASE construct on return. This avoids an irregularity in the syntax and semantics of argument association. Consider the following statement:

```
CALL SUBR_NAME (X, Y, Z, *100, *200, *300)
```

The preceding statement can be replaced by the following code:

```
CALL SUBR_NAME (X, Y, Z, RETURN_CODE)
SELECT CASE (RETURN_CODE)
  CASE (1)
    ...
  CASE (2)
    ...
  CASE (3)
    ...
  CASE DEFAULT
    ...
END SELECT
```

For more information on alternate returns, see *Fortran Language Reference Manual, Volume 2*.

- Computed GO TO statement. The preferred alternative is the CASE construct, which is a generalized, easier to use, and more efficient means of expressing the same computation. For more information on the computed GO TO statement, see the *Fortran Language Reference Manual, Volume 1*.
- Statement functions. These are subject to a number of nonintuitive restrictions and are a potential source of error because their syntax is easily confused with that of an assignment statement. The preferred alternative is the internal function, which is a more generalized form of the statement function that completely superseded the statement function construct. For more information on statement functions, see the *Fortran Language Reference Manual, Volume 1*.
- DATA statements among executables. The statement ordering rules of previous Fortran standards allowed DATA statements to appear anywhere in a program unit after the specification statements. The ability to position DATA statements amongst executable statements is rarely used and is a potential source of error. For more information on the DATA statement, see the *Fortran Language Reference Manual, Volume 1*.
- Assumed character length functions. Assumed character length for functions is an irregularity in the language because the typical Fortran philosophy is that the attributes of a function result depend only on the actual arguments of the invocation and on any data accessible by the function through host or USE association. Some uses of this facility can be replaced with an automatic character length function in which the length of the function result is declared in a specification expression. Other uses can be replaced by the use

of a subroutine with arguments that correspond to the function result and the function arguments.

Note that dummy arguments of a function can be assumed character length. For more information on assumed character length functions, see the *Fortran Language Reference Manual, Volume 2*.

- Fixed source form. Fixed source form was designed when the principal machine-readable input medium for new programs was punched cards. Now that new and amended programs are typically entered at a keyboard with a screen display terminal, it is unnecessary overhead and is potentially error-prone to have to locate positions 6, 7, or 72 on a line. Free form source was designed expressly for this more modern technology. It is a simple matter to convert from fixed form to free form. For more information on fixed source form, see the *Fortran Language Reference Manual, Volume 1*.
- CHARACTER* form of CHARACTER declaration. The CHARACTER*length form of specifying character declarations is redundant. For more information on the CHARACTER* form, see the *Fortran Language Reference Manual, Volume 1*.

Character Set [3]

The ASCII character set contains the control and graphic characters shown in the following table. Numbers, letters, and special characters in the character set are identified by the letter "C" in the *Notes* column. All other characters are members of the auxiliary character set. The letter "A" identifies the characters that belong to the Fortran character set as defined by the standard. Letters in parentheses following the descriptions in the *Description* column indicate the following control character usage:

- "CC" stands for Communication Control
- "FE" stands for Format Effector
- "IS" stands for Information Separator

Table 3. Character set

Character	Octal	Decimal	Hex	Notes	Description
NUL	000	000	00		Null
SOH	001	001	01		Start of heading (CC)
STX	002	002	02		Start of text (CC)
ETX	003	003	03		End of text (CC)
EOT	004	004	04		End of transmission (CC)
ENQ	005	005	05		Enquiry (CC)
ACK	006	006	06		Acknowledge (CC)
BEL	007	007	07		Bell (audible signal)
BS	010	008	08		Backspace (FE)
HT	011	009	09	C	Horizontal tabulation (FE)
LF	012	010	0A		Line feed (FE)
VT	013	011	0B		Vertical tabulation (FE)
FF	014	012	0C		Form feed (FE)
CR	015	013	0D		Carriage return (FE)

Character	Octal	Decimal	Hex	Notes	Description
SO	016	014	0E		Shift out
SI	017	015	0F		Shift in
DLE	020	016	10		Data link escape (CC)
DC1	021	017	11		Device control 1
DC2	022	018	12		Device control 2
DC3	023	019	13		Device control 3
DC4	024	020	14		Device control 4 (stop)
NAK	025	021	15		Negative acknowledge (CC)
SYN	026	022	16		Synchronous idle (CC)
ETB	027	023	17		End of transmission block (CC)
CAN	030	024	18		Cancel
EM	031	025	19		End of medium
SUB	032	026	1A		Substitute
ESC	033	027	1B		Escape
FS	034	028	1C		File separator (IS)
GS	035	029	1D		Group separator (IS)
RS	036	030	1E		Record separator (IS)
US	037	031	1F		Unit separator (IS)
space	040	032	20	A, C	(blank)
!	041	033	21	A, C	Exclamation point
"	042	034	22	A, C	Quotation mark
#	043	035	23		Number sign
\$	044	036	24	A, C	Dollar sign (currency symbol)
%	045	037	25	A, C	Percent
&	046	038	26	A, C	Ampersand
'	047	039	27	A, C	Apostrophe (single quote)
(050	040	28	A, C	Opening (left) parenthesis
)	051	041	29	A, C	Closing (right) parenthesis

Character	Octal	Decimal	Hex	Notes	Description
*	052	042	2A	A, C	Asterisk
+	053	043	2B	A, C	Plus
,	054	044	2C	A, C	Comma (cedilla)
-	055	045	2D	A, C	Minus (hyphen)
.	056	046	2E	A, C	Period (decimal point)
/	057	047	2F	A, C	Slant (slash, virgule)
0	060	048	30	A, C	Zero
1	061	049	31	A, C	One
2	062	050	32	A, C	Two
3	063	051	33	A, C	Three
4	064	052	34	A, C	Four
5	065	053	35	A, C	Five
6	066	054	36	A, C	Six
7	067	055	37	A, C	Seven
8	070	056	38	A, C	Eight
9	071	057	39	A, C	Nine
:	072	058	3A	A, C	Colon
;	073	059	3B	A, C	Semicolon
<	074	060	3C	A, C	Less than
=	075	061	3D	A, C	Equal
>	076	062	3E	A, C	Greater than
?	077	063	3F	A, C	Question mark
@	100	064	40	C	"At" sign. Reserved for internal use. Not a valid character on IRIX systems.
A	101	065	41	A, C	Uppercase letter
B	102	066	42	A, C	Uppercase letter
C	103	067	43	A, C	Uppercase letter
D	104	068	44	A, C	Uppercase letter

Character	Octal	Decimal	Hex	Notes	Description
E	105	069	45	A, C	Uppercase letter
F	106	070	46	A, C	Uppercase letter
G	107	071	47	A, C	Uppercase letter
H	110	072	48	A, C	Uppercase letter
I	111	073	49	A, C	Uppercase letter
J	112	074	4A	A, C	Uppercase letter
K	113	075	4B	A, C	Uppercase letter
L	114	076	4C	A, C	Uppercase letter
M	115	077	4D	A, C	Uppercase letter
N	116	078	4E	A, C	Uppercase letter
O	117	079	4F	A, C	Uppercase letter
P	120	080	50	A, C	Uppercase letter
Q	121	081	51	A, C	Uppercase letter
R	122	082	52	A, C	Uppercase letter
S	123	083	53	A, C	Uppercase letter
T	124	084	54	A, C	Uppercase letter
U	125	085	55	A, C	Uppercase letter
V	126	086	56	A, C	Uppercase letter
W	127	087	57	A, C	Uppercase letter
X	130	088	58	A, C	Uppercase letter
Y	131	089	59	A, C	Uppercase letter
Z	132	090	5A	A, C	Uppercase letter
{	133	091	5B		Opening (left) brace
\	134	092	5C		Reverse slant (backslash)
}	135	093	5D		Closing (right) brace
^	136	094	5E		Caret (circumflex)
_	137	095	5F	A, C	Underline
'	140	096	60		Grave accent

Character	Octal	Decimal	Hex	Notes	Description
a	141	097	61	A, C	Lowercase letter
b	142	098	62	A, C	Lowercase letter
c	143	099	63	A, C	Lowercase letter
d	144	100	64	A, C	Lowercase letter
e	145	101	65	A, C	Lowercase letter
f	146	102	66	A, C	Lowercase letter
g	147	103	67	A, C	Lowercase letter
h	150	104	68	A, C	Lowercase letter
i	151	105	69	A, C	Lowercase letter
j	152	106	6A	A, C	Lowercase letter
k	153	107	6B	A, C	Lowercase letter
l	154	108	6C	A, C	Lowercase letter
m	155	109	6D	A, C	Lowercase letter
n	156	110	6E	A, C	Lowercase letter
o	157	111	6F	A, C	Lowercase letter
p	160	112	70	A, C	Lowercase letter
q	161	113	71	A, C	Lowercase letter
r	162	114	72	A, C	Lowercase letter
s	163	115	73	A, C	Lowercase letter
t	164	116	74	A, C	Lowercase letter
u	165	117	75	A, C	Lowercase letter
v	166	118	76	A, C	Lowercase letter
w	167	119	77	A, C	Lowercase letter
x	170	120	78	A, C	Lowercase letter
y	171	121	79	A, C	Lowercase letter
z	172	122	7A	A, C	Lowercase letter
[173	123	7B		Opening (left) bracket
	174	124	7C		Vertical line

Character	Octal	Decimal	Hex	Notes	Description
]	175	125	7D		Closing (right) bracket
~	176	126	7E		Overline (tilde, general accent)
DEL	177	127	7F		Delete

Extensions and Differences [4]

This chapter describes the differences various Fortran implementations. This chapter is divided into the following sections:

- Fortran 95 standard differences and incompatibilities with FORTRAN 77 implementations
- CF90 extensions to Fortran 95
- MIPSpro 7 Fortran 90 extensions to Fortran 95
- CF90, MIPSpro Fortran 77 and MIPSpro 7 Fortran 90 differences

4.1 Fortran 95 Standard Differences and Incompatibilities With FORTRAN 77 Implementations

This section discusses the following topics:

- Fortran 95 and the `G` edit descriptor output differences
- Fortran 95 and list-directed output differences
- Delimited and undelimited character strings in list-directed I/O
- List-directed I/O and floating-point zero

4.1.1 Fortran 95 and the `G` Edit Descriptor Output Differences

The format of a floating-point zero written with a `G` edit descriptor is different in Fortran 95.

The floating-point zero was written with an `Ew.d` edit descriptor in FORTRAN 77, but it is written with an `Fw.d` edit descriptor in the CF90 and MIPSpro 7 Fortran 90 compilers.

The `G` edit descriptor has been expanded to be a general edit descriptor that can read or write any data type including character, integer, and logical data. FORTRAN 77 allows only floating-point data types.

Fortran 95 is specific about the rounding of floating-point values with the `G` format. The change in rules may cause asterisks in the output field for some

floating-point values. Other values will be written as an *Ew.d*-formatted value by the CF90 and MIPSpro 7 Fortran 90 compilers.

Consider the following code fragment:

```
DOUBLE PRECISION AVD, BVD, CVD
AVD = 0.0D0
WRITE(6, 1) AVD
1  FORMAT(G28.2)
END
```

It generates the following output when compiled by CF90:

```
0.0
```

4.1.2 Fortran 95 and List-directed Output Differences

Fortran 95 requires a separator between noncharacter data and character data in list-directed output. FORTRAN 77 disallows a separator in this instance.

Consider the following example output list:

```
'This is a one(',1,')
```

This output list generates different output under the two standards:

- Fortran 95 output:

```
>This is a one( 1 )
```

- FORTRAN 77 output:

```
>This is a one(1)
```

4.1.3 Delimited and Undelimited Character Strings in List-directed I/O

FORTAN 77 compilers support only delimited character string input to a list-directed item that will be stored to a list item of type character.

4.1.4 List-directed I/O and Floating-point Zero

Fortran 95 specifies a different form of output constant for a floating-point zero in list-directed output records. Consider the following program:

```
      PRINT *,0.0
      PRINT 1, 0.0
1     FORMAT(1X,G12.2)
      END
```

The preceding code generates the following output for the CF90 compiler:

```
% f90 tt.f
% a.out
 0.E+0
   0.0
```

4.2 CF90 extensions to Fortran 95

This section discusses the following topics:

- List-directed I/O and Hollerith constants
- Differences in the B, O, and Z edit descriptors

4.2.1 List-directed I/O and Hollerith Constants

Fortran 95 does not support Hollerith constants in list-directed input files, but the CF90 and MIPSpro 7 Fortran 90 compilers provides this as an extension. Note that Hollerith data is an deleted feature. See Chapter 6, page 137, for information on deprecated features and preferred alternatives.

4.2.2 Differences in the B, O, and z Edit Descriptors

The B, O, and Z edit descriptors are available in Fortran 95. They are limited to integer I/O list items.

Signed octal and hexadecimal values are not allowed in Fortran 95. The CF90 and MIPSpro 7 Fortran 90 compilers allow signed input values, but they write only unsigned values.

If the size of the value is less than w in ow or zw on output, Fortran 95 requires blank padding on the left. If the edit descriptor $ow.m$ or $zw.m$ is used, the field must contain at least m digits. The $.m$ form is one way to get leading zeros with Fortran 95. If the size of the value is greater than w in ow or zw on output, the CF90 and MIPSpro 7 Fortran 90 compilers fill the field with asterisks. The CF90 and MIPSpro 7 Fortran 90 compilers provide the Fortran 95 form of ow and zw output.

4.3 MIPSpro 7 Fortran 90 extensions to Fortran 95

The MIPSpro 7 Fortran 90 compiler supports the following extensions to Fortran 95:

- Formatted I/O:
 - Q edit descriptor returns the number of characters remaining in the record.
 - The dollar sign (\$) and backslash (\) edit descriptors suppress the newline character at the end of a record.
 - The field width of a data edit descriptor other than the A format cannot be present. The data edit descriptors are B, D, E, F, G, I, L, O, and Z. The field width defaults to a size chosen by the I/O library. The Fortran 95 standard allows a zero for the field width for the I, B, O, and Z data edit descriptors.
- Internal I/O:
 - ENCODE statement
 - DECODE statement
- NAMELIST I/O:
 - Skip unmatched namelist groups on input without error.
 - Accept either an ampersand (&) or dollar sign (\$) as a prefix to the namelist group name on input. If the dollar sign is used, the slash (/) or \$END can be used to terminate the namelist input group. If the ampersand is used, the slash or \$END can be used to terminate the namelist input group.
- Unformatted I/O:
 - Specify FORM= 'SYSTEM' to ensure that no record headers exist in the file. This is available on all systems but is only necessary where the f77 layer is the default for an unformatted file.
 - Specify FORM= 'BINARY' to ensure that no record headers exist in the file.

