

CONVEX FORTRAN Language Reference Manual

Tenth Edition



CONVEX



606

Convex Computer Corporation
3000 Waterview Parkway
P.O. Box 833851
Richardson, TX 75083-3851
United States of America
(214)497-4000



CONVEX FORTRAN

Language Reference Manual



Order No. DSW-037

Tenth Edition
November 1992

CONVEX Press
Richardson, Texas
United States of America

CONVEX FORTRAN Language Reference Manual

Order No. DSW-037

Copyright © 1992 CONVEX Computer Corporation
All rights reserved.

This document is copyrighted. All rights reserved. This document may not, in whole or part, be copied, duplicated, reproduced, translated, electronically stored, or reduced to machine readable form without prior written consent from CONVEX Computer Corporation.

Although the material contained herein has been carefully reviewed, CONVEX does not warrant it to be free of errors or omissions. CONVEX reserves the right to make corrections, updates, revisions, or changes to the information contained herein. CONVEX does not warrant the material described herein to be free of patent infringement.

UNLESS PROVIDED OTHERWISE IN WRITING WITH CONVEX COMPUTER CORPORATION (CONVEX), THE PROGRAM DESCRIBED HEREIN IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES. THE ABOVE EXCLUSION MAY NOT BE APPLICABLE TO ALL PURCHASERS BECAUSE WARRANTY RIGHTS CAN VARY FROM STATE TO STATE. IN NO EVENT WILL CONVEX BE LIABLE TO ANYONE FOR SPECIAL, COLLATERAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING ANY LOST PROFITS OR LOST SAVINGS, ARISING OUT OF THE USE OR INABILITY TO USE THIS PROGRAM. CONVEX WILL NOT BE LIABLE EVEN IF IT HAS BEEN NOTIFIED OF THE POSSIBILITY OF SUCH DAMAGE BY THE PURCHASER OR ANY THIRD PARTY.

CONVEX, the CONVEX logo ("C") CXdb, CXpa, and CXmetrics are registered trademarks of CONVEX Computer Corporation.

CONVEX C100 Series, C200 Series, C3 Series, C3200 Series, C3400 Series, C3800 Series, C1, C120, C210, C220, C230, C240, C3200, C3400, C3420, C3440, C3460, C3480, C3820, C3840, C3860, C3880, ConvexOS and CXwindows are trademarks of CONVEX Computer Corporation.

COVUE is a registered trademark of CONVEX Computer Corporation. COVUE consists of the COVUEbatch, COVUEbinary, COVUEedt, COVUElib, COVUEnet, and COVUEshell products.

Cray is a registered trademark of Cray Research, Inc.

Sun FORTRAN is a trademark of Sun Microsystems, Inc.

UNIX is a trademark of UNIX System Laboratories, Inc.

VAX and VMS are trademarks of Digital Equipment Corporation.

Printed in the United States of America

Revision Information for

CONVEX FORTRAN Language Reference Manual

Edition	Document No.	Description
Tenth	720-002230-007	Released with CONVEX FORTRAN software V8.0, November 1992. Split <i>CONVEX FORTRAN Guide</i> back into its component books. Added masked array assignments, POINTERS, Cray word pointer arithmetic, EOSHIFT and CSHIFT intrinsics. Reorganized chapters and appendixes. Extensively revised content. Added Fortran 90 compatibility appendix. Enhanced and enlarged index.
Ninth	720-002230-004	Released with CONVEX FORTRAN V7.0.0.1, August, 1991.
Eighth Rev. 2	720-002230-003	Released with CONVEX FORTRAN V6.1.0.1, October, 1990.
Eighth Rev. 1	720-002230-001	Released with CONVEX FORTRAN V6.0, March, 1990.
Eighth	720-000050-204	Released with CONVEX FORTRAN V5.1, May 1989.
Seventh	720-000050-203	Released with CONVEX FORTRAN V5.0 November 1988.
Sixth Rev. 1	720-000050-202	Released with CONVEX FORTRAN V4.1, May 1988.
Sixth	720-000050-201	Released with CONVEX FORTRAN V4.0, November 1987.
Fifth	720-000050-200	Released with CONVEX FORTRAN V3.0, May 1987.
Fourth	720-000099-000	Released with CONVEX FORTRAN V2.2, September 1986.
Third	720-000150-000	Released with CONVEX FORTRAN V2.0, April 1986.
Second	720-000150-000	Released with CONVEX FORTRAN V1.7, September 1985.
First	720-000150-100	Released with CONVEX FORTRAN V1.0, February 1985. First release of the manual.



Contents

How to use this manual	xvii
Purpose and audience	xvii
Organization	xvii
Scope	xviii
Notational conventions	xviii
Notes and cautions	xix
Associated documents	xx
Technical assistance	xxi
The contact utility	xxi

1 Introduction	1
Types of programs	1
FORTRAN character set	1
Comment line	2
FORTRAN statements	2
Character-per-column formatting	4
Statement label field	4
Initial line	5
Continuation line	5
Statement text field	5
Debug statements	5
Compiler directives	6
ANSI-standard formatting	6
Tab-key formatting	6
Order of statements and lines	6
Symbolic names	7
OPTIONS statement	7
INCLUDE statement	8

2 Data types, constants, and variables	9
Data types	9
Conversion of data types	12
Constants	13
Integer constants	13
Real constants	13
Complex constants	14
Octal constants	15
Hexadecimal constants	16
Hollerith constants	17

Logical constants	18
Character constants	19
Variables	19

3 Arrays and substrings 21

Arrays	21
Array declaration	22
Allocatable array declarations	25
Fortran 90 allocatable arrays	26
Referencing array elements	27
Array storage	28
Character substrings	28

4 Expressions 31

Arithmetic expressions	31
Operator precedence	32
Data type priority	32
Relational expressions	33
Logical expressions	34
Character expressions	35

5 Specification statements 37

PROGRAM statement	38
COMMON statement	38
IMPLICIT statement	39
PARAMETER statement	41
Standard PARAMETER statement	41
Alternate PARAMETER statement	42
Type-declaration statements	43
Numeric type-declaration statements	43
CHARACTER type-declaration statements	44
RECORD type-declaration statements	45
POINTER statement	45
The LOC function	46
Dynamic memory allocation	47
DIMENSION statement	49
ALLOCATABLE statement	50
EQUIVALENCE statement	51
Equivalencing arrays	51
Equivalencing substrings	53
Using EQUIVALENCE in common blocks	53
NAMELIST statement	54
EXTERNAL statement	55
INTRINSIC statement	56
SAVE statement	56

6	DATA statement	59
	DATA statement form	59
	Implied-DO	61
	DATA statement extensions	62
7	Assignment statements	65
	Character conversions	66
	Fortran 90 array assignments	66
	Array-valued expressions	67
	Masked array assignment	67
	Data conversion rules	70
	ASSIGN statement	73
8	Control statements	75
	GOTO statements	75
	Unconditional GOTO statement	75
	Computed GOTO statement	76
	Assigned GOTO statement	77
	IF statements	78
	Arithmetic IF statement	78
	Logical IF statement	79
	Block IF statement	80
	Nested block IF statements	83
	Short-circuit evaluation of conditionals	83
	DO statement	84
	Nested DO loops	86
	Extended-range DO loops	86
	DO WHILE statement	87
	END DO statement	88
	CONTINUE statement	89
	CALL statement	89
	RETURN statement	89
	STOP statement	90
	PAUSE statement	90
	END statement	91
9	Input/output statements	93
	Records	94
	Formatted records	94
	Unformatted records	95
	ENDFILE record	95
	Files	95
	Internal files	95
	Units	96

Accessing files	97
Sequential access	97
Direct access	98
I/O statement format	98
Input/output lists	99
Implied-DO lists	99
Specifiers	100
Unit specifier	100
Format specifier	101
Record specifier	101
Status specifier	102
Error specifier	102
End-of-file specifier	103
Namelist specifier	103
READ statement	104
External sequential-access READ statements	105
Formatted	105
Unformatted	105
List-directed	106
Namelist-directed	106
External direct-access READ statements	107
Formatted	107
Unformatted	107
Internal READ statements	108
Sequential access	108
Direct-access	108
ACCEPT statement	109
WRITE statement	110
Sequential-access WRITE statements	111
Formatted	111
Unformatted	111
List-directed	112
Namelist-directed	112
External direct-access WRITE statements	112
Formatted	112
Unformatted	113
Internal WRITE statements	113
Sequential access	113
Direct access	114
PRINT and TYPE statements	114
Special input/output statements	115
ENCODE statement	115
DECODE statement	117
FIND statement	118
Auxiliary input/output statements	119
OPEN statement	119
ACCESS keyword	122
ASSOCIATEVARIABLE keyword	122
BLANK keyword	123

BLOCKSIZE keyword	123
CARRIAGECONTROL keyword	124
DEFAULTFILE keyword	124
DISPOSE keyword	125
ERR keyword	126
FILE keyword	126
FORM keyword	126
IOSTAT keyword	127
MAXREC keyword	127
NOSPANBLOCKS keyword	128
READONLY keyword	128
RECL keyword	128
RECORDTYPE keyword	129
SHARED keyword	129
STATUS keyword	129
UNIT keyword	130
CLOSE statement	131
INQUIRE statement	132
File-positioning statements	135
REWIND statement	136
BACKSPACE statement	136
ENDFILE statement	137
Binary data file format conversions	137
When to use the conversion feature	138
-dfc option	139
Conversion using OPEN statement	139
Restrictions on conversions	140
Error handling using data format conversions	141
User-defined conversions	142
Sample conversion routine	143
User-supplied conversion routine names	144
Conversion using a shell variable	146