4.4 CF90, MIPSpro 7 Fortran 90, and MIPSpro Fortran 77 Differences

This section describes differences between the various Fortran compilers supported on IRIX, UNICOS, and UNICOS/mk systems.

4.4.1 MIPSpro 7 Fortran 90 and CF90 Compiler Differences

The following sections describe various differences found when compiling Fortran programs with the MIPSpro 7 Fortran 90 compiler and the CF90 compiler.

4.4.1.1 Numerical Model Differences

The model differences are as follows:

- The model for the CF90 `REAL(KIND=16)` data type on CRAY T90 systems that support IEEE floating-point arithmetic is different from the model for the MIPSpro 7 Fortran 90 compiler. This means that the results of math functions, arithmetic calculations, I/O, and other library routines are different for this particular data type.
- The internal size of `INTEGER(KIND=1)`, `INTEGER(KIND=2)`, `LOGICAL(KIND=1)`, and `LOGICAL(KIND=2)` on the MIPSpro 7 Fortran 90 compiler is actually 1 and 2 bytes, respectively. The CF90 compiler treats these kind type parameters as `INTEGER(KIND=4)` and `LOGICAL(KIND=4)`.
- The default sizes of the MIPSpro 7 Fortran 90 integer, real, and logical data types are 32 bits. This differs from the CF90 default of 64 bits. The default data type sizes for the MIPSpro 7 Fortran 90 compiler may be incorrect for routines such as `IRTC(3I)` and `SHMEM`.
- The MIPSpro 7 Fortran 90 compiler does not support Cray character pointers.
- Pointer arithmetic is in default numeric storage units when using the CF90 compiler. Pointer arithmetic is in bytes when using the MIPSpro 7 Fortran 90 compiler.

For more information on the model, see the `models(3I)` man page.

4.4.1.2 Fortran Statement Differences

The Fortran 95 statement differences are as follows:

- When using the MIPSpro 7 Fortran 90 compiler, the execution of the `STOP` statement does not cause the word `STOP` to be written to `stdout` unless there is an argument to the `STOP` statement. The CF90 compiler always writes `STOP` to `stdout`.
- When using the MIPSpro 7 Fortran 90 compiler, the initialization of entities in a common block in a `DATA` statement can only be done in one program unit. That is, if a common block contains two variables initialized in a `DATA`

statement, those DATA statements must be in one program unit. The load indicates the presence of multiple initializations, and only one initialization is done.

With the CF90 compiler, different variables can be initialized in DATA statements in separate program units.

4.4.1.3 Function and Procedure Differences

The CF90 typeless functions (such as MASK(3I), SHIFTL(3I), SHIFTR(3I), SHIFT(3I), CVM(3I), and so on) are typed as integer functions by the MIPSpro 7 Fortran 90 compiler. Conversion occur in expressions involving a mixture of floating point and integer functions. When called by the CF90 compiler, these functions are typeless and no conversion occurs when there is a mixture of floating point and these typeless functions.

4.4.1.4 Modules Differences

When using the MIPSpro 7 Fortran 90 compiler, the compilation of Fortran 95 modules creates a *file.mod* for each module in the source file and creates a *file.o* for any module procedures.

To load compiled module procedures, specify *module.o* on the command line.

When using the CF90 compiler, compiling modules creates one *file.o* that contains all the Fortran 95 modules in the source file.

4.4.1.5 I/O Library Differences

The I/O library differences are as follows:

- Direct access formatted output files cannot be read as sequential formatted files by MIPSpro 7 Fortran 90 programs unless an `assign(1)` command with `-s unblocked`, `-F cachea`, or `-F cache` is supplied for the particular file.
- The set of I/O library errors begins at 4000 for MIPSpro 7 Fortran 90 programs. The error numbers begin at 1000 for CF90 programs.
- The `FILENV` environment variable must be set for MIPSpro 7 Fortran 90 programs when using the `assign(1)` command. For CF90 users, this environment variable need not be set.

4.4.1.6 Library Function and Procedure Differences

The library function and intrinsic procedure differences are as follows:

- The `CRI_IEEE_DEFINITIONS` module is available for the MIPSpro 7 Fortran 90 compiler, but the preferred name is `FTN_IEEE_DEFINITIONS` for the IEEE module and the interface to the IEEE procedures.
- The `MAXVAL(3I)` intrinsic procedure returns negative infinity for a zero-sized input array when called from a MIPSpro 7 Fortran 90 program and returns `-HUGE(3I)` when called from a CF90 program. A request for interpretation has been submitted to the Fortran standards committee.
- The `MINVAL(3I)` intrinsic procedure returns positive infinity for a zero-sized input array when called from a MIPSpro 7 Fortran 90 program and returns `+HUGE(3I)` when called from a CF90 program. A request for interpretation has been submitted to the Fortran standards committee.

4.4.1.7 Math Library Differences

The math library differences are as follows:

- The math routines from the MIPSpro 7 Fortran 90 compiler are referenced from the compiler. The results of the math routines from the MIPSpro 7 Fortran 90 compiler may differ from the results returned by the math routines for the CF90 compiler.
- Signaling of errors during references to the MIPSpro 7 Fortran 90 compiler math routines is not turned off. For the CF90 compiler, the math routines turn off signaling of errors and detect input data errors through source code checks.

4.4.2 MIPSpro FORTRAN 77 and MIPSpro 7 Fortran 90 Compiler Differences

The following sections describe various differences found when compiling Fortran programs with the MIPSpro FORTRAN 77 compiler and the MIPSpro 7 Fortran 90 compiler.

4.4.2.1 Intrinsic Function and Subroutine Differences

The MIPSpro FORTRAN 77 compiler supports the `TIME` intrinsic function and the MIPSpro 7 Fortran 90 compiler does not. The Fortran 95 standard defines the `DATE_AND_TIME(3I)` function, and its use is recommended when using the MIPSpro 7 Fortran 90 compiler.

4.4.2.2 DATA Statement Initialization Differences

The Fortran 95 standard explicitly disallows multiple explicit initializations of the same variable or part of a variable. Doing so results in undefined behavior.

Some codes initialize the same local variable or part of a variable in a DATA statement. Some codes initialize data in COMMON blocks more than once, either in the same or in different program units.

The MIPSpro 7 Fortran 90 compiler, like many other implementations, allows COMMON blocks to be initialized in program units other than BLOCKDATA subprograms. Multiple initializations are not detected by the system. As a result, different processors may exhibit different behavior in cases of multiple initializations. For example, one processor may use the last value seen as the value of the initialized variable, while another may use the first value seen. Porting a code from one of these processors to another may result in differing results due to this difference.

Permitting multiple initializations of the same or part of a variable is not an extension. It is a user error that cannot be detected, in all cases, by the compiler. Behavior of multiple initializations is different across the IRIX, UNICOS, and UNICOS/mk platforms. For the program to be a standard conforming program with predictable results, you must remove multiple initializations.

4.4.2.3 I/O Record Length Differences

Fortran 95 standard I/O always specifies record lengths in I/O statements in bytes. By default, FORTRAN 77 direct-access unformatted I/O specifies the record length in words. This can cause incompatibilities when moving codes from FORTRAN 77 to Fortran 95 and vice versa. The `-byterecflen` option to the `f77(1)` command causes the FORTRAN 77 compiler to interpret all record lengths in bytes.

4.4.2.4 Special File Formats Differences

The MIPSpro FORTRAN 77 compiler permits you to specify the following two special modes in the `FORM=` clause of the `OPEN` statement:

- `FORM="BINARY"`, which permits reading and writing binary data from character variables.
- `FORM="SYSTEM"`, which allows input ignoring record boundaries.

These special modes are permitted in early releases of the MIPSpro 7 Fortran 90 compiler. However, neither form is supported by the Fortran 95 standard or by

the MIPSpro 7 Fortran 90 compiler, releases 7.2 and later. Either type of file access can also be achieved by using the `read(2)` and `write(2)` IRIX kernel functions.

4.4.2.5 Miscellaneous Differences

The following miscellaneous items in the MIPSpro Fortran 77 compiler are not in the MIPSpro 7 Fortran 90 compiler:

- Keyed access (`ACCESS= 'KEYED'`):
 - OPEN statement extensions:
 - `ACCESS= 'KEYED'` specifier
 - `KEY= ()` specifier
 - `RECORDSIZE=` specifier (same as `RECL`)
 - INQUIRE statement extension:
 - `KEYED=` specifier
 - READ statement extensions:
 - `KEY=` specifier
 - `KEYID=` specifier
 - DELETE statement
 - REWRITE statement
 - UNLOCK statement
- Direct access (`ACCESS = 'DIRECT'`):
 - OPEN statement extension:
 - `ASSOCIATEVARIABLE=` specifier to retain the `REC=` information.
 - `MAXREC=` specifier contains the maximum number of records allowed for a direct access file. If not present, no limit exists.
 - `RECORDSIZE=` specifier (same as `RECL`).
 - DEFINE FILE statement for a direct access unformatted file.
 - FIND statement for a direct access file.

- Record length of one byte indicates a binary data stream without record boundaries.
- Efficient reading of direct unformatted record in reverse order.
- Sequential access (ACCESS='SEQUENTIAL'). Extension for ENDFILE to allow multifile file by using a tab character rather than an EOF.
- stdin, stdout, and stderr I/O:
 - ACCEPT statement on stdin only. Same as a READ statement with an asterisk (*) as the unit.
 - OPEN statement extensions:


```
FILE='SYS$INPUT', 'SYS$OUTPUT', 'SYS$ERROR'
```
 - TYPE statement on stdout only; same as a PRINT statement.
 - Allow the double asterisk (**) to indicate stderr unit; that is, specifying WRITE (**, ...).
 - Special compile time option to allow reread from stdin after EOF is -vms_stdin.
- Unformatted I/O:
 - Allow unformatted I/O on an INTERNAL file.
 - Nonstandard aggregate reference in an I/O list is restricted to unformatted I/O and only one aggregate item is allowed.
- Formatted I/O:
 - Carriage control with an ASCII NUL to cause overprint with no advance; that is, it does not return to left margin after printing.
 - Variable format in which the width can be an expression that is evaluated at execution.
 - INQUIRE statement extension:


```
CARRIAGECONTROL= specifier
```
 - Allow the use of the letter Q for an exponent letter on input for a floating point value.
 - Special compile time option to interpret the carriage control character and replace it with the proper character in stdout.

- List-directed I/O:
 - No repeat counts are generated for list-directed output.
 - Allow the use of the letter Q for an exponent letter on input for a floating-point value.
- NAMELIST I/O:
 - No repeat counts are generated for namelist output.
 - Allow the use of the letter Q for an exponent letter on input for a floating point value.
- INQUIRE statement extensions:
 - ORGANIZATION= specifier to also contain ACCESS= specifications.
 - Use of different-sized variables for more specifiers than allowed in Fortran 95.
- OPEN statement extensions:
 - DEFAULTFILE= specifier to contain an alternate path or prefix for the opened unit. The specifier is concatenated with the string in the FILE= specifier or with the unit number if FILE= is absent.
 - DISP= specifier to designate how the file is handled after the file is closed. This is another form of the STATUS specifier on the CLOSE statement for Fortran 95.
 - FORM= 'BINARY' to indicate a formatted read with the A format on binary data. The implemented form of binary is unformatted I/O without record headers.
 - READONLY specifier indicates the file is read only. Fortran 95 provides the ACTION= specifier with more options.
 - RECORDTYPE= specifier to indicate type of records the file can contain. This may be present for direct, sequential, and keyed files.
 - SHARED= specifier indicates the file must be flushed after each record is written.
 - TYPE= specifier is the same as the Fortran 95 STATUS= specifier.

- Environment variable `FORTRAN_BUFFER_SIZE` to set buffer size and determine if direct I/O should be used.
- Interoperability with UNIX routines for I/O such as `fseek(3)`, `ftell(3)`, `flush(3)`, `fgetc(3)`, `fstat(2)`, and others. The PXF interface routines allow the use of these routines with a proper Fortran interface.
- Ability to query for the error status of the last I/O operation by unit.
- Allow MP-safe I/O at the subroutine level so non-MP-safe and MP-safe I/O can be mixed in one program but not in a program unit.

Data Representation and Storage [5]

This chapter shows how different data types are represented in storage and describes how the CRAY T90 and MIPSpro 7 Fortran 90 compilers use storage.

Numbers shown on the formats are bit positions, which represent powers of 2 in binary notation. Code that depends on internal representation is not portable and might not conform with the Fortran standard.

Note: Storage words are represented here with bits counted from the right, making bit 0 the low-order bit and bit 31 or 63 the high-order bit. This agrees with the convention used in the integer-type bit functions as well as the convention used in Cray hardware documentation. It does not agree with some conventions used in some other UNICOS and UNICOS/mk software documentation.

This chapter describes the machine representation of data. The last sections in this chapter describe storage issues, including overindexing.

5.1 Data Representation for UNICOS Systems

The following sections describe the representation of data on UNICOS systems, including CRAY T90 systems that support Cray floating-point arithmetic. These subsections do not describe data representation on CRAY T90 systems that support IEEE floating-point arithmetic. For information pertaining to CRAY T90 systems that support IEEE floating-point arithmetic, see Section 5.4, page 120.

5.1.1 Integer Type

All integer data is 64 bits (KIND=8), 2's complement.

When slower integer operations (`f90 -O nofastint`) are in effect, the range for `INTEGER(KIND=8)` operations is $-2^{63} < I < 2^{63}$ or approximately $-10^{18} < I < 10^{18}$.

When fast integer operations (`f90 -O fastint`) are in effect, which is the default, the range for `INTEGER(KIND=8)` operations is $-2^{46} < I < 2^{46}$ or approximately $-10^{13} < I < 10^{13}$.




Figure 1. Default 64-bit integers




Figure 2. Fast integer operations with `INTEGER (KIND=8)`, UNICOS systems (except CRAY T90 systems that support IEEE floating-point arithmetic)

To declare an entity to be of type integer, specify one of the following:

- `KIND=1`, `KIND=2`, `KIND=4`, or `KIND=8`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 1, 2, 4, or 8.

5.1.2 Real Type

Real (floating-point) numbers are represented in a packed representation of a binary mantissa and an exponent (power of 2). The bits in a Cray word are used to represent a real number as follows:




Figure 3. Real type

Notes on real data type representation:

The exponent is a power of 2, represented by a number that is 40000_8 higher than the actual value; this is called a *bias*. The effect of the bias is that the second bit in the word serves as the exponent's sign bit. This bit's usage is the inverse of the mantissa's sign bit, as follows:

Bit	Applies to	1 value indicates
63	Mantissa	Negative
62	Exponent	Positive

The exponent is represented by the second through sixth digits in an octal printout; these digits have the range 40000 through 57776_8 for a positive exponent, and 37777 through 20003_8 for a negative exponent.

When the bias is accounted for, the range of all exponents is as follows (notice the negative range is one smaller):

- 2^{-17775} to 2^{17776} (octal)
- or
- 2^{-8189} to 2^{8190} (decimal)

The mantissa is a 48-bit signed fraction. The sign of the mantissa is separated from the rest of the mantissa as shown in the preceding diagram. The mantissa is not complemented for negative values. That is, the mantissa for -10.0 is the same as for $+10.0$.

In terms of decimal values, the floating-point format of the CPU allows representation of numbers to about 15 significant decimal digits in the following approximate decimal range:

$$.367 \times 10^{-2465} < R < .273 \times 10^{2466}$$

A zero value is not biased and is represented as a word of all zeros.

Following are some sample numbers as represented within memory:

Decimal	Octal	Hexadecimal
10.0	040004500000000000000000	4004A0000000000000
-10.0	140004500000000000000000	C004A0000000000000

using it in a floating-point operation. To do this, add the unnormalized floating-point operand to 0. Compiler optimization suppresses an operation such as $X=X+0$. You can perform it with code such as the following:

```
DATA REALZERO /0./
X = X + REALZERO
```

5.1.3 Double-precision Type

A double-precision value is represented by 2 words. The first has the same format as the real type. The second word uses bits 0 through 47 as 48 additional bits of the mantissa. The other 16 bits of the second word must be zeros. Double-precision numbers can be in the following range:

- $2^{-8188} \text{ .LE. } R < 2^{8189}$
- or approximately
- $.367 \times 10^{-2465} < R < .273 \times 10^{2466}$




Figure 5. Double-precision type

To declare an entity to be of type double precision, specify one of the following:

- `REAL(KIND=16)`.
- `REAL(KIND=KIND(kind_expr))`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 16.

5.1.4 Single-precision Complex Type

A single-precision complex value is represented by 2 words, each of which has the same format as the real type. The first word represents the real part, and

the second represents the imaginary part. Each word has the same range as a real value.




Figure 6. Single-precision complex type

To declare an entity to be of single-precision complex type, specify one of the following:

- `KIND=4` or `KIND=8`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression that evaluates to 4 or 8.

Note that a complex data object with `KIND=4` has the same internal representation as a complex data object with `KIND=8`. Numeric inquiry functions on a complex data object with `KIND=4` return different values than on a complex data object with `KIND=8`. A numeric operation on a complex data object with `KIND=4` returns the same result as the same numeric operation on a complex data object with `KIND=8`.

5.1.5 Double-precision Complex Type

Values of double precision complex type are represented by 4 words. The first 2 words are the real part, and the last 2 words are the imaginary part. The real part and the imaginary part each have the same range as a double precision value.




Figure 7. Double-precision complex type (real portion)




Figure 8. Double-precision complex type (imaginary portion)

To declare an entity to be of double-precision complex type, specify one of the following:

- `KIND=16`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 16.

5.1.6 Character Type

Characters are represented by 8-bit ASCII codes packed eight per word.




Figure 9. Character type

The CF90 compiler does not support a nondefault character type. The only kind value supported is 1.

5.1.7 Logical Type

A logical variable uses one 64-bit word. Its value is true if the numeric value in the word is negative (typically, -1), and it is false if the numeric value in the word is nonnegative (typically, 0).

Note: Silicon Graphics does not guarantee a particular internal representation of logical values on any machine or system; the CF90 compiler is designed on the assumption that logical values will be used only as described in the Fortran standard. Therefore, it is not good programming practice to exploit gaps in type checking, such as between a function reference and its function value, to use logical values as numbers or vice versa.

To declare an entity to be of logical type, you can specify one of the following:

- `KIND=1`, `KIND=2`, `KIND=4`, or `KIND=8`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 1, 2, 4, or 8.

Note that logical entities with `KIND=1`, `KIND=2`, `KIND=4`, and `KIND=8` all occupy 64 bits.

5.1.8 Cray Character Pointers

Cray character pointers include a word address, bit offset, and bit length field.




Figure 10. 64-bit addressing for UNICOS systems (except CRAY T90 systems)




Figure 11. 64-bit addressing for CRAY T90 systems

5.2 Data Representation for IRIX systems

The following sections describe the representation of data on IRIX systems.


Note: On IRIX systems, `KIND=4` values are stored in 32 bits and can be packed two per word.

5.2.1 Integer Type

The following sections describe integer data representation of `KIND=1`, 2, 4, and 8 on IRIX systems.

5.2.1.1 `KIND=1`

Range: $-2^7 < I < 2^7$ or approximately $-10^2 < I < 10^2$

Figure 12. `INTEGER(KIND=1)` on IRIX systems

To declare 8-bit integers, specify one of the following:

- `KIND=1`.
- `KIND=KIND(kind_expr)`, where `kind_expr` is a scalar initialization expression with a kind type parameter that evaluates to 1.

5.2.1.2 KIND=2

Range: $-2^{15} < I < 2^{15}$ or approximately $-10^4 < I < 10^4$



Figure 13. INTEGER(KIND=2) on IRIX systems

To declare 16-bit integers, specify one of the following:

- KIND=2.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 2.

5.2.1.3 KIND=4

Range: $-2^{31} < I < 2^{31}$ or approximately $-10^9 < I < 10^9$




Figure 14. INTEGER(KIND=4) on IRIX systems

To declare 32-bit integers, specify one of the following:

- KIND=4.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 4.

5.2.1.4 KIND=8

Range: $-2^{63} < I < 2^{63}$ or approximately $-10^{18} < I < 10^{18}$




Figure 15. INTEGER(KIND=8) on IRIX systems

To declare 64-bit integers, specify one of the following:

- KIND=8.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 8.

5.2.2 Real Type

The following sections describe real data representation of KIND= 4, 8, and 16 on IRIX systems. Real (floating-point) numbers are represented in a packed representation of a sign, an exponent (power of 2), and a binary mantissa.

5.2.2.1 KIND=4

Range: -2^{-125} .LE. $I < 2^{128}$ or approximately -10^{-38} .LE. $I < 10^{38}$




Figure 16. REAL(KIND=4) on IRIX systems

To declare 32-bit reals, specify one of the following:

- KIND=4.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 4.

Notes on real data type representation:

The exponent is a power of 2, represented by a number that is 177_8 higher than the actual value; this is called a *bias*. The effect of the bias is that the second bit in the word serves as the exponent's sign bit. This bit's usage is the inverse of the mantissa's sign bit, as follows:

Bit	Applies to	1 value indicates
31	Mantissa	Negative
30	Exponent	Positive (> 0)

The exponent is represented by the second through ninth digits in a binary printout; these digits have the range 0111111_2 through 1111110_2 for a positive exponent, and 0000000_2 through 0111110_2 for a negative exponent.

When the bias is accounted for, the range of all exponents is as follows:

- 2^{-177} to 2^{177} (octal)
- or
- 2^{-127} to 2^{127} (decimal)

The mantissa is a 24-bit fraction with an assumed leading 1; that is, the leading 1 is not stored. The only exception is for the value 0, which has an assumed leading 0. The sign of the mantissa is separated from the rest of the mantissa as shown in the preceding diagram. The mantissa is not complemented for negative values. That is, the mantissa for -10.0 is the same as for $+10.0$.

In terms of decimal values, the 32-bit floating-point format allows representation of numbers to about 7 significant decimal digits in the following approximate decimal range:

$$1.18 \times 10^{-38} < R < 3.4 \times 10^{38}$$

A zero value is not biased and is represented as a word of all zeros.

The following are some sample numbers as represented within memory:

Decimal	Octal	Hexadecimal
10.0	010110000000	41200000
-10.0	030110000000	C1200000

Decimal	Octal	Hexadecimal
0.1	007563146315	3DCCCCD
-0.1	027563146315	BDCCCCD




Figure 17. Binary version of 10.0

The leftmost bit, with a 0 value, indicates a positive mantissa; that is, the real value is positive. The next 8 bits (1000010, or decimal 130) are the exponent. Subtracting the bias of 127 yields an exponent of 3, meaning that the binary fraction in the mantissa is multiplied by 2^3 . To express it another way, the binary point is moved 3 bits to the right from the mantissa’s highest bit. Interpreted this way, the first 4 bits of the mantissa, [1]010, indicate the real decimal value 10.0 (remember that there is an assumed 1 to the left of the mantissa in the IEEE floating-point format with a binary point to its immediate right). You can display other values by printing them with formats O11, Z8, or B32.

5.2.2.2 KIND=8

Double precision, `REAL(KIND=8)`, values are represented in 2 words on IRIX systems.

Range: -2^{-1021} .LE. $1 < 2^{1024}$ or approximately -10^{-308} .LE. $1 < 10^{308}$




Figure 18. `REAL(KIND=8)` on IRIX systems

To declare 64-bit reals, specify one of the following:

- `KIND=8`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 8.

5.2.2.3 `KIND=16`

Quad precision, `REAL(KIND=16)`, values are represented in 4 words on IRIX systems. For more information on quad precision representation IRIX systems, see `math(3M)`.

Range: -2^{-967} .LE. $I < 2^{1023}$ or approximately -10^{-292} .LE. $I < 10^{308}$




Figure 19. `REAL(KIND=16)` on IRIX systems

To declare 128-bit reals, specify one of the following:

- `KIND=16`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 16.

5.2.3 Complex Type

The following sections describe complex data representation of `KIND=4`, 8, and 16 on IRIX systems. A complex value has two parts. The first part represents the real part, and the second represents the imaginary part. Each word has the same range as a real value.

5.2.3.1 `KIND=4`

A single-precision, `KIND=4`, complex value is represented by 2 words. The first word represents the real part, and the second represents the imaginary part. Each word has the same range as a real value.

Range: $-2^{-125} \cdot \text{LE. I} < 2^{128}$ or approximately $-10^{38} \cdot \text{LE. I} < 10^{38}$




Figure 20. `COMPLEX(KIND=4)` on IRIX systems (real portion)




Figure 21. `COMPLEX(KIND=4)` on IRIX systems (imaginary portion)

To declare an entity to be of single-precision, complex type, specify one of the following:

- `KIND=4`.
- `KIND=KIND(kind_expr)`, where `kind_expr` is a scalar initialization expression with a kind type parameter that evaluates to 4.

5.2.3.2 KIND=8

A double-precision, KIND=8, complex value is represented by 4 words. The first 2 words represent the real part, and the second 2 words represent the imaginary part. Each word has the same range as a real value.

Range: -2^{-1021} .LE. $I < 2^{1024}$ or approximately -10^{308} .LE. $I < 10^{308}$




Figure 22. COMPLEX(KIND=8) on IRIX systems (real portion)




Figure 23. COMPLEX(KIND=8) on IRIX systems (imaginary portion)

To declare an entity to be of double-precision, complex type, specify one of the following:

- KIND=8.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 8.

5.2.3.3 KIND=16

A quad precision, KIND=16, complex value is represented by 8 words. The first 4 words represent the real part, and the second 4 words represent the imaginary part. Each word has the same range as a real value.

Range: -2^{-967} .LE. I < 2^{1023} or approximately -10^{-292} .LE. I < 10^{308}

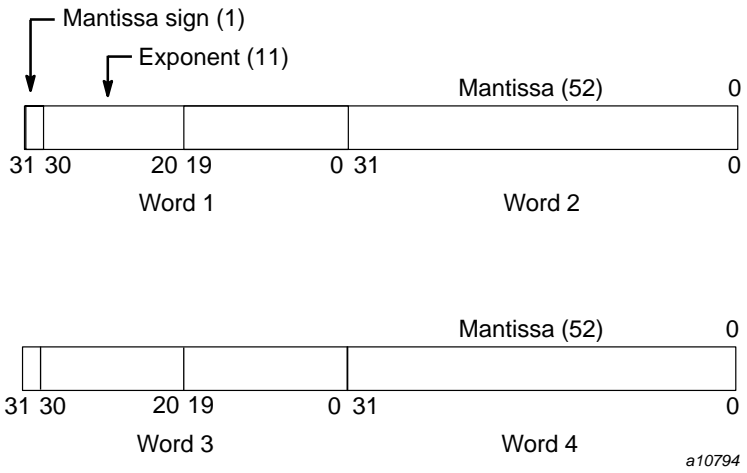


Figure 24. COMPLEX(KIND=16) on IRIX systems (real portion)




Figure 25. COMPLEX(KIND=16) on IRIX systems (imaginary portion)

To declare an entity to be of quad precision, complex type, specify one of the following:

- KIND=16.

- $KIND=KIND(kind_expr)$, where $kind_expr$ is a scalar initialization expression with a kind type parameter that evaluates to 16.

5.2.4 Character Type

Characters are represented by 8-bit ASCII codes. On IRIX systems, the codes are stored in 1 byte.




Figure 26. Character type

The MIPSpro 7 Fortran 90 compiler does not support a nondefault character type. The only kind value supported is 1.

5.2.5 Logical Type

Logical entities specified as $KIND=1$, $KIND=2$, and $KIND=4$ occupy 32 bits on IRIX systems. Logical entities specified as $KIND=8$ occupy 64 bits on IRIX systems. Its value is true if the numeric value in the word is one (1). Its value is false if the numeric value in the word is zero (0).

Note: Silicon Graphics does not guarantee a particular internal representation of logical values on any machine or system; the MIPSpro 7 Fortran 90 compiler is designed on the assumption that logical values will be used only as described in the Fortran standard. Therefore, it is not good programming practice to use logical values as numbers or vice versa.

To declare an entity to be of logical type, you can specify one of the following:

- $KIND=1$, $KIND=2$, $KIND=4$, or $KIND=8$.
- $KIND=KIND(kind_expr)$, where $kind_expr$ is a scalar initialization expression with a kind type parameter that evaluates to 1, 2, 4, or 8.

5.2.6 Cray Character Pointers (Deferred Implementation)

Cray character pointers include a byte address and a byte length field.




Figure 27. 32-bit addressing on IRIX systems

5.3 Data Representation for UNICOS/mk Systems

The following sections describe the representation of data on UNICOS/mk systems.


Note: On UNICOS/mk systems, `KIND=4` values are stored in 32 bits and can be packed two per word.

5.3.1 Integer Type

The following subsections describe integer data representation of `KIND=1`, `2`, `4`, and `8` on UNICOS/mk systems.