10 Format specifications 149

FORMAT statement	149
FORMAT control	151
Repeat count	152
Descriptors	153
A descriptor	153
Apostrophe (') descriptor	155
Asterisk (*) descriptor	155
H descriptor	156
L descriptor	156
I descriptor	157
O descriptor	158
Z descriptor	159
F descriptor	160
E and D descriptors	162

G descriptor	164
B descriptors	166
P descriptor	168
S descriptors	169
R descriptor	170
X descriptor	171
T descriptors	171
Dollar sign (\$) descriptor	173
Q descriptor	174
Colon (:) descriptor	174
Slash (/) descriptor	175
Default field descriptor values	175
Comma field separator on input data	176
Runtime formats	177
Variable formats	177
List-directed formatting	179
List-directed input	179
Character input	180
Nulls and slashes	180
Namelist-directed input formatting	181
List-directed output	184
Namelist-directed output formatting	184
Carriage-control characters	186

11 Subprograms 187

BLOCK DATA subprogram	187
Procedures	188
Dummy and actual arguments	188
Variables as dummy arguments	189
Arrays as dummy arguments	190
Character arguments	192
Procedures as dummy arguments	194
Alternate return arguments	194
Functions	195
Intrinsic functions	195
Built-in functions	196
Statement functions	197
Function subprograms	199
Subroutine subprograms	201
ENTRY statement	202
RETURN statement	204

A Intrinsic and commonly used library routines 207

Generic and specific intrinsics	207
Notes	219
Commonly used library routines	222

B FORTRAN 66 compatibility	225
Compiling FORTRAN 66 programs	225
EXTERNAL statement	226
DO loop minimum iteration count	226
OPEN statement keywords	227
BLANK keyword	227
STATUS keyword	228
X descriptor	228
Format code separators	228

C Fortran 90 compatibility	229
Array sections	229
Automatic arrays	231
Allocatable arrays	232
The ALLOCATABLE statement	232
The ALLOCATE and DEALLOCATE statements	233
Fortran 90 array manipulation intrinsics	235
Vector and matrix multiply functions	235
DOT_PRODUCT	235
MATMUL	236
Reduction functions	236
ALL	237
ANY	237
COUNT	238
MAXVAL	238
MINVAL	239
PRODUCT	239
SUM	240
Construction functions	241
MERGE	241
PACK	242
SPREAD	242
UNPACK	243
Manipulation functions	243
CSHIFT	243
EOSHIFT	244
TRANPOSE	245
Location functions	245
MAXLOC	246
MINLOC	246
Array assignments	247
Array-valued expressions	247
Masked array assignment	248

D Cray FORTRAN compatibility	253
Compiler defaults	253
Unsupported Cray features	254
Cray POINTER support	254
Debugging code containing Cray pointers	255
Cray automatic arrays	256
Cray BUFFERIN, BUFFEROUT support	256
Related statements and routines	256
Restrictions	257
Cray unformatted file support	257
Supported Cray library routines	258
Supported Cray intrinsics	258
Cray TASK COMMON support	259
Cray Boolean octal constant support	259
Cray Hollerith constants	260

E VAX FORTRAN compatibility	261
Supported features	261
Unsupported features	262
Miscellaneous differences	264
VAX FORTRAN records	265
Structure declaration	265
Field declaration	266
VAX floating point data	267
Supported VAX intrinsics	269

F Sun FORTRAN compatibility	271
------------------------------------	------------

Index	273
--------------	------------

Figures

Figure 1	Required order of statements	7
Figure 2	FORTRAN example conversion routine	143
Figure 3	C example conversion routine	144

Tables

Table 1	FORTRAN fields	4
Table 2	Data types	9
Table 3	Storage requirements for data types	11
Table 4	Arithmetic operators	31
Table 5	Arithmetic operator precedence	32
Table 6	Data type priority	33
Table 7	Logical operator precedence	34
Table 8	Array locations	52
Table 9	Conversion of expressions	71
Table 10	Data transfer I/O statements	93
Table 11	OPEN statement keywords	120
Table 12	RECORDTYPE defaults	129
Table 13	INQUIRE specifiers	134
Table 14	Data format conversion routine names	138
Table 15	User-supplied conversion routine names	145
Table 16	Shell variable attributes	146
Table 17	Character assignment for numeric I/O list elements	154
Table 18	Data conversion based on magnitude	165
Table 19	BLANK specifier defaults	167
Table 20	Default field descriptors	176
Table 21	List-directed output formats	184
Table 22	Vertical format control	186
Table 23	Built-in functions and defaults for argument lists	197
Table 24	Generic and specific intrinsics	207
Table 25	Unformatted Cray files readable by CONVEX FORTRAN	257
Table 26	VAX floating point data format names	268
Table 27	Supported Sun FORTRAN escape sequences	271

How to use this manual

Purpose and audience

This manual is a reference for the CONVEX FORTRAN programming language and is designed to provide a thorough working definition of the language. CONVEX FORTRAN is an implementation of the ANSI FORTRAN 77 standard that includes many extensions and enhancements not present in the standard. This manual encompasses both the *American National Standard programming language FORTRAN* (ANSI X3.9-1978) document and the CONVEX extensions to FORTRAN 77.

It is assumed throughout that you are an experienced FORTRAN programmer. For further discussion of the CONVEX FORTRAN language and other CONVEX software, refer to the "Associated documents" section in this preface.

If you are unfamiliar with the ConvexOS operating system, also consult the bibliography. Although a detailed knowledge of the operating system is not necessary for an understanding of this document, some familiarity with the system is beneficial.

Organization

This manual is organized as follows:

- Chapter 1 discusses FORTRAN program elements and program unit format.
- Chapter 2 discusses data types, constants and variables.
- Chapter 3 describes arrays and substrings.
- Chapter 4 discusses expressions.
- Chapter 5 discusses specification statements.
- Chapter 6 discusses the DATA statement.
- Chapter 7 describes the assignment statement and defines values used in a program.

- Chapter 8 describes functions and operations of control statements.
- Chapter 9 discusses files, units, input/output (I/O) statement components, data transfer I/O statements, and auxiliary I/O statements.
- Chapter 10 defines format specification descriptors and carriage-control options and separators.
- Chapter 11 discusses functions and operations of subprograms.
- Appendix A lists generic and specific intrinsic functions in table form, including the number of arguments, generic name, specific name, type of argument, and type of result. It also lists commonly used library routines.
- Appendix B discusses FORTRAN 66 compatibility.
- Appendix C discusses Fortran 90 compatibility.
- Appendix D discusses Cray FORTRAN compatibility.
- Appendix E discusses VAX FORTRAN compatibility.
- Appendix F discusses Sun FORTRAN compatibility.
- An index is included at the end of this manual.

Scope

This manual covers CONVEX FORTRAN Version 8.0, which runs under ConvexOS Version 10.0 or higher. The CONVEX FORTRAN compiler runs on all CONVEX hardware platforms, including C1, C2 and C3 Series architectures.

Notational conventions

The following conventions are used throughout this document:

- Brackets ([]) designate optional entries.
- A caret (^) represents the space character.
- Horizontal ellipsis (. . .) shows repetition of the preceding item(s). In an example, horizontal ellipsis indicates that statements are omitted.
- Vertical ellipsis shows continuation of a sequence where not all of the statements in an example are shown.

- References to the online man pages appear in the form `f c (1F)`, where the name of the manual page is followed by its section number enclosed in parentheses.
- *Italics* within text denote user-supplied variable information.
- Monospaced text, like this, is used to denote screen output, code examples, non-variable text in command forms, file names, utility names, and, in general, text that appears exactly as shown on output or must be entered exactly as shown on input.
- Within command sequences set apart from text, *italics* indicate user-supplied variable information. Substitute actual information for the italicized words. For example, the command sequence

`ld [options] [object files] [libraries]`

instructs you to type the command `ld`, followed by your choice of options, object files, or libraries.

- *Text describing CONVEX extensions to the FORTRAN language appear in this type style.* This convention is not followed in appendices that specifically deal with non-standard FORTRAN compatibility modes.
- Command sequences or forms which are CONVEX extensions appear in *monospaced italics, like this text.*

Notes and cautions

Notes and cautions are set apart from standard text. Examples are shown below.

Note

A note highlights information of a supplemental nature. The note immediately precedes or follows the highlighted information.

Caution

A caution highlights procedures or information necessary to avoid damage to equipment, damage to software, loss of data, or invalid test results.

Associated documents

The following documents, available from CONVEX Computer Corporation, are recommended to the CONVEX FORTRAN programmer:

- *CONVEX FORTRAN User's Guide* (DSW-037) describes how to compile and run CONVEX FORTRAN programs under the ConvexOS operating system, and describes the features and utilities included with the CONVEX FORTRAN compiler.
- *CONVEX FORTRAN Optimization Guide* (DSW-034) describes the different types of optimization available in CONVEX FORTRAN and shows you how to use optimization directives and options.
- *CONVEX Interlanguage Programming Guide* (DSW-043) outlines techniques for calling CONVEX FORTRAN routines from programs written in other languages, and for calling routines written in other languages from CONVEX FORTRAN.
- *CONVEX Application Compiler User's Guide* (DSW-401) (optional product) describes how to use the CONVEX Application Compiler to optimize programs.
- *ConvexOS Primer* (DSW-133) has basic self-instruction for learning and using the ConvexOS operating system.
- *ConvexOS Man Pages for Programmers* (DSW-332) contains copies of Sections 2 through 5 of the online man pages. These man pages are primarily concerned with operating system information for programmers.
- *ConvexOS Man Pages for Users* (DSW-331) contains copies of Sections 1 and 7 of the online man pages. These man pages are primarily concerned with operating system information for users.
- *CONVEX adb Debugger User's Guide* (DSW-009), a tutorial and reference manual, describes the functions and operations of the CONVEX adb debugger.
- *CONVEX Consultant User's Guide* (DSW-025) (optional product) describes the functions and operations of the CONVEX csd debugger, postmortem dump (pmd) utility, and the gprof profiler.
- *CONVEX Compiler Utilities User's Guide* (DSW-096) describes the CONVEX loader and the CONVEX assembler.

- *CONVEX COVUEshell Reference Manual* (DSW-136) (optional product) describes COVUEshell. COVUEshell is an optional CONVEX product that provides a VMS-type interface, giving the user access to a subset of Digital Command Language (DCL) commands.
- *CONVEX Performance Analyzer (CXpa) User's Guide* (DSW-251) (optional product) describes how to use the interactive profiler at the various levels that are available.
- *CONVEX CXdb Concepts* (DSW-471) and the *CONVEX CXdb Reference* (DSW-472) (optional products) describe all aspects of the optional CXdb debugger, which is briefly described in Chapter 4 of the *CONVEX FORTRAN User's Guide*.
- *CXmetrics User's Guide* (DSW-475) (optional product) describes software metrics data and explains how to use CONVEX CXmetrics to report on this data.

Other documents of interest include the following:

- *American National Standard programming language FORTRAN*, ANSI X3.9-1978. This book is the definition of standard FORTRAN 77, which is fully supported in CONVEX FORTRAN Version 8.0.
- ISO/IEC 1539:1991, the International Fortran Standard. This book is the definition of standard International Fortran, which is identical to the ANSI Fortran 90 programming language, ANSI X3.198-1992, certain features of which are supported in CONVEX FOTRAN V8.0.

Technical assistance

If you have questions that are not answered by the documentation, contact the CONVEX Technical Assistance Center (TAC). To contact the TAC, use one of the following phone numbers:

- Within the continental U.S., call (800) 952-0379.
- Outside the continental U.S., contact the local CONVEX office.

The contact utility

The TAC recommends using the contact utility to report a hardware, software, or documentation problem. The contact utility is an interactive program that helps the TAC track reports and route them to the CONVEX personnel most qualified to fix a problem.

After you invoke `contact`, it prompts you for information about the problem. When you finish your report, `contact` mails it to the TAC electronically. The TAC notifies you within 48 hours that your report has been received.

Using `contact` requires:

- A UNIX-to-UNIX Communication Protocol (UUCP) connection to the TAC
- The full path name of the program or utility in question
- The version number of the program or utility in question

Refer to the `contact(1)` man page for complete details.

Introduction

CONVEX FORTRAN is a high-level language that increases programmer productivity, maximizes software portability, and enhances the speed of execution using global and local optimization, vectorization, and parallelization techniques. CONVEX FORTRAN includes standard FORTRAN functions as defined by the American National Standard FORTRAN 77 (ANSI X3.9-1978) and unique CONVEX extensions. *This sentence illustrates the type style that is used to describe CONVEX extensions throughout this document.*

Types of programs

An executable program consists of a main program and, optionally, one or more subprograms. A program unit is defined as a sequence of FORTRAN statements that ends with an END statement. *A program unit can also include an OPTIONS statement, and comment lines.*

A main program can begin with a PROGRAM statement but cannot begin with a FUNCTION, SUBROUTINE, or BLOCK DATA statement. A subprogram must begin with a FUNCTION, SUBROUTINE, or BLOCK DATA statement.

FORTRAN character set

The standard FORTRAN 77 character set consists of the following characters:

Uppercase letters	A through Z
Digits	0 through 9
Special characters	blank = + - * / () , . \$ ' :

The CONVEX extended character set includes the following:

<i>Lowercase letters</i>	a - z
<i>Exclamation mark</i>	!
<i>Percent sign</i>	%
<i>Ampersand</i>	@
<i>Quotation mark</i>	"
<i>Underscore</i>	_
<i>Left angle bracket</i>	<
<i>Right angle bracket</i>	>
<i>Pound sign</i>	#
<i>Semicolon</i>	;
<i>Tab</i>	

Additional ASCII printable characters can appear in FORTRAN statements only as part of character or Hollerith constants. You can, however, use all printable ASCII characters in comment lines.

Blanks (spaces) can be used to improve readability of a program. Blanks are ignored unless they appear within a character string or a Hollerith constant or as an editing specification.

Comment line

A comment line has no effect on the actual execution of a program. It is used for documenting program action, identifying processes, or improving program readability. You can place a comment line anywhere in a program unit, even before the initial line or between continuation lines. You cannot continue a comment line using the continuation indicator.

The letter C or an asterisk (*) in column 1 of a line indicates a comment line. *Also, an exclamation point (!) in any column except column 6 indicates that the remainder of the line is comment text.* You can begin the comment text anywhere on the line following the comment indicator. A line containing only blanks is also a comment line.

FORTRAN statements

FORTRAN statements are classified as executable or nonexecutable. Executable statements specify action; they form an execution sequence in an executable program. Nonexecutable statements indicate characteristics, arrangement, and initial values of variables; contain editing information; classify program units; and designate entry points within subprograms.

If you are entering your CONVEX FORTRAN program from a terminal, you can enter statement lines of any length as long as you do not exceed the compiler's statement length limit (refer to the *CONVEX FORTRAN User's Guide*, Appendix F, for the statement length limit) or use the `-72` compiler option.

Executable statements include the following:

- Arithmetic, logical, statement label (ASSIGN), and character assignment
- Unconditional GOTO, assigned GOTO, and computed GOTO
- Arithmetic IF and logical IF
- Block IF ELSE IF, ELSE, and END IF
- CONTINUE
- STOP and PAUSE
- DO and ENDDO
- WHERE, ELSEWHERE, and ENDWHERE
- ALLOCATE and DEALLOCATE
- READ, WRITE, ACCEPT, TYPE, and PRINT
- REWIND, BACKSPACE, ENDFILE, OPEN, CLOSE, and INQUIRE
- CALL and RETURN
- END

Nonexecutable statements include the following:

- PROGRAM, FUNCTION, SUBROUTINE, ENTRY, and BLOCK DATA
- DIMENSION, COMMON, EQUIVALENCE, IMPLICIT, PARAMETER EXTERNAL, INTRINSIC, ALLOCATABLE and SAVE
- INTEGER, REAL, DOUBLE PRECISION, COMPLEX, LOGICAL, and CHARACTER
- DATA
- FORMAT
- Statement function

Each FORTRAN statement is divided into fields providing for statement label, continuation indicator, statement text, and a sequence number. Table 1 gives general rules for entering items into fields.

Table 1
FORTRAN fields

Column	Purpose	Use
1	Comment or debugging statement indicator	The letter <i>C</i> , <i>exclamation point (!)</i> , or asterisk indicates a comment line. <i>The letter D designates a debugging statement.</i>
1-5	Statement label or compiler directive	A statement label has one to five digits. <i>C\$DIR indicates a compiler directive.</i>
6	Initial or continuation line	A zero or blank indicates an initial line; any other character indicates a continuation line; <i>for tab formatting, a tab followed by a digit 1 through 9 indicate continuation.</i>
7-end of line	FORTRAN statements	This field contains the actual FORTRAN statement. <i>At any point in the field, except within character or Hollerith constants, an exclamation point indicates that the remaining text is a comment.</i>

You can use either character-per-column or *tab-character formatting*. *Tab formatting is convenient for terminal entry.*

Character-per-column formatting

This section describes the column-by-column contents of each field of a FORTRAN statement. Optional information that can be entered in the various fields is also discussed.

Statement label field

A statement label references statements in a program unit. Although any statement can have a label, only labeled executable statements and `FORMAT` statements can be referenced by other statements. Two statements in a program unit cannot have the same label.

The statement label, which must be a decimal integer, can be positioned in any column, 1 through 5. Blanks and leading zeros are ignored; for example, 5, 05, and 00005 are the same label.

Initial line

An initial line begins a single FORTRAN statement. A 0 or a blank (space) in column 6 indicates an initial line. An initial line cannot be a comment line; however, it can have a statement label. If you do not label the line, leave columns 1 through 5 blank.

A FORTRAN statement may be comprised of a single initial line or an initial line and all following continuation lines (up to the next initial line). Comments and/or blank lines are allowable in between initial lines and continuation lines, among continuation lines, and in between statements.

Continuation line

A continuation line is a line with no statement label that contains blanks in columns 1-5, and any character other than a blank or 0 in column 6. A continuation line can also begin with a single tab character followed by a digit 1-9.

Statement text field

The statement text begins in column 7 and continues to the end of the line *or until an exclamation mark is encountered, except within character or Hollerith constants.* (Refer to the *CONVEX FORTRAN User's Guide*, Appendix F, for the maximum allowed line length.) The interpretation of a statement is not affected by spaces and tabs except when they appear within character and Hollerith constants. To continue a statement on the next line, use the continuation indicator.

Two statements, separated by a semicolon (;), can appear on the same line. The character following the semicolon is treated as column 7 of the second statement. Thus, the statement following the semicolon cannot be a comment, specify a continuation field, or contain a label.

Debug statements

With CONVEX FORTRAN you can place a debugging statement indicator, the letter D, in column 1 of the statement label field. You can add a statement label in the remaining columns of the label field. To continue a debugging statement over more than one line, begin each new continuation line with a D in column 1 and a continuation indicator.

You can treat debugging statements as comments or as source text to the compiler. If you use the compiler command-line option -dc, the statements are treated as source text to the compiler; omitting this option causes the statements to be treated as comments.