5.3.1.1 `KIND=1`, `KIND=2`, or `KIND=4`

Range: $-2^{31} < I < 2^{31}$ or approximately $-10^9 < I < 10^9$

Figure 28. Integer `KIND=1`, `2`, or `4` on UNICOS/mk systems

To declare 32-bit integers, specify one of the following:

- `KIND=1`, `KIND=2`, or `KIND=4`.
- `KIND=KIND(kind_expr)`, where `kind_expr` is a scalar initialization expression with a kind type parameter that evaluates to `1`, `2`, or `4`.

5.3.1.2 KIND=8

Range: $-2^{63} < I < 2^{63}$ or approximately $-10^{18} < I < 10^{18}$




Figure 29. INTEGER(KIND=8) on UNICOS/mk systems

To declare 64-bit integers, specify one of the following:

- KIND=8.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 8.

5.3.2 Real Type

The following sections describe real data representation of KIND=4 and 8. Real (floating-point) numbers are represented in a packed representation of a sign, an exponent (power of 2), and a binary mantissa.

5.3.2.1 KIND=4

Range: $-2^{-125} \text{ .LE. } I < 2^{128}$ or approximately $-10^{-38} \text{ .LE. } I < 10^{38}$




Figure 30. REAL(KIND=4) on UNICOS/mk systems

To declare 32-bit reals, specify one of the following:

- KIND=4.

- $KIND=KIND(kind_expr)$, where $kind_expr$ is a scalar initialization expression with a kind type parameter that evaluates to 4.

Notes on real data type representation:

The exponent is a power of 2, represented by a number that is 177_8 higher than the actual value; this is called a *bias*. The effect of the bias is that the second bit in the word serves as the exponent's sign bit. This bit's usage is the inverse of the mantissa's sign bit, as follows:

Bit	Applies to	1 value indicates
31	Mantissa	Negative
30	Exponent	Positive (> 0)

The exponent is represented by the second through ninth digits in a binary printout; these digits have the range 0111111_2 through 1111110_2 for a positive exponent, and 0000000_2 through 0111110_2 for a negative exponent.

When the bias is accounted for, the range of all exponents is as follows:

- 2^{-177} to 2^{177} (octal)
- or
- 2^{-127} to 2^{127} (decimal)

The mantissa is a 24-bit fraction with an assumed leading 1; that is, the leading 1 is not stored. The only exception is for the value 0, which has an assumed leading 0. The sign of the mantissa is separated from the rest of the mantissa as shown in the preceding diagram. The mantissa is not complemented for negative values. That is, the mantissa for -10.0 is the same as for $+10.0$.

In terms of decimal values, the 32-bit floating-point format allows representation of numbers to about 7 significant decimal digits in the following approximate decimal range:

$$1.18 \times 10^{-38} < R < 3.4 \times 10^{38}$$

A zero value is not biased and is represented as a word of all zeros.

The following are some sample numbers as represented within memory:

Decimal	Octal	Hexadecimal
10.0	010110000000	41200000
-10.0	030110000000	C1200000
0.1	007563146315	3DCCCCCD
-0.1	027563146315	BDCCCCCD

0100000100100000000000000000000000000000
 ↑
 Bit 22 a10801

Figure 31. Binary version of 10.0

The leftmost bit, with a 0 value, indicates a positive mantissa; that is, the real value is positive. The next 8 bits (1000010, or decimal 130) are the exponent. Subtracting the bias of 127 yields an exponent of 3, meaning that the binary fraction in the mantissa is multiplied by 2^3 ; to express it another way, the binary point is moved 3 bits to the right from the mantissa's highest bit. Interpreted this way, the first 4 bits of the mantissa, [1]010, indicate the real decimal value 10.0; remember that there is an assumed 1 to the left of the mantissa in the IEEE floating-point format with a binary point to its immediate right. You can display other values by printing them with formats O11, Z8, or B32.

5.3.2.2 KIND=8

Range: -2^{-1021} .LE. $I < 2^{1024}$ or approximately -10^{-308} .LE. $I < 10^{308}$




Figure 32. REAL(KIND=8) on UNICOS/mk systems

To declare 64-bit reals, specify one of the following:

- `KIND=8`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 8.

5.3.3 Complex Type

The following sections describe complex data representation of `KIND=4` and `KIND=8` on UNICOS/mk systems. A complex value has two parts. The first part represents the real part, and the second represents the imaginary part. Each word has the same range as a real value.

5.3.3.1 `KIND=4`

A `KIND=4` complex value consists of 2 parts. The first part represents the real portion, and the second represents the imaginary portion. Each part has the same range as a 32-bit (or `KIND=4`) real value.

Range: -2^{-125} .LE. $1 < 2^{128}$ or approximately -10^{-38} .LE. $1 < 10^{38}$




Figure 33. `COMPLEX(KIND=4)` on UNICOS/mk systems (real portion)




Figure 34. `COMPLEX(KIND=4)` on UNICOS/mk systems (imaginary portion)

To declare an entity to be of complex type with a total length of 64 bits, specify one of the following:

- `KIND=4`.
- `KIND=KIND(kind_expr)`, where `kind_expr` is a scalar initialization expression with a kind type parameter that evaluates to 4.

5.3.3.2 `KIND=8`

A single-precision, `KIND=8`, complex value is represented by 2 words. The first word represents the real part, and the second represents the imaginary part. Each word has the same range as a 64-bit (or `KIND=8`) real value.

Range: -2^{-1021} .LE. $1 < 2^{1024}$ or approximately -10^{-308} .LE. $1 < 10^{308}$




Figure 35. `COMPLEX(KIND=8)` on UNICOS/mk systems (real portion)



Figure 36. `COMPLEX(KIND=8)` on UNICOS/mk systems (imaginary portion)

To declare an entity to be of single-precision, complex type, specify one of the following:

- `KIND=8`.

- $KIND=KIND(kind_expr)$, where $kind_expr$ is a scalar initialization expression with a kind type parameter that evaluates to 8.

5.3.4 Character Type

Characters are represented by 8-bit ASCII codes. On UNICOS/mk systems, the codes are packed 8 per word.




Figure 37. Character type

The CF90 compiler does not support a nondefault character type. The only kind value supported is 1.

5.3.5 Logical Type

A logical variable uses one word. Its value is true if the numeric value in the word is nonzero, and it is false if the numeric value in the word is zero.

Note: Silicon Graphics does not guarantee a particular internal representation of logical values on any machine or system; the CF90 compiler is designed on the assumption that logical values will be used only as described in the Fortran standard. Therefore, it is not good programming practice to use logical values as numbers or vice versa.

To declare an entity to be of logical type, you can specify one of the following:

- $KIND=1$, $KIND=2$, $KIND=4$, or $KIND=8$.
- $KIND=KIND(kind_expr)$, where $kind_expr$ is a scalar initialization expression with a kind type parameter that evaluates to 1, 2, 4, or 8.

On UNICOS/mk systems, all $KIND=1$, 2, and 4 occupy 32 bits. The $KIND=8$ specification occupies 64 bits.

5.3.6 Cray Character Pointers

Cray character pointers include a byte address and a byte length field. On UNICOS/mk systems, character pointers are 128-bit objects, as follows:




Figure 38. Cray character pointers on UNICOS/mk systems

5.4 Data Representation for CRAY T90 Systems That Support IEEE Floating-point Arithmetic

The following sections describe the representation of data on CRAY T90 systems that support IEEE floating-point arithmetic.

5.4.1 Integer Type

The following sections describe integer data representation of `KIND=1`, `2`, `4`, and `8` on CRAY T90 systems that support IEEE floating-point arithmetic.

5.4.1.1 `KIND=1`, `KIND=2`, or `KIND=4`

Range: $-2^{31} < I < 2^{31}$ or approximately $-10^9 < I < 10^9$




Figure 39. Integer `KIND=1`, `2`, or `4` on CRAY T90 systems that support IEEE floating-point arithmetic

To declare 32-bit integers, specify one of the following:

- `KIND=1`, `KIND=2`, or `KIND=4`.

- $KIND=KIND(kind_expr)$, where $kind_expr$ is a scalar initialization expression with a kind type parameter that evaluates to 1, 2, or 4.

5.4.1.2 $KIND=8$

By default, the range for $INTEGER(KIND=8)$ operations is $-2^{63} < I < 2^{63}$ or approximately $-10^{18} < I < 10^{18}$. When fast integer operations are specified on the `f90(1)` command line, the range for $INTEGER(KIND=8)$ operations is $-2^{50} < I < 2^{50}$ or approximately $-10^{15} < I < 10^{15}$.




Figure 40. Default $INTEGER(KIND=8)$ on CRAY T90 systems that support IEEE floating-point arithmetic




Figure 41. Fast operations with $INTEGER(KIND=8)$ on CRAY T90 systems that support IEEE floating-point arithmetic

To declare 64-bit integers, specify one of the following:

- $KIND=8$.
- $KIND=KIND(kind_expr)$, where $kind_expr$ is a scalar initialization expression with a kind type parameter that evaluates to 8.

5.4.2 Real Type

The following sections describe real data representation of $KIND=4$, 8, and 16 on CRAY T90 systems that support IEEE floating-point arithmetic. Real (floating-point) numbers are represented in a packed representation of a sign, an exponent (power of 2), and a binary mantissa.

5.4.2.1 KIND=4 and KIND=8

Range: -2^{-1021} .LE. $1 < 2^{1024}$ or approximately -10^{-308} .LE. $1 < 10^{308}$




Figure 42. Real KIND=4 or 8 on CRAY T90 systems that support IEEE floating-point arithmetic

To declare 64-bit reals, specify one of the following:

- KIND=4 or KIND=8.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 4 or 8.

For additional information on how real data is represented on CRAY T90 systems that support IEEE floating-point arithmetic, see "Notes on real data type representation" in Section 5.3.2.1, page 114. The information presented there for UNICOS/mk systems applies to CRAY T90 systems that support IEEE floating-point arithmetic.

Note that a real data object with KIND=4 has the same internal representation as a real data object with KIND=8. Numeric inquiry functions on a real data object with KIND=4 return different values than on a real data object with KIND=8. A numeric operation on a real data object with KIND=4 returns the same result as the same numeric operation on a real data object with KIND=8.

5.4.2.2 KIND=16

Double precision, REAL(KIND=16), values are represented in 2 words on CRAY T90 systems that support IEEE floating-point arithmetic.

Range: -2^{-16381} .LE. $1 < 2^{16384}$ or approximately -10^{-4932} .LE. $1 < 10^{4932}$

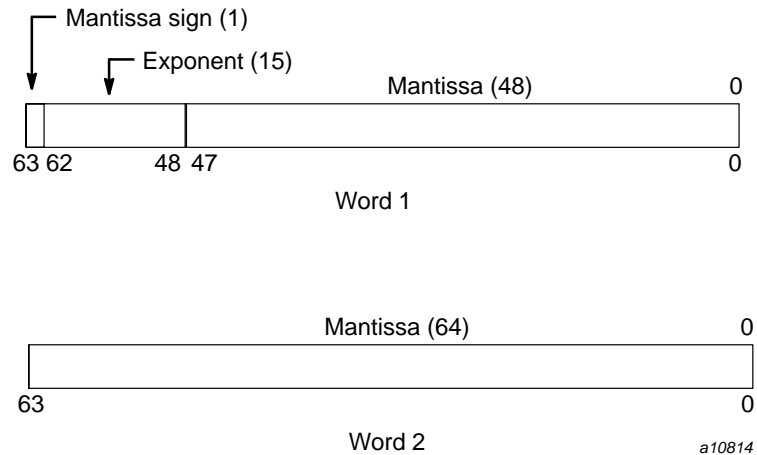


Figure 43. REAL(KIND=16) on CRAY T90 systems that support IEEE floating-point arithmetic

To declare 64-bit reals, specify one of the following:

- KIND=16.
- KIND=KIND(*kind_expr*), where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 16.

5.4.3 Complex Type

The following sections describe complex data representation of KIND=4, 8, and 16 on CRAY T90 systems that support IEEE floating-point arithmetic. A complex value has two parts. The first part represents the real part, and the second represents the imaginary part. Each word has the same range as a real value.

5.4.3.1 KIND=4 and KIND=8

A single-precision, KIND=4 or KIND=8, complex value is represented by 2 words. The first word represents the real part, and the second represents the imaginary part. Each word has the same range as a real value.

Range: -2^{-1021} .LE. $1 < 2^{1024}$ or approximately -10^{-308} .LE. $1 < 10^{308}$




Figure 44. Complex `KIND=8` or `4` on CRAY T90 systems that support IEEE floating-point arithmetic (real portion)




Figure 45. Complex `KIND=8` or `4` on CRAY T90 systems that support IEEE floating-point arithmetic (imaginary portion)

To declare an entity to be of single-precision, complex type, specify one of the following:

- `KIND=4` or `KIND=8`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 4 or 8.

Note that a complex data object with `KIND=4` has the same internal representation as a complex data object with `KIND=8`. Numeric inquiry functions on a complex data object with `KIND=4` return different values than on a complex data object with `KIND=8`. A numeric operation on a complex data object with `KIND=4` returns the same result as the same numeric operation on a complex data object with `KIND=8`.

5.4.3.2 `KIND=16`

A double-precision, `KIND=16`, complex value is represented by 4 words. The first two words represent the real part, and the second two words represent the imaginary part. Each word has the same range as a real value.

Range: -2^{-16381} .LE. $1 < 2^{16384}$ or approximately -10^{-4932} .LE. $1 < 10^{4932}$




Figure 46. `COMPLEX(KIND=16)` on CRAY T90 systems that support IEEE floating-point arithmetic (real portion)

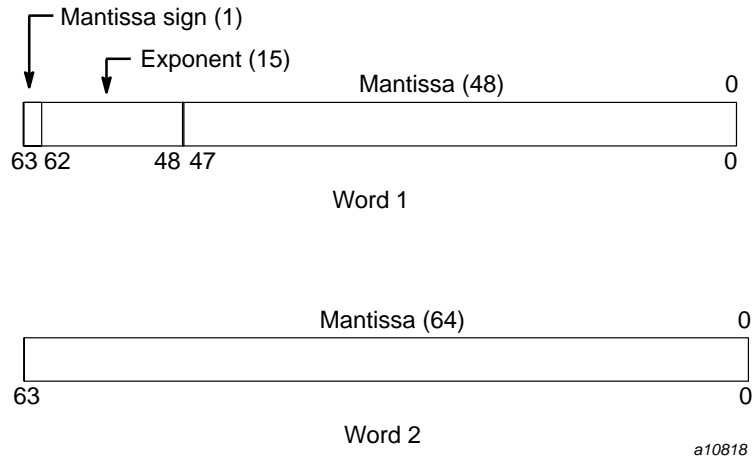


Figure 47. `COMPLEX(KIND=16)` on CRAY T90 systems that support IEEE floating-point arithmetic (imaginary portion)

To declare an entity to be of double-precision, complex type, specify one of the following:

- `KIND=16`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 16.