Compiler directives

A compiler directive provides information to the compiler or instructs the compiler to override certain conditions that inhibit optimization, vectorization, or parallelization. A compiler directive begins with `C$DIR` in columns 1 through 5 and must fit on one line. Refer to the CONVEX FORTRAN User's Guide, Appendix A, for more information.

ANSI-standard formatting

The ANSI FORTRAN standard specifies a restricted form of character-per-column formatting and states that all lines are 72 characters in length. You can achieve the same effect with CONVEX FORTRAN by specifying the `-72` option on the compiler command line. If `-72` is specified, any line shorter than 72 characters is padded with blanks and any characters beyond 72 are ignored. A tab counts as one character.

Tab-key formatting

Tab-key formatting is a shorthand method of skipping to the various fields of a CONVEX FORTRAN statement. As with character-per-column formatting, the statement label must appear in the first 5 columns of the line. The next character is the tab character. If the character immediately following the tab is a digit from 1 to 9, this specifies a continuation as if the character had been in column 6 in character-per-column format.

Any other character after the tab is considered to be the first character of the statement field, as if it had been entered in column 7 in character-per-column format.

Order of statements and lines

Figure 1 shows the order required for a CONVEX FORTRAN program unit. You can mix statements separated by vertical lines but not statements separated by horizontal lines. For example, you can intersperse `DATA` statements with statement function definitions and with executable statements. You cannot, however, mix statement function definitions with executable statements.

Figure 1
Required order of statements

Comment lines and <i>INCLUDE</i> statements	<i>OPTIONS</i> statement			
	PROGRAM, FUNCTION, SUBROUTINE, or BLOCK DATA statement			
	NAMELIST, FORMAT, and ENTRY statements	IMPLICIT NONE statement		
		IMPLICIT statements		PARAMETER statements
		DATA statements	Other specification statements	
			Statement function definitions	
			Executable statements	
		END statement		

Symbolic names

Variables, arrays, and functions have symbolic names that identify them in a program. A symbolic name must start with a letter and can be followed by any number of uppercase letters (A-Z), digits (0-9), lowercase letters (a-z), underscores (_), or dollar signs (\$), up to the maximum length allowed for a symbolic name (refer to the CONVEX FORTRAN User's Guide, Appendix F, for the maximum name length).

The compiler converts letters specified in lowercase (a-z) to uppercase; thus the symbolic names "ABC" and "abc" are the same. Because dollar signs are used in CONVEX-supplied software, CONVEX recommends that you not use dollar signs in symbolic names.

OPTIONS statement

You can use the *OPTIONS* statement to include options that are not specified on the FORTRAN command line or to override options that are specified on the command line. The options remain in effect only within the program unit in which they are defined. If used, the *OPTIONS* statement must be the first statement in a program unit.

The *OPTIONS* statement has the form

OPTIONS option [option...]

Any options that can be specified on the command line can be specified in the *OPTIONS* statement.

Example:

```
OPTIONS -O2 -r8
```

This statement specifies that the following program unit is to be compiled at optimization level `-O2` and that all `REAL` data that is not explicitly sized is to be compiled as `DOUBLE PRECISION`.

**INCLUDE
statement**

The `INCLUDE` statement causes the compiler to insert source code from the specified file into the program being compiled. The contents of the file are inserted at the place where the `INCLUDE` statement appears. The statement has the form

```
INCLUDE 'filename'
```

or

```
#include "filename"
```

where `filename` is the path name of the file from which source code is to be read. The first form of the statement (without the `#` sign) is the recommended form; however, it cannot be used with the FORTRAN preprocessor (`fpp`).

When the compiler reaches the end of the included file, compilation resumes with the statement following the `INCLUDE` statement. The included file can itself contain an `INCLUDE` statement; `INCLUDE` statements can be nested up to the system limit as described in Appendix F of the CONVEX FORTRAN User's Guide. The `INCLUDE` statement can appear anywhere within a program unit.

Note

The `INCLUDE` statement operates somewhat differently under COVUEshell. Please refer to the CONVEX COVUEshell Reference Manual for more details.

Data types, constants, and variables

2

The symbolic name associated with each constant, variable, array and function contained in a FORTRAN program is assigned a data type, either implicitly or explicitly. This chapter discusses CONVEX FORTRAN data types and their application to constants and variables.

Data types

The *American National Standard programming language FORTRAN* defines five data types: CHARACTER, COMPLEX, INTEGER, LOGICAL, and REAL. *CONVEX FORTRAN supports the RECORD derived data type as an extension. Use of the RECORD type is discussed in detail in Appendix E, "VAX FORTRAN compatibility."*

Table 2 shows the ANSI FORTRAN data types available in CONVEX FORTRAN, along with their CONVEX equivalents.

Table 2
Data types

ANSI data type...	Includes types...
CHARACTER	CHARACTER*n, CHARACTER*(*)
COMPLEX	COMPLEX*8, COMPLEX*16, DOUBLE COMPLEX
INTEGER	INTEGER*1 (BYTE), INTEGER*2, INTEGER*4, INTEGER*8
LOGICAL	LOGICAL*1, LOGICAL*2, LOGICAL*4, LOGICAL*8
REAL	REAL*4, REAL*8, DOUBLE PRECISION, REAL*16 [†]

[†]REAL*16 is supported in native mode only.

In CONVEX FORTRAN, INTEGER*4 corresponds to the standard INTEGER data type, REAL*4 to REAL, COMPLEX*8 to COMPLEX, and REAL*8 to DOUBLE PRECISION. COMPLEX is an ordered pair of real values representing the real and imaginary parts of a complex number. *The DOUBLE COMPLEX (or COMPLEX*16) data type differs from COMPLEX*8 in that its parts are double-precision rather than single-precision. BYTE is a synonym for INTEGER*1.*

For LOGICAL and INTEGER data types, the storage requirement can be controlled by the -i2, -i4, -i8, -p8 or -pd8 compiler options; the default is four bytes. For REAL and COMPLEX, the storage requirements can be controlled by the -r4, -r8, -p8, or -pd8 compiler options.

For the CHARACTER data type, *n* can have an integer value ranging from 1 to the maximum length permitted by the system (refer to Appendix F, "System limits," of the CONVEX FORTRAN User's Guide). The notation CHARACTER* (*) specifies an assumed-length character string.

Table 3 shows the storage requirements for each data type. Defaults appear inside parentheses, but the defaults change when you use the -rn, -in, -p8, -pd8, or -cfc compiler options. Refer to Chapter 1 of the CONVEX FORTRAN User's Guide for more information on compiler options.

Table 3
Storage requirements for data types

Data type	Storage requirements (bytes)
LOGICAL	1, 2, (4), or 8
LOGICAL*1	1
LOGICAL*2	2
LOGICAL*4	4
LOGICAL*8	8
INTEGER	1, 2, (4), or 8
INTEGER*1, BYTE	1
INTEGER*2	2
INTEGER*4	4
INTEGER*8	8
REAL	(4), 8, or 16
REAL*4	4
REAL*8	8
REAL*16	16
COMPLEX	(8) or 16
COMPLEX*8	8
COMPLEX*16	16
DOUBLE PRECISION	(8) or 16
DOUBLE COMPLEX	(16)
CHARACTER	1
CHARACTER*len	len
CHARACTER* (*)	
RECORD	Varies

Some data types can be assigned either explicitly or implicitly. Type-declaration statements are used for explicit typing; refer to the "Type-declaration statements" section of Chapter 5, "Specification statements," for details. Symbolic names that are not explicitly typed are typed implicitly according to their first letter; names beginning with the letters I through N are of the

default INTEGER type; all others are of the default REAL type. The IMPLICIT statement allows you to override these implicit type assignments; refer to the "IMPLICIT statement" section of Chapter 5, "Specification statements," for more information.

Conversion of data types

Where data types differ between the variable or array element on the left side and the expression on the right side of an assignment statement, CONVEX FORTRAN converts the expression on the right to the same data type as that on the left side. The compiler uses the following rules for conversion:

- *Treats LOGICALS as INTEGERS of the same length.*
- *Converts INTEGERS to longer types by sign extension and to shorter types by truncation. Truncation of significant bits causes an integer overflow.*
- *Converts INTEGER values to REAL values by truncation. For example, 82762035 is too large to fit in a REAL*4 without loss of precision, so just enough rightmost bits of the binary representation are truncated to make it fit. After conversion, the value becomes 82762033.0.*
- *Converts REAL values to INTEGER values by truncation. Rounding is not performed. For example, I = 5.9 assigns the value 5 to I.*
- *Converts a REAL value to a lower precision REAL value (for example, REAL*8 to REAL*4) by rounding to the lower precision.*
- *Converts a REAL value to a higher precision REAL value by zero-extending the mantissa. Converting a REAL value to a higher precision, however, does not increase the accuracy of the value.*
- *Converts COMPLEX values to other noncomplex numeric data types by converting the real part only. For example, R = (15.6d0, 7.5d6) assigns the value 15.6 to R.*
- *Converts noncomplex values to COMPLEX by converting first to the appropriate precision REAL value to get the real part, and then assigning 0.0 or 0.0d0 to the imaginary part.*
- *Handles COMPLEX-to-COMPLEX conversions by converting the real and imaginary parts separately, that is, as two REAL conversions.*

Note

Optimizing code containing type conversions (including implicit type conversions) can cause some code to vectorize poorly or not at all.

Constants

A constant is an arithmetic or logical value or a character string. It does not change during program execution. The form in which a constant is expressed determines the value and data type. The `PARAMETER` statement assigns a symbolic name to a constant. Refer to Chapter 5, "Specification statements," for more information on the `PARAMETER` statement.