5.4.4 Character Type

Characters are represented by 8-bit ASCII codes. The codes are packed 8 per word.




Figure 48. Character type

The CF90 compiler does not support a nondefault character type. The only kind value supported is 1.

5.4.5 Logical Type

A logical variable uses one word. Its value is true if the numeric value in the word is nonzero, and it is false if the numeric value in the word is zero.

Note: Silicon Graphics does not guarantee a particular internal representation of logical values on any machine or system; the CF90 compiler is designed on the assumption that logical values will be used only as described in the Fortran standard. Therefore, it is not good programming practice to use logical values as numbers or vice versa.

To declare an entity to be of logical type, you can specify one of the following:

- `KIND=1`, `KIND=2`, `KIND=4`, or `KIND=8`.
- `KIND=KIND(kind_expr)`, where *kind_expr* is a scalar initialization expression with a kind type parameter that evaluates to 1, 2, 4, or 8.

On CRAY T90 systems that support IEEE floating-point arithmetic, all `KIND=1`, `KIND=2`, and `KIND=4` specifications occupy 32 bits. The `KIND=8` specification occupies 64 bits.

5.4.6 Cray Character Pointers

Cray character pointers are two words in length. The first word includes an offset and an address. The second word includes the byte length field.




Figure 49. Cray character pointer for CRAY T90 systems that support IEEE floating-point arithmetic

5.5 Storage Issues

This section describes how the CF90 and MIPSpro 7 Fortran 90 compilers use storage, including how these compilers accommodate programs that use overindexing.

Note: The information in this section assumes that you are using the default data representations.

On UNICOS/mk systems, specifying `-i 32` or `-s default32` on the `f90(1)` command line changes the storage and data representation of all noncharacter data types. This affects data that is storage sequence-associated. Mixing data types is not recommended when these command line options are used.

On IRIX systems, the following options to the `f90(1)` command affect storage and data representation:

- `-d16` changes default double precision and double complex to 128 bits
- `-i4` changes default integer and logical to 32 bits
- `-i8` changes default integer and logical to 64 bits
- `-n32` and `-64` change pointer sizes and the maximum amount of addressable memory
- `-r4` changes default real and complex to 32 bits/64 bits
- `-r8` changes default real and complex to 64 bits/128 bits

5.5.1 Storage Units and Sequences

A *numeric storage unit* can be one of the following:

- A word on UNICOS and UNICOS/mk systems of 64 bits.
- A word on IRIX systems of 32 bits.

A *character storage unit* is an 8-bit byte.

A *storage sequence* is a contiguous group of storage units with a consecutive series of addresses. Each array and each common block is stored in a storage sequence. The size of a storage sequence is the number of storage units it contains. Two storage sequences are *associated* if they share at least one storage unit.

All nondefault types have an unspecified storage unit. The `-s default32` option on the `f90(1)` command line changes the number of bits in a numeric storage unit for UNICOS/mk systems. There is no longer a relationship between storage units after this command line option is used.

The following list shows the storage units for the default types on UNICOS systems:

<u>Type</u>	<u>Storage units</u>
Integer	1
Real (single precision)	1
Real (double precision)	2
Complex	2
Logical	1

Complex values occupy twice the storage of real values. The real portion of the complex value occupies the first half of the total storage; the imaginary portion of the complex value occupies the second half of the total storage, as follows:

- On UNICOS and UNICOS/mk systems, a double precision or complex value ($KIND=4$ or $KIND=8$) uses a storage sequence of two numeric storage units. The first storage unit contains the most significant bits of a double-precision value or the real part of a complex value. The second storage unit contains the least significant bits of a double-precision value or the imaginary part of a complex value. Double precision and double complex data types are not supported on UNICOS/mk systems.

On IRIX systems, a double-precision value uses a storage sequence of 8 or 16 bytes. Depending on the $KIND=$ specification, a complex value uses 8, 16, or 32 bytes. The first half of the bytes used contains the most significant bits of a double-precision value or the real part of a complex value. The last half of the bytes used contains the least significant bits of a double-precision value or the imaginary part of a complex value.

- On UNICOS and UNICOS/mk systems, a double-complex value occupies 4 words of storage; the first 2 words contain the real part of the complex value, and the second 2 words contain the imaginary part.

On IRIX systems, a double-complex value occupies 16 bytes of storage; the first 8 bytes contain the real part of the complex value, and the second 8 bytes contain the imaginary part.

On IRIX systems, a quad precision complex value occupies 32 bytes of storage; the first 16 bytes contain the real part of the complex value, and the second 16 bytes contain the imaginary part.

A character value is represented as an 8-bit ASCII code, packed 8 characters per word on UNICOS and UNICOS/mk systems; this value is packed 4 characters per byte on IRIX systems. The storage size depends on the length specification of the value.

ANSI/ISO: The Fortran standard does not specify the relationship between storage units and computer words, and it does not specify any relation between default numeric and character storage units.

5.5.2 Static and Stack Storage

With static storage, any variable that is allocated memory occupies the same address throughout program execution. Allocation is determined before program execution.

Code using static storage can be used with Autotasking, multitasking, and macrotasking if variables in static storage conform to the following guidelines:

- Loops are Autotasked regardless of the presence of variables in static or stack storage. Scoping is controlled by the presence of `PRIVATE` or `SHARED` parameters on the `DOALL` Autotasking directive. If a subroutine that contains static data is called from within an autotasked loop, static data is treated as shared data, which means that the static data must be protected by `GUARD` and `ENDGUARD` Autotasking directives.
- Variables in static storage can be read when loops are multitasked and macrotasked. If a loop modifies variables in static storage, you must use guards (`GUARD` and `ENDGUARD` Autotasking directives) or locks (`LOCKON ()` and `LOCKOFF ()` calls) to protect the variables.

For more information on Autotasking directives, see the *CF90 Commands and Directives Reference Manual*, or the *MIPSpro 7 Fortran 90 Commands and Directives Reference Manual*. For more information on locks, see the `LOCKON(3F)` or `LOCKOFF(3F)` man pages.

Stack storage is the default storage allocation for the CF90 compiler on UNICOS and UNICOS/mk systems. On IRIX systems, stack storage is the MIPSpro 7 Fortran 90 default for all subprograms, but static storage is the default for items that require 256 bits of storage in a main program. The stack is an area of memory where storage for variables is allocated when a subprogram or procedure begins execution. These variables are released when execution completes. The stack expands and contracts as procedures are entered and exited. Autotasking and recursion require a stack.

When stack storage is used, the value of a variable is not saved between invocations of a subprogram unless it is specified in a `SAVE` or `DATA` statement. When `f90 -e v` (UNICOS and UNICOS/mk systems) or `f90 -static` (IRIX systems) is specified, all user variables are treated as if they appeared in a `SAVE`

statement. When `-e v` or `-static` is in effect, compiler-generated temporary variables and the calling sequence are still allocated to the stack.

Note: If `f90 -e i` is specified, variables are reset for each invocation of a subprogram, even in static storage. Therefore, the `SAVE` or `DATA` statement is necessary to preserve the value of a variable between invocations. This information applies only to UNICOS and UNICOS/mk systems.

The way in which the amount of memory available for the stack is determined depends on your platform. On UNICOS and UNICOS/mk systems, it is determined by the `STACK` directive, available with the `segldr(1)` or `cld(1)` loaders; see the `segldr(1)` or `cld(1)` man pages for more information. On IRIX systems, you can use the `limit(1)` command to change the amount of stack space that a program is allowed; see the `limit(1)` man page for more information.

A *heap* is memory that, like a stack, is dynamically allocated; it is used internally.

The CF90 and MIPSpro 7 Fortran 90 compilers allocate variables to storage according to the following criteria:

- Variables in common blocks are always allocated in the order in which they appear in `COMMON` statements.
- Data in modules are statically allocated.
- User variables that are defined or referenced in a program unit, and that also appear in `SAVE` or `DATA` statements, are allocated to static storage, but not necessarily in the order shown in your source program.
- Other referenced user variables are assigned to the stack. If `-e v` (UNICOS and UNICOS/mk systems) or `-static` (IRIX systems) is specified on the `f90(1)` command line, referenced variables are allocated to static storage. This allocation does not necessarily depend on the order in which the variables appear in your source program.
- Compiler-generated variables are assigned to a register or to memory (to the stack or heap), depending on how the variable is used. Compiler-generated variables include `DO`-loop trip counts, dummy argument addresses, temporaries used in expression evaluation, argument lists, and variables storing adjustable dimension bounds at entries.
- Automatic objects may be allocated to either the stack or to the heap, depending on how much stack space is available when the objects are allocated.

- Heap or stack allocation can be used for `TASK COMMON` variables and some compiler-generated temporary data such as automatic arrays and array temporaries.

Note: Unreferenced user variables not appearing in `COMMON` statements are not allocated.

5.5.3 Dynamic Memory Allocation (UNICOS Systems Only)

Many FORTRAN 77 programs contain a memory allocation scheme that expands an array in a common block located in central memory at the end of the program. This practice of expanding a blank common block or expanding a dynamic common block (sometimes referred to as *overindexing*) causes conflicts between user management of memory and the dynamic memory requirements of UNICOS libraries. It is recommended that you modify programs rather than expand blank common blocks, particularly when migrating from other environments.

Figure 50, page 133, shows the structure of a program under the UNICOS operating system in relation to expanding a blank common block. In both figures, the user area includes code, data, and common blocks.

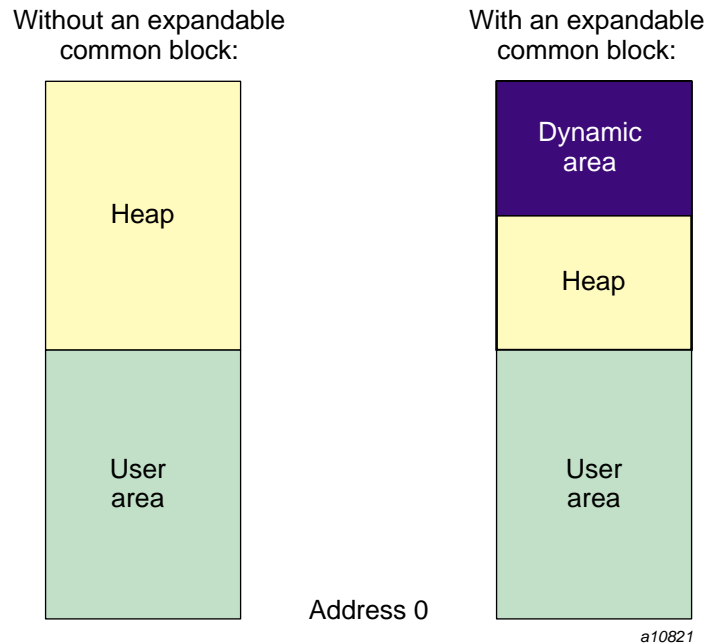


Figure 50. Memory use under UNICOS

There are two ways to change your code. The standard method, shown in Section 5.5.3.1 is preferred.

5.5.3.1 Changing Your Code: Standard Method

You can use the `ALLOCATE` statement to dynamically allocate an array. Use the following three-step process:

1. For arrays that expand in a common block, define Fortran allocatable arrays in a Fortran module.
2. Replace the common block definition in all source files that use the global array with a `USE` statement.
3. Use the `ALLOCATE` statement in place of any calls to the `MEMORY` routine.

Original code:

```
PROGRAM TEST
C Puts array X in blank common:
```

```

COMMON X(1)
...
C Adds 100000 words to blank common:
CALL MEMORY ('UC',100000)
...
DO 10, I=1,100000
X(I) = RANF()
10 CONTINUE
...

```

Converted code (after steps 1 and 2):

```

MODULE GLOBAL_DATA ! STEP 1
REAL, SAVE, ALLOCATABLE :: X(:)
END MODULE
...
PROGRAM TEST
USE GLOBAL_DATA ! STEP 2
LIMIT = 100000
ALLOCATE (X(LIMIT)) ! STEP 3
...
DO 10 I = 1,LIMIT
X(I) = RANF()
10 CONTINUE
...
END

```

5.5.3.2 Changing Your Code: Nonstandard Method

The nonstandard way to change your program is by using the following two-step process:

1. For arrays that expand in a common block, define Cray pointers that point to the first address in each array.
2. Change any calls to memory to calls to library routine HPALLOC(3).

Original code:

```

PROGRAM TEST
C Puts array X in blank common:
COMMON X(1)
...
C Adds 100000 words to blank common:
CALL MEMORY ('UC',100000)

```



```
...
      DO 10, I=1,100000
         X(I) = RANF()
10     CONTINUE
...
```

Converted code (after steps 1 and 2):

```
      PROGRAM TEST
      COMMON /WORK/ IPTR
...
C Establish array location at runtime:
      POINTER (IPTR,X(1))
...
C Effective common block size:
      CALL HPALLOC (IPTR,100000,ERRCODE,0)
...
      DO 10 I=1,100000
         X(I) = RANF()
10     CONTINUE
...
```


Outmoded Features [6]

This chapter describes outmoded Fortran features that the CF90 and MIPSpro 7 Fortran 90 compilers support. These features have been replaced by alternatives that enhance the portability of CF90 and MIPSpro 7 Fortran 90 programs. None of the outmoded features described in this chapter were part of any Fortran standard; they were extensions supported in older Cray Research compilers. The outmoded features and their preferred alternatives are listed in Table 4.

Table 4. Outmoded features and preferred alternatives

Outmoded feature	Preferred alternative
Hollerith data	Character data.
ENCODE and DECODE	Internal files.
Asterisk character constant delimiters in formats	Apostrophe or quotation mark delimiters.
$[-b]X$ edit descriptor	TL edit descriptor, 1X.
A descriptor used for noncharacter data and R descriptor	Character type and other conventional matchings of data with descriptors.
EOF, IEOF, and IOSTAT functions	End-of-file specifier (END=) or status specifier (IOSTAT=).
Initialization using long strings	Replace the numeric target with a character item. Replace a Hollerith constant with a character constant.
IMPLICIT UNDEFINED	IMPLICIT NONE
Type statements with $*n$	Standard type statements (KIND=).
Two-branch arithmetic IF	IF construct or IF statement.
TASK COMMON statement	TASKCOMMON compiler directive.
Indirect logical IF	IF construct or IF statement.
Nested loops ending with a single, labeled END DO	One END DO statement for each loop.
DOUBLE COMPLEX statement and related specific intrinsic function names	COMPLEX (KIND=) statement and standard intrinsic functions. See Section 6.9, page 152, for more information.