Integer constants

An `INTEGER` constant consists of the digits 0 through 9 and, possibly, a leading sign. The sign is optional for a positive constant but required for a negative constant. Leading zeros have no effect on the value.

Examples:

Valid	Invalid	Reason
248	24.8	Has decimal point
54	5E8	Uses exponential notation
12333	12,000	Has comma

*Integer constants take the default `INTEGER` precision. If the default precision is 2 or 4 bytes and the constant is too large, it is treated as an 8-byte constant; if it is too large to be represented in 8 bytes, the compiler truncates the constant to fit and issues a warning. The `-in`, `-cfc`, `-p8`, and `-pd8` compiler options affect the default precision. Refer to Chapter 1 of the *CONVEX FORTRAN User's Guide* for more information.*

Real constants

A real constant consists of an optional positive sign or required negative sign, digits (0-9), a decimal point, and an optional exponent. The exponent is represented as the letter E followed by an integer that denotes the power of 10. You can place the decimal point anywhere in the string (for example, 2.1, .2, 678912.). When you specify an exponent, the decimal is optional; 7.E3 is the same as 7E3.

A *DOUBLE PRECISION* constant is similar to *REAL* except that an exponent is required and the exponent letter *D* is used. *REAL*16* data requires the exponent letter *Q*. If the *Q* is omitted, the value defaults to *REAL*.

REAL and *DOUBLE PRECISION* constants take the default precision for their corresponding variable types. If the constant is too large, the compiler issues a fatal error. The *-rn*, *-cfc*, *-p8*, and *-pd8* compiler options affect the default precision. Refer to Chapter 1 of the *CONVEX FORTRAN User's Guide* for more information.

Examples:

Valid	Invalid	Reason
2500.	2500	Decimal point missing
+2.0E2	2.0E	Exponent field missing
5E4	5,000	Has comma
4E-2		
3.0E4		
5.4Q4		

Complex constants

Both *COMPLEX*8* and *COMPLEX*16* constants consist of a pair of real constants separated by a comma and enclosed in parentheses. The first constant is the real part and the second constant is the imaginary part. A *COMPLEX*8* constant is a pair of *INTEGER*2*, *INTEGER*4*, or *REAL*4* constants. A *COMPLEX*16* constant is an ordered pair of integer, *REAL*4* or *REAL*8* constants.

Example:

Valid	Invalid	Reason
(1.6405D0, -1.6405D0)	(1.640D)	Second constant missing

COMPLEX constants take the default *COMPLEX* precision. If the constant is too large, the compiler issues a fatal error. The *-cfc*, *-p8*, and *-pd8* compiler options affect the default precision. Refer to Chapter 1 of the *CONVEX FORTRAN User's Guide* for more information.

Octal constants

An octal constant consists of one or more octal digits enclosed in apostrophes and followed by the letter *o*. An octal digit can range from 0 to 7. An octal constant has the form

`'cc...c'o`

or

`O'cc...c'`

where *c* represents an octal digit.

Examples:

Valid	Invalid	Reason
<code>'765'o</code>	<code>'835'o</code>	8 not in range 0 to 7
<code>'123'o</code>	<code>123O</code>	Missing apostrophes

If you specify the `-vfc` compiler option (refer to Appendix E, "VAX FORTRAN compatibility"), the compiler accepts octal constants of the form:

`"cc...c`

where *c* represents an octal digit.

If you specify the `-cfc` compiler option (refer to Appendix D, "Cray FORTRAN compatibility"), the compiler accepts octal constants of the form:

`cc...cB`

where *c* represents an octal digit. This is compatible with Cray's Boolean octal constant form; refer to Appendix D, "Cray FORTRAN compatibility," for details.

Hexadecimal constants

A hexadecimal constant consists of one or more hexadecimal digits enclosed in apostrophes and followed or preceded by the letter X. A hexadecimal digit can range from 0 to 9, A to F (or a to f). A hexadecimal constant has the form:

`'cc...c'X`

or

`X'cc...c'`

or

`Z'cc...c'`

where *c* represents a hexadecimal digit.

Examples:

<i>Valid</i>	<i>Invalid</i>	<i>Reason</i>
<code>'1A6'X</code>	<code>'FFG'X</code>	G not in range 0 - 9, A - F
<code>'123'X</code>	<code>'12.4'X</code>	Decimal point not allowed
<code>'FFFFFFFF'X</code>	<code>1AB2X</code>	No apostrophe
<code>X'66F'</code>	<code>X129A'</code>	Missing first apostrophe
<code>'abc123ff'X</code>		

Octal and hexadecimal constants assume data types depending on how they are used. The following conditions apply:

- When octal or hexadecimal constants are used as actual arguments, no data type is assumed.
- When you use either of the constants with a binary operator, the data type of the constant matches the data type of the other operand.
- When a specific data type is required, that type is assumed for the constant.
- In any other context, the type is `INTEGER*4` for the constant.

When the number of digits required exceeds the length of the constant, the leftmost places are filled with zeros. When the length of the constant exceeds the number of digits required, the excess digits are truncated on the left; an error message results if any of the truncated digits are nonzeros.

Example:

```
INTEGER*4 i,j
```

Truncation/extension occurs as shown below:

```
i = '12'X           (same as '00000012'X)
j = '7777ffffffff0076'X (same as 'ffff0076'X)
```

Hollerith constants

Hollerith constants are strings of printable ASCII characters preceded by a character count and the letter H. A Hollerith constant has the form:

```
nHcc...c
```

where

n specifies the number of the characters in the constant (including spaces and tabs).

c is a printable ASCII character.

The value of *n* must be an unsigned positive integer greater than zero.

Example:

Valid	Invalid	Reason
4HHe1p	0H	Must contain at least one character

If you specify the `-cfc` compiler option (refer to Appendix D, "Cray FORTRAN compatibility"), the compiler accepts Hollerith constants of the form:

```
nLcc...c
```

where *n* and *c* take the same values as their standard-form counterparts.

The Cray form left-justifies the constant in memory and zero-fills its storage space to the right. You must supply the `-cfc` option when using this form. Refer to Appendix D, "Cray FORTRAN compatibility," for more information.

Hollerith constants assume data type according to the context in which they are used:

- When a specific data type is required, that type is assumed for the constant.
- When the constant is used as an argument, no data type is assumed.
- When the constant is used with a binary operator, the data type of the constant is that of the other operand. Although the data type is that of the other operand, the bit pattern is taken from the Hollerith constant.
- If you pass a Hollerith constant to a subroutine, you must pass it into a dummy argument of type `INTEGER`, `REAL`, or `LOGICAL`. Passing it into a `CHARACTER` variable will cause erroneous behavior. Refer to the ANSI FORTRAN 77 standard, page C-3, for more information.
- In any other context, the constant assumes an `INTEGER*4` data type, unless you change the default `INTEGER` length by using the `-i2`, `-i8`, `-p8`, `-pd8` or `-cfc` options.
- When a Hollerith constant continues across a line, you can use the `-72` option to imitate the behavior of other FORTRAN compilers (blank padding to column 72).

Note

If you continue a Hollerith constant on another line, you must use the `-72` option for compilation.

Logical constants

A logical constant represents the value true or false. The following forms are acceptable:

- `T` or `F`
- `.T.` or `.F.`
- `.TRUE.` or `.FALSE.`

Character constants

A character constant is a string of printable ASCII characters enclosed within delimiting apostrophes. The value of the character constant includes characters, spaces, and tabs between the delimiting apostrophes. The delimiting apostrophes are not part of the value, but every string must begin and end with them. Within a string, use two consecutive apostrophes (' ') to represent one apostrophe.

Examples:

```
'final'  
'two''s complement'
```

You can delimit a character constant with quotation marks (") instead of apostrophes. In either case, the beginning delimiter must be the same as the ending delimiter. When quotation marks are the delimiters, use two consecutive quotation marks (" ") within a string to represent one quotation mark.

Examples:

```
"begin"  
"two's complement"  
'double"quote'  
'bad string"           (invalid)
```

Note

If you continue a character constant on a second line, you must use the `-z2` option for compilation.

Variables

A variable represents a value that can be changed during program execution by an assignment or `READ` statement. You can assign an initial value to a variable with a `DATA` statement (refer to Chapter 6, "DATA statement") or a type-declaration statement (refer to Chapter 5, "Specification statements"). A variable is associated with a storage location. Whenever a variable is used, the current value in the storage location is referenced.

Variable types are classified and assigned according to the methods discussed in the "Data types" section of this chapter.

Multiple variables can be associated with the same storage location by using `COMMON` statements, `EQUIVALENCE` statements, or actual arguments and dummy arguments in subprogram references. The `COMMON` statement allows two or more variables in different program units to share the same storage unit. The `EQUIVALENCE` statement allows variables in the same program unit to share the same storage unit (refer to Chapter 5, "Specification statements").

An array is a set of storage locations of the same data type identified by a single name and accessed individually by a subscript. A substring is a contiguous portion of a character string. This chapter discusses both arrays and substrings in detail.

Arrays

Conceptually, a one-dimensional array is a column of elements, with each element accessible by its subscript. A two-dimensional array can be thought of as a table of elements containing rows and columns, with each element accessible by a pair of subscripts. *Under the -f90 option, you can reference an entire array by using the array name, and you can reference sections of the array with subscript ranges.* An array can have from one to seven dimensions, with the number of subscripts used to reference individual elements equal to the number of dimensions. Refer to the "Array subscripts" section of this chapter for more information.

All the values in an array have the same data type and any value assigned is converted to the data type of the array. A `DATA` statement can be used to define an array element or an entire array before program execution. During execution, an array element is defined with an assignment or input statement; the entire array can be defined with an input statement, *or, through use of the -f90 flag, with an assignment statement.*

The number of dimensions of an array is referred to as the array's rank. A rank definition has the form rank n , where n is the number of dimensions. Related to this is the array's shape, which describes the absolute number of elements for each dimension. Shape has the form (m_1, m_2, \dots, m_n) , where m is the number of elements in the given dimension, and n is the rank of the array. Arrays having the same shape are said to be

conformable. Scalar values are conformable with any shape array; when a scalar is used in this context, it can be thought of as filling an array of the proper shape with its value.

Rank, shape and conformability are important in discussing Fortran 90 array support, which CONVEX FORTRAN provides in a limited capacity through use of the -f90 compiler option. Supported Fortran 90 array features are discussed under applicable topics in this chapter; all Fortran 90 features present in CONVEX FORTRAN are discussed in detail in Appendix C, "Fortran 90 compatibility."

Array declaration

DIMENSION, COMMON, ALLOCATABLE, POINTER or type statements allow array declarations. An array declarator defines the name of the array within the program unit, the number of dimensions, and the upper and lower bounds of elements in each dimension. For multidimensional arrays, separate the dimension declarators with commas.

The -f90 and -cfc compiler options enable the use of automatic arrays. Any array declared local to a subroutine with one or more non-constant dimensions is considered an automatic array. Automatic arrays are discussed briefly in this section; for more information, see Appendix C, "Fortran 90 compatibility."

Array declarations have the following form:

```
type arrname ([lb:] ub [, ...])
```

or

```
DIMENSION arrname ([lb:] ub [, ...]) [, arrname...]
```

where

type

is the type of the array elements. Any valid CONVEX FORTRAN type can be used. If the DIMENSION form is used, the array will take an implicit type based on its name if it is not explicitly typed elsewhere.

arrname

is the name of the array. Multiple arrays can be dimensioned in a single DIMENSION statement. *If the array is to be automatic, arrname must be local to a subroutine.*

lb

is the lower dimension bound. The default is 1.

For static arrays, *lb* must be a constant. For allocatable arrays, *lb* can be a constant, variable or expression.

ub

is the upper dimension bound. *ub* must be greater than or equal to *lb*.

For static arrays, *ub* must be a constant. For allocatable arrays, *ub* can be a constant, variable or expression.

If you do not specify a lower bound, the default value is 1, and the upper bound is the number of elements given for the dimension. To use a lower bound that is not 1, you must specify both bounds. The bounds values can be positive, negative, or zero. Separate lower- and upper-bound values with a colon.

*For automatic arrays, *ub* and/or *lb* for at least one dimension must either consist of or be derived from an integer argument passed to the subroutine in which the automatic array resides.*

Different dimensions of a multidimensional automatic array can be declared to be different sizes based on one or more dimension arguments passed to the subroutine.

The type of array and the product of the subscripts determine the number of storage units allocated to each array named in the statement. FORTRAN arrays can have from one to seven dimensions.

The examples below illustrate several forms of both static and automatic array declarations.

Example 1:

```
DIMENSION MYRAY (5)
```

This example dimensions a one-dimensional static array with five elements. The array is implicitly typed `INTEGER`.

Example 2:

```
DIMENSION MYRAY (0:4)
```

This example dimensions a one-dimensional static `INTEGER` array with five elements numbered 0 through 4.

Example 3:

```
COMMON RERAY (5,5)
```

This example allocates storage in an unnamed common block for a two-dimensional static array with 25 elements (5 rows \times 5 columns). The array is implicitly typed REAL.

Example 4:

```
COMMON MYRAY (-1:3,2:6)
```

This example specifies a two-dimensional static INTEGER array with 25 elements. Elements in the first dimension are numbered -1 through 3; elements in the second are numbered 2 through 6.

Example 5:

```
INTEGER A(5,5,5)
```

This example declares a three-dimensional static array with 125 elements (5 planes \times 5 rows \times 5 columns) of type INTEGER. The INTEGER type statement overrides implicit typing.

Example 6:

```
CHARACTER*8 MRAY (2)
```

This example declares a one-dimensional static CHARACTER array with two elements of eight bytes each.

The example below illustrates two ways of declaring multidimensional automatic arrays.

```
SUBROUTINE SUB (I, J, K)
.
.
.
DIMENSION IARR (I, J, 2*I) !IMPLICITLY TYPED INTEGER
REAL*8 X (J+I, K)
.
.
.
```

In this example the arrays IARR and X are local to the subroutine SUB. I, J and K are integer arguments passed to the subroutine that are used in allocating the arrays. The arrays are

automatically allocated at runtime on entry into the subroutine. *IARR* has one dimension of length *I*, one of length *J*, and one of length $2 * I$. *X* has one dimension of length $J + I$ and one of length *K*. Both are automatically deallocated on exit from the subroutine.

Automatic arrays can only be used in subroutines and functions; they are not allowed in *MAIN* programs or *BLOCK DATA* subprograms or in *COMMON* blocks, nor can they be used as dummy arguments.

Allocatable array declarations

The *-f90* compiler option allows the use of allocatable arrays. Allocatable arrays are declared with the *ALLOCATABLE* statement. The array's rank must be supplied either in the *ALLOCATABLE* statement, in a *DIMENSION* statement, or in the array's type declaration, which, if used, precedes the *ALLOCATABLE* statement. Rank definitions have the following form:

```
arrname ( : [ , : . . . ] )
```

where *arrname* is the name of the array, which is followed by parentheses containing one colon for each dimension of the array. Multiple colons are separated by commas. Code examples of this are given in Appendix C, "Fortran 90 compatibility." Note that the above example is not an executable statement, but rather a form for use within certain executable statements where a rank definition is required.

Allocatable array declarations have the following form:

```
[type arrexp ]  
ALLOCATABLE (arrexp [ , . . . ] )
```

where

type

is an optional type definition for the array. The form for this is identical to the form for a standard array type definition, except instead of specifying constant dimension parameters, you must either supply a rank definition or provide no parameters.

arrexp

is the array name or rank definition if it was not given in a preceding type statement. The array name may occur in both the type and `ALLOCATABLE` statements; the rank definition must occur in exactly one, or the compiler flags an error.

Refer to Appendix C, "Fortran 90 compatibility," for a detailed example of an allocatable array declaration.

Fortran 90 allocatable arrays

Use of the `-f90` or `-cfc` flag enables Fortran 90 allocatable arrays. Declaration of these arrays is discussed under the "Allocatable array declarations" section of this chapter.

Storage for allocatable arrays is dynamically allocated on the heap at runtime when an `ALLOCATE` statement is executed. The heap space must be deallocated through use of the `DEALLOCATE` statement when the allocatable array is no longer needed.

After declaring an allocatable array, you must allocate storage for it before you can use it. This is done with the `ALLOCATE` statement, which has the following form:

```
ALLOCATE (arrdef [, ...])
```

where *arrdef* is an array definition with dimension values or ranges for all dimensions. Dimension ranges are described under the preceding "Array declarations" section of this chapter.

When you are done using an allocatable array, deallocate the array's storage space using the `DEALLOCATE` command. You can deallocate multiple arrays with one statement, or, if you finish with them at different points in the program, individually. The `DEALLOCATE` statement has the following form:

```
DEALLOCATE (arrname [, ...])
```

where *arrname* is the name of the array to be deallocated. *arrname* must be free of subscripts.

You must manually deallocate arrays that are allocated local to a subroutine before exiting the subroutine.

To change the size or subscript range of a presently allocated allocatable array, you must first deallocate the array, then allocate it with the new information.

Refer to Appendix C, "Fortran 90 compatibility," for more information on the `ALLOCATE` and `DEALLOCATE` statements, including usage examples.

Referencing array elements

Individual array elements are referenced by subscripts. A pair of parentheses that enclose one to seven subscript expressions separated by commas constitutes a subscript. The subscript immediately follows the array name. Specify one subscript expression for each dimension defined for the array. For a two-dimensional array declared as `COMMON MYRAY (5, 5)`, a valid reference is `MYRAY (2, 4)`; an invalid reference is `MYRAY (5, 7)`, because 7 lies outside the subscript range of the second dimension of the array. A subscript expression can be any valid arithmetic expression, or a Fortran 90 array section.

Use of the `-f90` compiler flag allows you to reference array elements using array sections. An array section is a piece of an array bounded by specified elements in each dimension. Array sections can be as small as one element or as large as the entire array. An array section is specified with the following form:

```
arrname( [i] : [j] [ :step ] [, ... ] )
```

where

arrname
is the array name.

i
is a constant, variable, or expression describing the element at which the section starts for that particular dimension of the array. *i* defaults to the declared lower bound of that dimension of the array.

j
is a constant, variable, or expression describing the element at which the section ends for that particular dimension of the array. *j* defaults to the declared upper bound of that dimension of the array.

step
is a constant, variable, or expression describing the number of elements to step along that particular dimension when selecting elements for the array section. *step* the default is 1.

Given the defaults for each of these values, array section specifiers such as $X(:, :)$ and $X(:)$ are allowed. These particular examples are equivalent to the entire array, which can also be denoted by X .

For a more detailed discussion of array sections and for examples of their use, see Appendix C, "Fortran 90 compatibility."

Array storage

Even though array elements are arranged and referenced in dimensions, array storage in memory is linear. For example, a one-dimensional array is a column of figures, stored with the first element in the first storage location and the last element in the last storage location of the sequence. Multidimensional array elements are stored so that the value of the first subscript (leftmost) varies most rapidly. This is called the "order of subscript progression."