Outmoded feature	Preferred alternative
Bitwise intrinsic functions	Standard intrinsic functions. See Section 6.10, page 153, for more information.
CLOCK(3I), DATE(3I), and JDATE(3I) intrinsic functions	DATE_AND_TIME(3I) intrinsic subroutine.
DCOT(3M) intrinsic function	COT(3M) intrinsic function.
DFLOAT(3M) and DREAL(3M) intrinsic functions	REAL(3M) intrinsic function.
NUMARG(3I) intrinsic function	PRESENT(3I) intrinsic function for optional arguments.
RANF(3I) and RANGET(3I) intrinsic functions	RANDOM_NUMBER(3I) intrinsic subroutine.
RANSET(3I) intrinsic function	RANDOM_SEED(3I) intrinsic subroutine.
RTC(3I) intrinsic function	SYSTEM_CLOCK(3I) intrinsic subroutine.

6.1 Hollerith Type

Hollerith data is a sequence of any characters capable of internal representation as specified in Table 3, page 77. Its length is the number of characters in the sequence, including blank characters. Each character occupies a position within the storage sequence identified by one of the numbers 1, 2, 3, . . . indicating its placement from the left (position 1). Hollerith data must contain at least one character.

6.1.1 Hollerith Constants

A Hollerith constant is expressed in one of three forms. The first of these is specified as a nonzero integer constant followed by the letter H, L, or R and as many characters as equal the value of the integer constant. The second form of Hollerith constant specification delimits the character sequence between a pair of apostrophes followed by the letter H, L, or R. The third form is like the second, except that quotation marks replace apostrophes. For example:

```
Character sequence:   ABC 12
Form 1:               6HABC 12
Form 2:               'ABC 12'H
Form 3:               "ABC 12"H
```

Two adjacent apostrophes or quotation marks appearing between delimiting apostrophes or quotation marks are interpreted and counted by the compiler as

a single apostrophe or quotation mark within the sequence. Thus, the sequence `DON'T USE "*"` would be specified with apostrophe delimiters as `'DON' 'T USE "*" 'H`, and with quotation mark delimiters as `"DON'T USE """"H`.

Each character of a Hollerith constant is represented internally by an 8-bit code, with up to 32 such codes allowed. This limit corresponds to the size of the largest numeric type, `COMPLEX(KIND = 16)`. The ultimate size and makeup of the Hollerith data depends on the context. If the Hollerith constant is larger than the size of the type implied by context, the constant is truncated to the appropriate size. If the Hollerith constant is smaller than the size of the type implied by context, the constant is padded with a character dependent on the Hollerith indicator. When an `H` Hollerith indicator is used, the truncation and padding is done on the right end of the constant. The pad character is the blank character code (20).

Null codes can be produced in place of blank codes by substituting the letter `L` for the letter `H` in the Hollerith forms described above. The truncation and padding is also done on the right end of the constant, with the null character code (00) as the pad character.

Using the letter `R` instead of the letter `H` as the Hollerith indicator means truncation and padding is done on the left end of the constant with the null character code (00) used as the pad character.

All of the following Hollerith constants yield the same Hollerith constant and differ only in specifying the content and placement of the unused portion of the single 64-bit entity containing the constant:

Hollerith constant	Internal byte, beginning on bit:							
	0	8	16	24	32	40	48	56
<code>6HABCDEF</code>	A	B	C	D	E	F	20 ₁₆	20 ₁₆
<code>'ABCDEF'H</code>	A	B	C	D	E	F	20 ₁₆	20 ₁₆
<code>"ABCDEF" H</code>	A	B	C	D	E	F	20 ₁₆	20 ₁₆
<code>6LABCDEF</code>	A	B	C	D	E	F	00	00
<code>'ABCDEF'L</code>	A	B	C	D	E	F	00	00
<code>"ABCDEF"L</code>	A	B	C	D	E	F	00	00
<code>6RABCDEF</code>	00	00	A	B	C	D	E	F

Hollerith constant	Internal byte, beginning on bit:							
	0	8	16	24	32	40	48	56
'ABCDEF'R	00	00	A	B	C	D	E	F
"ABCDEF"R	00	00	A	B	C	D	E	F

A Hollerith constant is limited to 32 characters except when specified in a `CALL` statement, a function argument list, or a `DATA` statement. An all-zero computer word follows the last word containing a Hollerith constant specified as an actual argument in an argument list.

A character constant of 32 or fewer characters is treated as if it were a Hollerith constant in situations where a character constant is not allowed by the standard but a Hollerith constant is allowed by the CF90 and MIPSpro 7 Fortran 90 compilers. If the character constant appears in a `DATA` statement value list, it can be longer than 32 characters.

6.1.2 Hollerith Values

A *Hollerith value* is a Hollerith constant or a variable that contains Hollerith data. A Hollerith value is limited to 32 characters.

A Hollerith value can be used in any operation in which a numeric constant can be used. It can also appear on the right-hand side of an assignment statement in which a numeric constant can be used. It is truncated or padded to be the correct size for the type implied by the context.

6.1.3 Hollerith Relational Expressions

Used with a relational operator, the Hollerith value e_1 is less than e_2 if its value precedes the value of e_2 in the collating sequence and is greater if its value follows the value of e_2 in the collating sequence.

The following examples are evaluated as true if the integer variable `LOCK` contains the Hollerith characters `K`, `E`, and `Y` in that order and left-justified with five trailing blank character codes:

```

3HKEY.EQ.LOCK
'KEY'.EQ.LOCK
LOCK.EQ.LOCK
'KEY1'.GT.LOCK
'KEY0'H.GT.LOCK

```

6.2 Formatted I/O and Internal Files

A formatted I/O operation defines entities by transferring data between I/O list items and records of a file. The file can be on an external media or in internal storage.

The Fortran standard provides READ and WRITE statements for both formatted external and internal file I/O. This is the preferred method for formatted internal file I/O. It is the only method for list-directed internal file I/O.

The ENCODE and DECODE statements are an alternative to standard Fortran READ and WRITE statements for formatted internal file I/O.

An internal file in standard Fortran I/O must be declared as character, while the internal file in ENCODE and DECODE statements can be any data type. A record in an internal file in standard Fortran I/O is either a scalar character variable or an array element of a character array. The record size in an internal file in an ENCODE or DECODE statement is independent of the storage size of the variable used as the internal file. If the internal file is a character array in standard Fortran I/O, multiple records can be read or written with internal file I/O. The alternative form does not provide the multiple record capability.

6.2.1 ENCODE Statement

The ENCODE statement provides a method of converting or encoding the internal representation of the entities in the output list to a character representation. The format of the ENCODE statement is as follows:

<pre>ENCODE (<i>n</i>, <i>f</i>, <i>dest</i>) [<i>elist</i>]</pre>
--

<i>n</i>	Number of characters to be processed. Nonzero integer expression not to exceed the maximum record length for formatted records. This is the record size for the internal file.
<i>f</i>	Format identifier. It cannot be an asterisk.

- dest* Name of internal file. It can be a variable or array of any data type. It cannot be an array section, a zero-sized array, or a zero-sized character variable.
- elist* Output list to be converted to character during the ENCODE statement.

The output list items are converted using format *f* to produce a sequence of *n* characters that are stored in the internal file *dest*. The *n* characters are packed 8 characters per word on UNICOS and UNICOS/mk systems. The *n* characters are packed 4 characters per word on IRIX systems.

An ENCODE statement transfers one record of length *n* to the internal file *dest*. If format *f* attempts to write a second record, ENCODE repositions the current record position to the beginning of the internal file and begins writing at that position.

An error is issued when the ENCODE statement attempts to write more than *n* characters to the record of the internal file. If *dest* is a noncharacter entity and *n* is not a multiple of 8 (for UNICOS and UNICOS/mk systems) or 4 (for IRIX systems), the last word of the record is padded with blanks to a word boundary. If *dest* is a character entity, the last word of the record is not padded with blanks to a word boundary.

Example 1: The following example assumes a machine word length of 64 bits and uses the underscore character (`_`) as a blank:

```
      INTEGER ZD(5), ZE(3)
      ZD(1) = 'THIS ____'
      ZD(2) = 'MUST ____'
      ZD(3) = 'HAVE ____'
      ZD(4) = 'FOUR ____'
      ZD(5) = 'CHAR ____'
1     FORMAT(5A4)
      ENCODE(20,1,ZE)ZD
      DO 10 I=1,3
          PRINT 2,'ZE(' ,I,')="' ,ZE(I),'" '
10    CONTINUE
2     FORMAT(A,I2,A,A8,A)
      END
```


On UNICOS systems, the output is as follows:

```
>ZE( 1) = "THISMUST"
>ZE( 2) = "HAVEFOUR"
>ZE( 3) = "CHAR____"
```

Example 2: On IRIX systems, the comparable example would be as follows:

```
INTEGER ZD(5), ZE(3)
ZD(1) = 'TH__'
ZD(2) = 'IS__'
ZD(3) = '=4__'
ZD(4) = 'CH__'
ZD(5) = 'AR__'
1  FORMAT(5A2)
   ENCODE(10,1,ZE)ZD
   DO 10 I=1,3
   PRINT 2,'ZE(',I,')="'',ZE(I),'"'
10  CONTINUE
2  FORMAT(A,I2,A,A4,A)
   END
```

The output is as follows:

```
>ZE( 1) = "THIS"
>ZE( 2) = "=4CH"
>ZE( 3) = "AR__"
```

6.2.2 DECODE Statement

The DECODE statement provides a method of converting or decoding from a character representation to the internal representation of the entities in the input list. The format of the DECODE statement is as follows:

<pre>DECODE (<i>n</i>, <i>f</i>, <i>source</i>) [<i>dlist</i>]</pre>
--

n Number of characters to be processed. Nonzero integer expression not to exceed the maximum record length for formatted records. This is the record size for the internal file.

f Format identifier. It cannot be an asterisk.

source Name of internal file. It can be a variable or array of any data type. It cannot be an array section or a zero-sized array or a zero-sized character variable.

dlist Input list to be converted from character during the DECODE statement.

The input list items are converted using format *f* from a sequence of *n* characters in the internal file *source* to an internal representation and stored in the input list entities. If the internal file *source* is noncharacter, the internal file is assumed to be a multiple of 8 characters (for UNICOS and UNICOS/mk systems) or 4 characters (for IRIX systems).

Example 1: On UNICOS systems, an example of a DECODE statement is as follows:

```

      INTEGER ZD(4), ZE(3)
      ZE(1)='WHILETHI'
      ZE(2)='S HAS F'
      ZE(3)='IVE '
3     FORMAT(4A5)
      DECODE(20,3,ZE)ZD
      DO 10 I=1,4
          PRINT 2,'ZD(',I,')="'',ZD(I),'"'
10    CONTINUE
2     FORMAT(A,I2,A,A8,A)
      END

```

The output is as follows:

```

>ZD( 1)="WHILE  "
>ZD( 2)="THIS   "
>ZD( 3)="HAS    "
>ZD( 4)="FIVE   "

```

Example 2: On IRIX systems, an example of a DECODE statement is as follows:

```

      INTEGER ZD(5), ZE(4)
      ZE(1)='WHIL'
      ZE(2)='E_IT'
      ZE(3)='=4CH'
      ZE(4)='ARS_'
      ZE(5)='RS.+ '
3     FORMAT(5A3)
      DECODE(16,3,ZE)ZD
      DO 10 I=1,4

```

```

          PRINT 2, 'ZD( ', I, ' ) = " ', ZD(I), ' "'
10      CONTINUE
2       FORMAT(A, I2, A, A4, A)
      END

```

The output is as follows:

```

>ZD( 1) = "WHI_"
>ZD( 2) = "LE__"
>ZD( 3) = "IT=_ "
>ZD( 4) = "4CH_"
>ZD( 5) = "ARS_"

```

6.3 Edit Descriptors

The following sections show obsolete edit descriptors and outmoded uses of current descriptors.

6.3.1 Asterisk Delimiters

The asterisk was allowed to delimit a literal character constant. It has been replaced by the apostrophe and quotation mark.

$*h_1 h_2 \dots h_n*$

*** Delimiter for a literal character string

h Any ASCII character indicated by a C in Table 3, page 77 (that is, capable of internal representation)

Example:

```
*AN ASTERISK EDIT DESCRIPTOR*
```

6.3.2 Negative-valued x Descriptor

A negative value could be used with the X descriptor to indicate a move to the left. This has been replaced by the TL descriptor.

$[-b]x$

b Any nonzero, unsigned integer constant
x Indicates a move of as many positions as indicated by *b*

Example:

```
-55X ! Moves current position 55 spaces left
```

6.3.3 A and R Descriptors for Noncharacter Types

The *Rw* descriptor and the use of the *Aw* descriptor for noncharacter data are available primarily for programs that were written before a true character type was available. Other uses include adding labels to binary files and the transfer of data whose type is not known in advance.

List items can be of type real, integer, complex, or logical. For character use, the binary form of the data is converted to or from ASCII codes. The numeric list item is assumed to contain ASCII characters when used with these edit descriptors.

Complex items use two storage units and require two A descriptors, for the first and second storage units respectively.

The *Aw* descriptor works with noncharacter list items containing character data in essentially the same way as described in the *Fortran Language Reference Manual, Volume 1*. The *Rw* descriptor works like *Aw* with the following exceptions:

- Characters in an incompletely filled input list item are right-justified with the remainder of that list item containing binary zeros.
- Partial output of an output list item is from its rightmost character positions.

The following example shows the *Aw* and *Rw* edit descriptors for noncharacter data types:

```
INTEGER IA
LOGICAL LA
REAL RA
DOUBLE PRECISION DA
COMPLEX CA
CHARACTER*52 CHC
CHC='ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz'
READ(CHC,3) IA, LA, RA, DA, CA
3 FORMAT(A4,A8,A10,A17,A7,A6)
PRINT 4, IA, LA, RA, DA, CA
```

```

4      FORMAT(1x,3(A8,'-'),A16,'-',2A8)
      READ(CHC,5) IA, LA, RA
5      FORMAT(R2,R8,R9)
      PRINT 4, IA, LA, RA
      END

```

On UNICOS and UNICOS/mk systems, the output of this program would be as follows:

```

> ABCD      -EFGHIJKL-OPQRSTUVWXYZ-xyzabcdefghijklmnopqrst uvwxyz
      ^^^^^
> ooooooAB-CDEFGHIJ-LMNOPQRS-

```

The carat (^) indicates leading blanks in the use of the A edit descriptor. The lowercase letter o is used to indicate where binary zeros have been written with the R edit descriptor.

On IRIX systems, the output of this program would be as follows:

```

>      ABCD-      IJKL-      STUV-      fghijklm-      qrst      wxyz
      ^^^^^      ^^^^^      ^^^^^      ^^^^^^^^^      ^^^^^      ^^^^^
>      AB-      GHIJ-      PQRS-
      ^^^^^      ^^^^^      ^^^^^

```

The binary zeros are not printable characters, so the printed output simply contains the characters without the binary zeros.

6.4 Type Declaration with Data Length

Data type declarations that include the data length are outmoded. The CF90 and MIPSpro 7 Fortran 90 compilers recognize this usage in type statements, IMPLICIT statements, and FUNCTION statements, mapping these numbers onto lengths appropriate for the target machine.

Format:

```

type [ *n ] v [ , v ] ...

IMPLICIT type [ *n ] ( a1 [ -a2 ] [ , a1 [ -a2 ] ] ... )
      [ , type ... ] ...

[ type [ *n ] ] FUNCTION fun ([ d [ , d ] ... ] )

```

<i>type</i>	Can be INTEGER, REAL, DOUBLE PRECISION, COMPLEX, or LOGICAL.
<i>*n</i>	Data length as shown in Table 5, Table 6, page 149, and Table 7, page 149. Any other data length generates an error.
<i>v</i>	Name of a constant, variable, or array declarator.
<i>a_n</i>	A letter. A range of letters is denoted by the first and last letters of the range separated by a hyphen. A range (<i>a₁ – a_n</i>) has the same effect as a list of the letters (<i>a₁, a₂, ... a_n</i>).
<i>fun</i>	Name of the function subprogram.
<i>d</i>	Dummy argument representing a variable, array, or dummy procedure name.

The following tables show the data lengths for UNICOS, UNICOS/mk, and IRIX systems.

Note: On UNICOS systems, a 32-bit item or a 46-bit item is contained in a 64-bit word.

Table 5. Data length (UNICOS systems)

<i>type</i>	<i>n:</i>	*1	*2	*4	*8	*16	*32
INTEGER		64-bit	64-bit	64-bit	64-bit	Error	Error
REAL		Error	Error	64-bit single precision	64-bit single precision	128-bit double precision	Error
COMPLEX		Error	Error	Error	128-bit single precision	128-bit single precision	256-bit double precision
LOGICAL		64-bit	64-bit	64-bit	64-bit	Error	Error
DOUBLE PRECISION		Error	Error	Error	Error	128-bit double precision	Error

Table 6. Data length (UNICOS/mk systems)

<i>type</i>	<i>n:</i>	*1	*2	*4	*8	*16	*32
INTEGER		32-bit	32-bit	32-bit	64-bit	Error	Error
REAL		Error	Error	32-bit single precision ¹	64-bit double precision ²	64-bit double precision ³	Error
COMPLEX		Error	Error	Error	64-bit single precision ⁴	64-bit single precision ⁵	64-bit single precision ⁶
LOGICAL		32-bit	32-bit	32-bit	64-bit	Error	Error
DOUBLE PRECISION		Error	Error	Error	Error	64-bit single precision ⁷	Error

Table 7. Data length (IRIX systems)

<i>type</i>	<i>n:</i>	*1	*2	*4	*8	*16	*32
INTEGER		8-bit	16-bit	32-bit	64-bit	Error	Error
LOGICAL		8-bit	16-bit	32-bit	64-bit	Error	Error
REAL		Error	Error	32-bit	64-bit	128-bit	Error
COMPLEX		Error	Error	32-bit	64-bit	128-bit	Error
DOUBLE PRECISION		Error	Error	Error	64-bit	Error	Error

¹ This is an additional precision on a UNICOS/mk system.

² This is a single precision on a UNICOS/mk system.

³ 128-bit precision is not supported on UNICOS/mk systems.

⁴ This is an additional precision on a UNICOS/mk system.

⁵ 128-bit precision is not supported on UNICOS/mk systems.

⁶ 128-bit precision is not supported on UNICOS/mk systems.

⁷ 128-bit precision is not supported on UNICOS/mk systems.

6.5 DATA Statement Features

The DATA statement has the following outmoded features:

- A constant need not exist for each element of a whole array named in a *data_stmt_object_list* if the array is the last item in the list.
- A Hollerith or character constant can initialize more than one element of an integer or single-precision real array if the array is specified without subscripts.

Example 1: On a machine with 64-bit words, if an array is declared by `INTEGER A(2)`, the following DATA statements have the same effect:

```
DATA A /'1234567890123456'/  
DATA A /'12345678','90123456'/
```

Example 2: On a machine with 32-bit words, if an array is declared by `INTEGER A(2)`, the following DATA statements have the same effect:

```
DATA A /'12345678'/  
DATA A /'1234','5678'/
```

An integer or single-precision real array can be defined in the same way in a DATA implied-DO statement.

6.6 IF Statements

Outmoded IF statements are the two-branch arithmetic IF and the indirect logical IF.

6.6.1 Two-branch Arithmetic IF

A two-branch arithmetic IF statement transfers control to statement s_1 if expression e is evaluated as nonzero or to statement s_2 if e is zero. The arithmetic expression should be replaced with a relational expression, and the statement should be changed to an IF statement or an IF construct. This format is as follows:

<code>IF (e) s₁, s₂</code>
--

e Integer, real, or double-precision expression

s Label of an executable statement in the same program unit

Example:

```
IF ( I+J*K ) 100,101
```

6.6.2 Indirect Logical IF

An indirect logical IF statement transfers control to statement s_t if logical expression le is true and to statement s_f if le is false. An IF construct or an IF statement should be used in place of this outmoded statement. This format is as follows:

$\text{IF (} le \text{) } s_t \text{ } s_f$

le Logical expression

$s_t \text{ } s_f$ Labels of executable statements in the same program unit

Example:

```
IF ( X . GE . Y ) 148 , 9999
```

6.7 TASK COMMON Statement (UNICOS Systems Only)

When multitasking is used, some common blocks might need to be local to a task. The TASK COMMON statement declares all variables in a common block to be local to a task. If multiple tasks execute code containing the same task common block, each task will have a separate copy of the block. A common block cannot be declared both local common and task common. If a common block is declared local common in one routine and task common in another routine, the loader will generate an error.

A task common block can also be declared by the use of a COMMON statement with the TASKCOMMON compiler directive. The compiler directives are described in *CF90 Commands and Directives Reference Manual*. The directive is recommended over the TASK COMMON statement for better portability.

The keyword TASK must precede the keyword COMMON to establish a task common block. Task common blocks must be named. A task common block is allocated at task invocation.

The TASK COMMON statement has the following format:

TASK COMMON / <i>cb</i> / <i>member_list</i> [, / <i>cb</i> / <i>member_list</i>] ...

<i>cb</i>	Task common block name.
<i>member_list</i>	A variable name, array name, or array declarator. A member name must not be a subprogram dummy argument name.

Variables in *member_list* may appear in a DATA statement.

For information on using the `-a alloc` option to allocate storage from the f90(1) command line, see the f90(1) man page or the *CF90 Commands and Directives Reference Manual*.

6.8 Nested Loop Termination

Older Cray Research Fortran compilers allowed nested DO loops to terminate on a single END DO statement if the END DO statement had a statement label. The END DO statement is included in the Fortran standard. The Fortran standard specifies that a separate END DO statement must be used to terminate each DO loop, so allowing nested DO loops to end on a single, labeled END DO statement is an outmoded feature.

6.9 DOUBLE COMPLEX Statement (UNICOS Systems Only)

The DOUBLE COMPLEX statement is used to declare an item to be of type double complex. The format for the DOUBLE COMPLEX statement is as follows:

DOUBLE COMPLEX [, <i>attribute_list</i> ::] <i>entity_list</i>
--

Items declared as DOUBLE COMPLEX contain two double-precision entities.

When the `-d p` option is in effect, double complex entities are affected as follows:

- The nonstandard DOUBLE COMPLEX declaration is treated as a single-precision complex type.
- Double-precision intrinsic procedures are changed to the corresponding single-precision intrinsic procedures.

The `-e p` or `-d p` specification is used for all source files compiled with a single invocation of the `f90(1)` command. If a module is compiled separately from a program unit that uses the module, they both must be compiled with the same `-e p` or `-d p` specification.

Table 8 shows the CF90 double complex intrinsic functions and the preferred standard alternatives:

Table 8. Standard alternatives to CF90 double-complex functions

Double complex function	Fortran 90 standard alternative
CDABS	ABS(3)
CDCOS	COS(3)
CDEXP	EXP(3)
CDLOG	LOG(3)
CDSIN	SIN(3)
CDSQRT	SQRT(3)

6.10 Bitwise Logical Expressions

A *bitwise logical expression* (also called a masking expression) is an expression in which a logical operator operates on individual bits within integer, real, Cray pointer, or Boolean operands, giving a result of type Boolean. Each operand is treated as a single storage unit. This storage unit is a 64-bit word on UNICOS and UNICOS/mk systems; it is a 32-bit word on IRIX systems. The result is a single storage unit. Boolean values and bitwise logical expressions are contrasted to logical values and expressions.

Bitwise logical operators can also be written as functions; for example `A .AND. B` can be written as `AND(A, B)` and `.NOT. A` can be written as `COMPL(A)`.

The CF90 and MIPSpro 7 Fortran 90 compiler intrinsic functions that operate on Boolean values in bitwise fashion, such as shifting, parity count, and tallying 1s or leading 0s, are extensions to the Fortran standard. Generally, these bitwise functions have equivalent Fortran standard intrinsic procedures. Table 9 shows the bitwise functions and, where possible, their equivalent Fortran standard intrinsic procedures:

Table 9. Standard alternatives to CF90 and MIPSpro 7 Fortran 90 bitwise functions

Bitwise function	Fortran standard alternative
AND(3M)	IAND(3I)
COMPL(3I)	NOT(3I)
CSMG(3I)	MERGE(3I)
EQV(3M)	IEOR(3I)
MASK(3I)	IBSET(3I)
OR(3M)	IOR(3I)
NEQV(3M)	IEOR(3I)
SHIFT(3I)	ISHFT(3I), ISHFTC(3I)
SHIFTL(3I), LSHIFT(3I)	ISHFT(3I), ISHFTC(3I)
SHIFTR(3I), , RSHIFT(3I)	ISHFT(3I), ISHFTC(3I)
XOR(3M)	IEOR(3I)

If one operand is of type logical, then both operands must be of type logical; the operation performed, then, is a logical operation (not a masking operation). In a logical or masking operation, neither operand can be of type double precision or of type double complex.

Table 10, page 155, shows which data types can be used together in bitwise logical operations.

Table 10. Data types in bitwise logical operations

x_1, x_2	Integer	Real	Boolean	Pointer	Logical	Character
Integer	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Not valid	Not valid ¹
Real	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Not valid	Not valid ¹
Boolean	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Not valid	Not valid ¹
Pointer	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Masking operation, Boolean result.	Not valid	Not valid ¹
Logical	Not valid	Not valid	Not valid	Not valid	Logical operation logical result	Not valid
Character	Not valid ¹	Not valid ¹	Not valid ¹	Not valid ¹	Not valid	Not valid

Notes:

1. x_1 and x_2 represent operands for a logical or bitwise expression, using operators `.NOT.`, `.AND.`, `.OR.`, `.XOR.`, `.NEQV.`, and `.EQV.`
2. The entry “Not valid¹” indicates that if the operand is a character operand of 32 or fewer characters, the operand is treated as a Hollerith constant and is allowed.

Bitwise logical expressions can be combined with expressions of Boolean or other types by using arithmetic, relational, and logical operators. Evaluation of an arithmetic or relational operator processes a bitwise logical expression with no type conversion. Boolean data is never automatically converted to another type.

A bitwise logical expression performs the indicated logical operation separately on each bit. The interpretation of individual bits in bitwise *multiplication_exprs*,

summation_exprs, and general expressions is the same as for logical expressions. The results of binary 1 and 0 correspond to the logical results TRUE and FALSE, respectively, in each of the bit positions. These values are summarized as follows:

.NOT. 1100	1100	1100	1100	1100
=0011	.AND. 1010	.OR. 1010	.XOR. 1010	.EQV. 1010
	----	----	----	----
	1000	1110	0110	1001

CF90 Defined Externals [7]

This chapter describes global variables used by the CF90 compiler on UNICOS and UNICOS/mk systems

7.1 Conformance Checks

Additional `segldr(1)` and `clld(1)` directives for load time optimization and activating library features are described in the *Application Programmer's I/O Guide*.

Several `segldr(1)` directives are used to provide strict, intermediate, and minimal error checking of edit descriptors with input/output (I/O) list items during formatted READ and WRITE statements. The NOCHK versions provide the least error checking.

The version of NOCHK for formatted output is as follows:

```
% segldr -D EQUIV=$WNOCHK($WCHK)
```

The version of NOCHK for a formatted input is as follows:

```
% segldr -D EQUIV=$RNOCHK($RCHK)
```

For strict conformance to editing in FORTRAN 77, use the CHK77 versions, which are as follows:

```
% segldr -D EQUIV=$WCHK77($WCHK)
```

```
% segldr -D EQUIV=$RCHK77($RCHK)
```

For strict conformance to editing in Fortran 90, use the CHK90 versions, which are as follows:

```
% segldr -D EQUIV=$WCHK90($WCHK)
```

```
% segldr -D EQUIV=$RCHK90($RCHK)
```

The default checking is somewhat stricter than the NOCHK versions but is not as strict as the CHK77 and CHK90 versions.

7.2 Record Length

The RECL specifier in an OPEN statement can be used to specify the maximum record size for a file declared with sequential access. An alternate method is

provided through `segldr(1)` directives. If `RECL` is present, the values provided by these directives are ignored. The use of `RECL` for sequential access files is recommended.

To set the maximum output record length `x` for a file opened as a sequential formatted file, use the following specification:

```
SET=$WBUFLN:X  
COMMONS=$WFDCOM:X+9
```

The default size is 267.