Character substrings

A character substring is a sequence of contiguous characters that are part of a character variable or array element. A substring name identifies a character substring that can be assigned values and referenced. A character substring reference to a variable has the following form:

$$\text{var}[\textit{subscript}] ([\textit{expr1}] : [\textit{expr2}])$$

where \textit{var} is a character variable or array name, $\textit{subscript}$ is an array subscript (required only when \textit{var} is an array), $\textit{expr1}$ is an optional numeric expression indicating the leftmost character position of the substring, and $\textit{expr2}$ is an optional numeric expression indicating the rightmost character position of the substring. Character positions are numbered from left to right within a character variable or array element.

If $\textit{expr1}$ is omitted, a value of 1 is assumed; if $\textit{expr2}$ is omitted, the length of the character variable is assumed. The value of $\textit{expr1}$ must be positive and less than or equal to $\textit{expr2}$. The value of $\textit{expr2}$ must be less than or equal to the length of the string.

Example:

Given the following statements:

```
CHARACTER*14 NAME
NAME = 'CONVEX FORTRAN'
```

the reference:

`NAME(8:14)`

indicates a substring beginning with the position 8 (F) and ending in position 14 (N) of the variable `NAME`, giving the value of `FORTRAN` to the substring `NAME(8:14)`.

Example:

`EXAMPLE(1,5) (:3)`

indicates the substring begins with the first character position and ends with the third character position of the character array `EXAMPLE(1,5)`.

An expression is a combination of one or more operands and optional operators. During program execution, an expression specifies a computation or evaluation that produces either a scalar value *or an array value*. The operators determine the operations to be executed on the operands. An expression can be arithmetic, logical, relational, or character.

Arithmetic expressions

An arithmetic expression includes arithmetic operators and operands. Arithmetic operands include character, *Hollerith*, *octal*, and *hexadecimal* constants described previously, in addition to the standard FORTRAN 77 operands of numeric constants, numeric variables, numeric array elements, arithmetic expressions enclosed in parentheses, *array sections* or arithmetic function references. *The term numeric as used here includes logical data. The compiler treats logical data as integer data when it is used in an arithmetic context.*

The arithmetic operator specifies the computation to perform on the operands. This computation generates a numeric value. Arithmetic operators are listed in Table 4.

Table 4
Arithmetic operators

Operator	Function	Example
**	Exponentiation	C**2
*	Multiplication	C*2
/	Division	C/2
+	Addition	C+2
-	Subtraction	C-2
+	Unary plus (identity)	(+C)
-	Unary minus (negation)	(-C)

Operator precedence

When an expression has two or more operators and contains no parentheses, the operations are executed in the order given in Table 5.

Table 5
Arithmetic operator
precedence

Operator	Priority
**	Evaluated first
* and /	Evaluated second
+ and -	Evaluated last

Operators with the highest priority are processed before those with lower priority, except that parentheses within an expression cause the operation inside the parentheses to be performed first.

Examples:

```
6 * 2**2 - 5           ! Yields a value of 19
3 + 4 * 3 - 9/3        ! Yields a value of 12
(3 + 4) * 3 - 9/3      ! Yields a value of 18
```

When an expression has two or more operators of equal precedence, evaluation occurs in left-to-right order, except for exponentiation, which is evaluated right to left. CONVEX FORTRAN, however, can execute operations in differing orders as long as the order remains algebraically equivalent to left-to-right order of evaluation.

If more than one set of operators appears within parentheses, the operators are evaluated according to the normal order of precedence, unless overridden by parentheses within parentheses. In nested expressions, the innermost expression enclosed within parentheses is evaluated first.

Data type priority

Where operands of different data types are combined in an arithmetic expression, the higher-ranked argument determines the type, and the greater precision argument determines the precision. Data type ordering is shown in Table 6.

Table 6
Data type priority

Order	Data type
1 (lowest)	LOGICAL*1, *2, *4, *8
2	INTEGER*1 (BYTE), *2, *4, *8
3	REAL*4 (REAL), *8 (DOUBLE PRECISION), *16
4 (highest)	COMPLEX*8 (COMPLEX), *16 (DOUBLE COMPLEX)

*When LOGICAL and INTEGER types are combined, the resulting type is INTEGER, and the precision is the highest specified, whether *1, *2, *4, or *8. Even REAL*8 and COMPLEX*8 yielding COMPLEX*16 is consistent with this rule, if you regard the precision of COMPLEX*8 as 4 and of COMPLEX*16 as 8 (the precision of the real and imaginary parts). The CONVEX FORTRAN extensions in data types are LOGICAL*1, LOGICAL*2, LOGICAL*8, INTEGER*1, INTEGER*2, INTEGER*8, REAL*16, and COMPLEX*16.*

The data types of arithmetic expressions follow certain conventions:

- LOGICAL entities are treated as INTEGERS when used in an arithmetic context.
- REAL operations are performed only if one or more of the operands is REAL.
- COMPLEX operations involving COMPLEX*8 and REAL*8 elements are evaluated as COMPLEX*16 operations; therefore, the REAL*8 element is not rounded.

*These conventions also apply to arithmetic operations where one of the operands is a constant. Additional precision is used for the constant if a real or complex constant is used in an expression of higher precision. In such expressions, the effect is as if a REAL*8 representation of the constant had been given.*

Relational expressions

A relational expression compares either the value of two arithmetic expressions or the value of two character expressions that produce a logical value of true or false. The two expressions are separated by a relational operator (.LT., .LE., .EQ., .NE., .GT., and .GE.). Each relational operator must include delimiting periods.

Logical expressions

A logical expression consists of one logical operand or a combination of logical operands and logical operators. After evaluation, a logical expression produces a logical value of true or false. Logical operands in CONVEX FORTRAN can be any of the following:

- LOGICAL or *INTEGER* constant
- LOGICAL or *INTEGER* variable
- LOGICAL or *INTEGER* array element
- LOGICAL or *INTEGER* expression enclosed in parentheses
- LOGICAL or *INTEGER* function reference
- Relational expression

The evaluation of a logical expression that has two or more logical operators is based on operator precedence as shown in Table 7.

Table 7
Logical operator precedence

Precedence	Operator
Lowest	.EQV., .NEQV., (.XOR.) .OR. [†] .AND. [†]
Highest	.NOT.

[†] Evaluation of the .OR. and .AND. operators is short-circuited. Refer to Chapter 8, "Control statements," for more information.

The logical operator .XOR. is the same as .NEQV. Operators on the same level of precedence are interpreted from left to right. Arithmetic rules for operator precedence apply for evaluation. Mixed operations are evaluated first by arithmetic rules of precedence, next by relational operations, and last by logical operations. Logical operands appearing in IF conditionals are short circuited; refer to the "Short circuit evaluation of conditionals" section of Chapter 8, "Control statements."

The expression in the following example yields a value of FALSE:

```
K = 6
M = 2
N = 5
(K .LE. N) .AND. (N .GT. M)
```

As stated in the ANSI standard, logical operators used on logical values produce values of type `LOGICAL`. *Logical operators operating on integer values produce values of type `INTEGER`. The logical operation is carried out bit by bit on the corresponding bits of the internal binary representation of the integer elements. When a logical operator combines integer and logical values, the logical value is first converted to an integer. The operation is then carried out for the two integer elements; the resulting data type is `INTEGER`.*

Character expressions

The evaluation of a character expression results in a string of type `CHARACTER`. Within a character expression, two slashes can be used to specify concatenation. Concatenation produces a string that is a combination of the operand strings and executes from left to right.

Parentheses have no effect on the value of a character expression. If spaces are included in the expression, the spaces are part of the value.

Examples:

```
'MY' //'EXAMPLE'    ! Yields a value of MYEXAMPLE  
'MY ' //'EXAMPLE'  ! Yields a value of MY EXAMPLE
```


Specification statements are nonexecutable and appear before the first executable statement in a program unit. These statements define the type of variable or array, stipulate storage requirements for each variable based on its type, indicate the dimension of arrays, define storage sharing, and assign initial values to variables and arrays. Specification statements include:

- PROGRAM
- COMMON
- IMPLICIT
- PARAMETER
- Type-declaration statements:
 - Numeric
 - Character
 - Record
- *POINTER*
- DIMENSION
- *ALLOCATABLE*
- EQUIVALENCE
- *NAMELIST*
- EXTERNAL
- INTRINSIC
- SAVE

If you specify the `-sfc` (Sun FORTRAN) compiler option, you can use the `STATIC` and `AUTOMATIC` statements. For further information, refer to Appendix F, "Sun FORTRAN compatibility."

The DATA statement, which assigns initial values to variables, arrays, and array elements, is classified as an initialization statement. The DATA statement is described in Chapter 6, "DATA statement."

PROGRAM statement

The PROGRAM statement can be used to assign a name to the main program unit; its use is optional. If used, a PROGRAM statement must always be the first statement in the program *unless an OPTIONS statement is also included, in which case the PROGRAM statement immediately follows the OPTIONS statement.*

The PROGRAM statement has the form:

```
PROGRAM pgm
```

where *pgm* is the symbolic name for the main program.

Do not use the symbolic main program name as a name for an external procedure, block data subprogram, or common block in the same executable program. Also, do not use a symbolic main program name as a local name in the main program.

You cannot reference the main program from itself or from a subprogram. An executable program has only one main program.

COMMON statement

The COMMON statement allows variables or arrays in a main program or subprogram to share the same storage location with variables and arrays in other subprograms. These blocks of storage are called common blocks. Common blocks can be named or unnamed; unnamed blocks are called blank common. The block specification determines storage order of variables and arrays. Named common blocks of the same name can be of different sizes in different program units of an executable program.