To set the maximum input record length `x` for a file opened as a sequential formatted file, use the following specification:

```
SET=$RBUFLN:X  
COMMONS=$RFDCOM:X+9
```

The default size is 267.

argument keyword

The name of a dummy (or formal) argument. This name is used in the subprogram definition; it also may be used when the subprogram is invoked to associate an actual argument with a dummy argument. Using argument keywords allows the actual arguments to appear in any order. The Fortran 90 standard specifies argument keywords for all intrinsic procedures. Argument keywords for user-supplied external procedures may be specified in a procedure interface block.

array

(1) A data structure that contains a series of related data items arranged in rows and columns for convenient access. The C shell and the `awk(1)` command can store and process arrays. (2) In Fortran 90, an object with the `DIMENSION` attribute. It is a set of scalar data, all of the same type and type parameters. The rank of an array is at least 1, and at most 7. Arrays may be used as expression operands, procedure arguments, and function results, and they may appear in input/output (I/O) lists.

association

An association permits an entity to be referenced by different names in a scoping unit or by the same or different names in different scoping units. Several kinds of association exist. The principal kinds of association are pointer association, argument association, host association, use association, and storage association.

automatic variable

A variable that is not a dummy argument but whose declaration depends on a nonconstant expression (array bounds and/or character length).

Autotasking

A trademarked process of Cray Research that automatically divides a program into individual tasks and organizes them to make the most efficient use of the computer hardware.

bottom loading

An optimization technique used on some scalar loops in which operands are prefetched during each loop iteration for use in the next iteration. The operand is available as soon as the first loop instruction executes. A prefetch is performed even during the final loop iteration, before the loop's final jump test has been performed.

cache

In a processing unit, a high-speed buffer storage that is continually updated to contain recently accessed contents of main storage. Its purpose is to reduce access time. In disk subsystems, a method the channel buffers use to buffer disk data during transfer between the devices and memory.

cache line

On Cray MPP systems, a cache line consists of four quad words, which is the maximum size of a hardware message.

CIV

A constant increment variable is a variable that is incremented only by a loop invariant value (for example, in a loop with index J , the statement $J = J + K$, in which K can be equal to 0, J is a CIV).

constant

A data object whose value cannot be changed. A named entity with the `PARAMETER` attribute is called a named constant. A constant without a name is called a literal constant.

construct

A sequence of statements that starts with a `SELECT CASE`, `DO`, `IF`, or `WHERE` statement and ends with the corresponding terminal statement.

control construct

An action statement that can change the normal execution sequence (such as a `GO TO`, `STOP`, or `RETURN` statement) or a `CASE`, `DO`, or `IF` construct.

critical region

On Cray MPP systems, a synchronization mechanism that enforces serial access to a piece of code. Only one PE may execute in a critical region at a time.

data entity

A data object, the result of the evaluation of an expression, or the result of the execution of a function reference (also called the function result). A data entity always has a type.

data object

A constant, a variable, or a part of a constant or variable.

declaration

A nonexecutable statement that specifies the attributes of a data object (for example, it may be used to specify the type of a variable or function result or the shape of an array).

definition

This term is used in two ways. (1) A data object is said to be defined when it has a valid or predictable value; otherwise, it is undefined. It may be given a valid value by execution of statements such as assignment or input. Under certain circumstances, it may subsequently become undefined. (2) Procedures and derived types are said to be defined when their descriptions have been supplied by the programmer and are available in a program unit.

derived type

A type that is not intrinsic (a user-defined type); it requires a type definition to name the type and specify its components. The components may be of intrinsic or user-defined types. An object of derived type is called a structure. For each derived type, a structure constructor is available to specify values. Operations on objects of derived type must be defined by a function with an interface and the generic specifier `OPERATOR`. Assignment for derived type objects is defined intrinsically, but it may be redefined by a subroutine with the `ASSIGNMENT` generic specifier. Data objects of derived type may be used as procedure arguments and function results, and they may appear in input/output (I/O) lists.

designator

Sometimes it is convenient to reference only part of an object, such as an element or section of an array, a substring of a character string, or a component of a structure. This requires the use of the name of the object followed by a selector that selects a part of the object. A name followed by a selector is called a **designator**.

entity

(1) In Open Systems Interconnection (OSI) terminology, a layered protocol machine. An entity in a layer performs the functions of the layer in one computer system, accessing the layer entity below and providing services to the layer entity above at local service access points. (2) In Fortran 90, a general term used to refer to any Fortran 90 concept (for example, a program unit, a common block, a variable, an expression value, a constant, a statement label, a construct, an operator, an interface block, a derived type, an input/output (I/O) unit, a name list group, and so on).

executable construct

A statement (such as a GO TO statement) or a construct (such as a DO or CASE construct).

expression

A set of operands, which may be function invocations, and operators that produce a value.

extent

A structure that defines a starting block and number of blocks for an element of file data.

function

Usually a type of operating-system-related function written outside a program and called in to do a specific function. Smaller and more limited in capability than a utility. In a programming language, a function is usually defined as a closed subroutine that performs some defined task and returns with an answer, or identifiable return value.

The word "function" has a more specific meaning in Fortran than it has in C. In C, it refers to any called code; in Fortran, it refers to a subprogram that returns a value.

generic specifier

An optional component of the `INTERFACE` statement. It can take the form of an identifier, an `OPERATOR (defined_operator)` clause, or an `ASSIGNMENT (=)` clause.

heap

A section of memory within the user job area that provides a capability for dynamic allocation. See the `HEAP` directive in SR-0066.

inlining

The process of replacing a user subroutine or function call with the definition itself. This saves subprogram call overhead and may allow better optimization of the inlined code. If all calls within a loop are inlined, the loop becomes a candidate for vectorization and/or tasking.

intrinsic

Anything that the language defines is intrinsic. There are intrinsic data types, procedures, and operators. You may use these freely in any scoping unit. Fortran programmers may define types, procedures, and operators; these entities are not intrinsic.

local

(1) A type of scope in which variables are accessible only to a particular part of a program (usually one module). (2) The system initiating the request for service. This term is relative to the perspective of the user.

multitasking

(1) The parallel execution of two or more parts of a program on different CPUs; these parts share an area of memory. (2) A method in multiuser systems that incorporates multiple interconnected CPUs; these CPUs run their programs simultaneously (in parallel) and shares resources such as memory, storage devices, and printers. This term can often be used interchangeably with **parallel processing**.

name

A term that identifies many different entities of a program such as a program unit, a variable, a common block, a construct, a formal argument of a

subprogram (dummy argument), or a user-defined type (derived type). A name may be associated with a specific constant (named constant).

operator

(1) A symbolic expression that indicates the action to be performed in an expression; operator types include arithmetic, relational, and logical. (2) In Fortran 90, an operator indicates a computation that involves one or two operands. Fortran 90 defines several intrinsic operators (for example, +, -, *, /, ** are numeric operators, and .NOT., .AND., and .OR. are logical operators). Users also may define operators for use with operands of intrinsic or derived types.

overindexing

The nonstandard practice of referencing an array with a subscript not contained between the declared lower and upper bounds of the corresponding dimension for that array. This practice sometimes, but not necessarily, leads to referencing a storage location outside of the entire array.

parallel processing

Processing in which multiple processors work on a single application simultaneously.

pointer

(1) A data item that consists of the address of a desired item. (2) A symbol that moves around a computer screen under the control of the user.

procedure

(1) A named sequence of control statements and/or data that is saved in a library for processing at a later time, when a calling statement activates it; it provides the capability to replace values within the procedure. (2) In Fortran 90, procedure is defined by a sequence of statements that expresses a computation that may be invoked as a subroutine or function during program execution. It may be an intrinsic procedure, an external procedure, an internal procedure, a module procedure, a dummy procedure, or a statement function. If a subprogram contains an ENTRY statement, it defines more than one procedure.

procedure interface

In Fortran 90, a sequence of statements that specifies the name and characteristics of one or more procedures, the name and attributes of each

dummy argument, and the generic specifier by which it may be referenced if any. See **generic specifier**.

In FORTRAN 77 and Fortran 90, a **generic function** is one whose output value data type is determined by the data type of its input arguments. In FORTRAN 77, the only generic functions allowed are those that the standard defines. In Fortran 90, programmers may construct their own generic function by creating "generic interface," which is like a regular procedure interface, except that it has a "generic specifier" (the name of the generic function) after the keyword `INTERFACE`.

reduction loop

A loop that contains at least one statement that reduces an array to a scalar value by doing a cumulative operation on many of the array elements. This involves including the result of the previous iteration in the expression of the current iteration.

reference

A data object reference is the appearance of a name, designator, or associated pointer in an executable statement that requires the value of the object. A procedure reference is the appearance of the procedure name, operator symbol, or assignment symbol in an executable program that requires execution of the procedure. A module reference is the appearance of the module name in a `USE` statement.

scalar

(1) In Fortran 90, a single object of any intrinsic or derived type. A structure is scalar even if it has a component that is an array. The rank of a scalar is 0. (2) A nonvectorized, single numerical value that represents one aspect of a physical quantity and may be represented on a scale as a point. This term often refers to a floating-point or integer computation that is not vectorized; more generally, it also refers to logical and conditional (jump) computation.

scope

The region of a program in which a variable is defined and can be referenced.

scoping unit

Part of a program in which a name has a fixed meaning. A program unit or subprogram generally defines a scoping unit. Type definitions and procedure

interface bodies also constitute scoping units. Scoping units do not overlap, although one scoping unit may contain another in the sense that it surrounds it. If a scoping unit contains another scoping unit, the outer scoping unit is referred to as the host scoping unit of the inner scoping unit.

search loop

A loop that can be exited by means of an IF statement.

sequence

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

shared

Accessible by multiple parts of a program. Shared is a type of scope.

shell variable

A name representing a string value. Variables that are usually set only on a command line are called **parameters** (positional parameters and keyword parameters). Other variables are simply names to which a user (user-defined variables) or the shell itself may assign string values. The shell has predefined shell variables (for example, HOME). Variables are referenced by prefixing the variable name by a \$ (for example, \$HOME).

software pipelining

Software pipelining is a compiler code generation technique in which operations from various loop iterations are overlapped in order to exploit instruction-level parallelism, increase instruction issue rate, and better hide memory and instruction latency. As an optimization technique, software pipelining is similar to bottom loading, but it includes additional, and more efficient, scheduling optimizations.

Cray compilers perform safe bottom loading by default. Under these conditions, code generated for a loop contains operations and stores associated with the present loop iteration and contains loads associated with the next loop iteration. Loads for the first iteration are generated in the loop preamble.

When software pipelining is performed, code generated for the loop contains loads, operations, and stores associated with various iterations of the loop. Loads and operations for first iterations are generated in the preamble to the

loop. Operations and stores for last iterations of loop are generated in the postamble to the loop.

statement keyword

A keyword that is part of the syntax of a statement. Each statement, other than an assignment statement and a statement function definition, begins with a statement keyword. Examples of these keywords are `IF`, `READ`, and `INTEGER`. Statement keywords are not reserved words; you may use them as names to identify program elements.

stripmining

A single-processor optimization technique in which arrays, and the program loops that reference them, are split into optimally-sized blocks, termed strips. The original loop is transformed into two nested loops. The inner loop references all data elements within a single strip, and the outer loop selects the strip to be addressed in the inner loop. This technique is often performed by the compiler to maximize the usage of cache memory or as part of vector code generation.

structure

A language construct that declares a collection of one or more variables grouped together under one name for convenient handling. In C and C++, a structure is defined with the `struct` keyword. In Fortran 90, a derived type is defined first and various structures of that type are subsequently declared.

subobject

Parts of a data object may be referenced and defined separately from other parts of the object. Portions of arrays are array elements and array sections. Portions of character strings are substrings. Portions of structures are structure components. Subobjects are referenced by designators and are considered to be data objects themselves.

subroutine

A series of instructions that accomplishes a specific task for many other routines. (A subsection of a user-written program of varying size and, therefore, function. It is written within the program. It is not a subsection of a routine.) It differs from a main routine in that one of its parameters must specify the location to which to return in the main program after the function has been accomplished.

TKR

An acronym that represents attributes for argument association. It represents the data type, kind type parameter, and rank of the argument.

type parameter

Two type parameters exist for intrinsic types: kind and length. The kind type parameter `KIND` indicates the decimal range for the integer type, the decimal precision and exponent range for the real and complex types, and the machine representation method for the character and logical types. The length type parameter `LEN` indicates the length of a character string.

variable

(1) A name that represents a string value. Variables that usually are set only on a command line are called parameters. Other variables are simply names to which the user or the shell may assign string values. (2) In Fortran 90, data object whose value can be defined and redefined. A variable may be a scalar or an array. (3) In the shell command language, a named parameter. See also **shell variable**.

A

ASCII character set, 77
ASSIGN statement, 70

B

Backus-Naur Form, 1
Bitwise logical expressions, 153
BNF syntax summary, 1

C

Character
 Hollerith, 138
Character data representation
 CRAY T90 (IEEE) systems, 126
 IRIX systems, 112
 UNICOS systems, 101
 UNICOS/mk systems, 119
Character set, 77
Compiler differences, 83
Complex data representation
 CRAY T90 (IEEE) systems, 123
 IRIX systems, 109
 UNICOS/mk systems, 117
Complex type (single precision), internal
 representation, 109
Constraints, 5
Cray character pointer data representation
 CRAY T90 (IEEE) systems, 127
 IRIX systems, 112
 UNICOS systems, 102
 UNICOS/mk systems, 120

D

Data
 type
 Hollerith, 138
DATA statement, 150
DECODE statement, 143
Decremental features, 69
Defined externals, 157
Differences (between compilers), 83
DOUBLE COMPLEX statement, 152
Double-precision complex data representation
 UNICOS systems, 100
Double-precision data representation
 UNICOS systems, 99
Dynamic memory allocation, 132

E

Edit descriptors
 outmoded, 145
ENCODE statement, 141
Extensions, 83
Externals (defined), 157

F

Formatted
 I/O and internal files, 141
Fortran
 keywords, 1

G

Global variables, 157

H

Hollerith type, 138

I

I/O

formatted, 141

IF statement, 150

Integer data representation

CRAY T90 (IEEE) systems, 120

IRIX systems, 103

UNICOS systems, 95

UNICOS/mk systems, 113

IRIX system data representation, 103

K

Keywords, 1

L

Logical data representation

CRAY T90 (IEEE) systems, 126

IRIX systems, 112

UNICOS systems, 102

UNICOS/mk systems, 119

M

Memory allocation, 132

O

Obsolescent features, 69, 73

Outmoded features, 137

P

PAUSE statement, 69

R

Real data representation

CRAY T90 (IEEE) systems, 121

IRIX systems, 105

UNICOS systems, 96

UNICOS/mk systems, 114

Real type, internal representation, 105

S

Single-precision complex data representation

UNICOS systems, 99

Stack storage, 130

Static storage, 130

Storage, 127

Syntax summary (in BNF), 1

T

TASKCOMMON statement, 151

U

UNICOS data representation, 95