The COMMON statement has the form:

```
COMMON [ /[cbn]/ ] nlist [ [, ] /[cbn] /nlist ] ...
```

where

cbn

is a symbolic name for a common block. If you do not specify a symbolic name (blank common), the first pair of slashes is optional.

nlist

is a list of variable names, array names, and array declarators.

An entity name (name in *nlist*) can appear only once in a COMMON statement within a program unit. If a common block name appears twice in the same program unit, the effect is as if the *nlist* of the second appearance followed the first *nlist*. You can use a common block name (*cbn*) more than once in a COMMON statement and in a program unit. A common block can have the same name as any local entity except a constant, intrinsic function, or a variable name that is also a function name. If you give a common block and variable the same name, all references to the name, except when it appears surrounded by slashes (/) in COMMON and SAVE statements, indicate the variable. Thus, SAVE X refers to the variable, while SAVE /X/ refers to the common block.

Data items appearing in COMMON block lists should be ordered from largest to smallest; quad-word (16 byte) items should appear first, followed by double-word items, single-word items, and finally partial-word items. Failure to order the data items this way will cause the compiler to issue a warning message, and, if left uncorrected, will impede performance.

Arrays and variables in COMMON can be initialized in DATA statements in program units other than BLOCK DATA subprograms. A variable in COMMON, however, can be initialized only in one program unit, although different variables in the same COMMON block can be initialized in different program units.

CONVEX FORTRAN supports Cray TASK COMMON statements when the -cfc flag is specified. Refer to Appendix D, "Cray FORTRAN compatibility," for more information.

IMPLICIT statement

The IMPLICIT statement allows you to override the implied data typing of symbolic names within a program unit. The IMPLICIT statement has the forms:

```
IMPLICIT typ (a[,a]...) [, [typ] (a[,a]...)]...
```

or

IMPLICIT NONE

where

typ

is an INTEGER [**len*], REAL [**len*], DOUBLE PRECISION, DOUBLE COMPLEX, COMPLEX [**len*], LOGICAL [**len*], or CHARACTER [**len*] data type.

a

is one letter or a range of letters; the range is expressed as first letter of range, minus sign (-), last letter of range (for example, F-H).

len

is an optional length specifier for the data type.

If used, the *IMPLICIT* statement must precede any other specification statements in the program unit. If this statement is not used, variable names that begin with the letters I through N imply type INTEGER; all others imply type REAL.

The IMPLICIT NONE form of the statement overrides all implicit defaults except intrinsic function types. When using IMPLICIT NONE, you must explicitly declare the data types of all symbolic names in the program unit. If you specify IMPLICIT NONE, no other IMPLICIT statement can be included in the program unit.

Examples:

```
IMPLICIT COMPLEX(F,H-J)
C ANY NAME BEGINNING WITH THE LETTER F OR
C ANY OF THE LETTERS H, I, J IS TYPE COMPLEX
```

```
IMPLICIT LOGICAL(L)
C ANY NAME BEGINNING WITH THE LETTER L IS TYPE
C LOGICAL (VALUE OF .TRUE. OR .FALSE.)
```

```
IMPLICIT CHARACTER*8(C)
C ANY NAME BEGINNING WITH C IS TYPE CHARACTER
C WITH THE LENGTH OF THE CHARACTER ENTITY
C BEING 8
```

```
IMPLICIT REAL(A-H) , (O-Z)
C ANY NAME BEGINNING WITH THE LETTERS A
C THROUGH H OR O THROUGH Z IS TYPE REAL
```

PARAMETER statement

CONVEX FORTRAN supports two types of PARAMETER statements. The first type is the standard FORTRAN PARAMETER statement; the second type provides compatibility with compilers supplied by other vendors.

Standard PARAMETER statement

The PARAMETER statement assigns a symbolic name to a constant to be used within the program unit. It has the following form:

```
PARAMETER (name = exp[, name = exp] ...)
```

where:

name
is a symbolic name.

exp
is a constant or constant expression.

The *name* specified references that constant in other statements in the program unit. You must have previously defined any symbolic constant names that appear in an expression. A constant is named only once in a program, although you can use the symbolic name in subsequent DATA statements or expressions in the same program.

You can use the PARAMETER statement to define an entire format specification; however, it must not appear as a part of a format specification. Also, it cannot be used as part of another constant *except as either the real or imaginary part of a complex constant.*

Examples:

```
PARAMETER (SMITH = 1)
PARAMETER (BOSQUE = 10, HAYS = 100)
COMPLEX LAMAR
PARAMETER (LAMAR = (BOSQUE, HAYS))
```

To determine the data type associated with each constant, use the implicit naming convention or an explicit type-declaration statement before the PARAMETER statement. Constants of data

type `POINTER` cannot appear in `PARAMETER` statements. If the length for a character-type constant is different from the default length, you must specify the length before the first appearance of the symbolic name.

Alternate `PARAMETER` statement

An alternate form of the `PARAMETER` statement is available under the `-vfc` compiler option. This `PARAMETER` statement also assigns a symbolic name to a constant, but its list is not bounded by parentheses, and the form of the constant determines the data type of the variable. The alternate `PARAMETER` statement does not conform to the ANSI standard. It has the following form:

```
PARAMETER p=c [, p=c ]...
```

where `p` is a symbolic name and `c` is a constant, the symbolic name of a constant, or a compile-time constant expression.

The data type is not determined by the implicit or explicit typing of the symbolic name, but by the form of the constant. Once you have defined a symbolic name as a constant, you can use it wherever a constant is allowed. You cannot, however, use the symbolic name of a constant as part of another constant, except as a real or imaginary part of a complex constant.

The symbolic name of a constant assumes the data type of its corresponding expression. Therefore you cannot specify the data type of a parameter in a type-declaration statement, and the initial letter of the constant name does not affect the data type.

Example:

```
PARAMETER BAKER=3, XRAY=45.4, ALPHA=XRAY*BAKER
```

To use the alternate `PARAMETER` statement, you must specify the `-vfc` compiler option. (Refer to Appendix E, "VAX FORTRAN compatibility".)

Caution

Without the `-vfc` option, the compiler will not flag an error when it encounters a single-argument `PARAMETER` statement without parentheses, but will consider the statement an assignment statement. The compiler will flag an error if it encounters a multi-argument `PARAMETER` statement without parenthesis in absence of the `-vfc` option.

Type-declaration statements

Type statements are numeric, logical, CHARACTER, or RECORD type-declarations. Type statements override implicit typing and IMPLICIT statements. You can also use these statements to specify array dimensions.

Both numeric and CHARACTER type-declaration statements can initialize data by including values bounded by slashes (/) in the statement. Place the values after the symbolic name of the variable or array to be initialized. Initial values are assigned in the same way that they are assigned in DATA statements.

The following subsections include examples of type-declaration statements.

Numeric type-declaration statements

Type-declaration statements have the form:

```
type v [/clist/] [, v [/clist/]] . . .
```

where

type

is any data type specifier except CHARACTER or RECORD.

v

is the symbolic name of a constant, variable, array, statement function or function subprogram, or array declarator.

clist

is a list of constants. (Refer to the "Data statement form" section of Chapter 6, "DATA statement.")

*You can follow the symbolic name with a data-type length specifier written as *s, where s is one of the acceptable lengths for the data type being declared. This overrides the length attribute that the statement implies for the specified item. When you use data type-length specifiers with an array declarator, they can be placed immediately after data type or immediately after the array name, as in the following example:*

```
REAL*8 X(100), Y(100)
```

or

```
REAL X*8(100), Y(100)
```

Both of these declarations define *X* as a 100-element array of *REAL*8* values; in the first, however, *Y* is also declared as a *REAL*8* array. In the second, the data-type length specifier is specific to *X*, and provided the default *REAL* size hasn't been changed with an *IMPLICIT* statement or a compiler option, *Y* will be a *REAL*4* (default *REAL*) array.

You can assign initial values to variables or arrays with */clist/*, which initializes the variable or array immediately preceding it. For arrays only, the *clist* can consist of more than one element. If you initialize an array using */clist/*, every element in the array must be assigned a value, as in the following example:

```
REAL*8 PI/7.43562D0/,E/3.33D0/,QARRAY(10)/5*0.0,5*1.0/
```

CHARACTER type-declaration statements

CHARACTER type-declaration statements, like the numeric type, use the *clist* provision, but in the following form:

```
CHARACTER [*len[,]] v[*len][/clist/][,v[*len][/clist/]]...
```

where

v

is the symbolic name of a constant, variable, array, statement function or function subprogram, or array declarator.

len

is an unsigned integer constant, an *INTEGER* constant expression enclosed in parentheses, or an asterisk enclosed in parentheses. The value of *len* specifies the length of the character data elements. When an array is being declared, the length must appear after the array dimension.

clist

is a list of constants, as in the *DATA* statement.

As for numeric type-declaration statements, the */clist/* assigns initial values to the variable or the array immediately preceding it. For arrays only, *clist* can contain more than one element. Where this is the case, it must contain a value for every element in the array.

The following example specifies an array DEMO consisting of fifty 16-character elements, an array DUMMY comprising twenty 9-character elements, and a variable DRAFT, which is 5 characters long *with an initial value of ABCDE*:

```
CHARACTER*16 DEMO(50), DUMMY(20)*9, DRAFT*5 /'ABCDE' /
```

If you do not specify the length of an item (`CHARACTER*len`), its length is *len*—the default length specification for that item. If you specify the length, that length overrides the length specified in `CHARACTER*len`.

If you specify the length as an `*`, for example, `CHARACTER* (*)`, a function name or dummy argument assumes the length specification of the corresponding function reference or actual argument, and a symbolic PARAMETER assumes the length of the actual constant. This is known as an assumed-length character argument.

RECORD type-declaration statements

CONVEX FORTRAN supports the RECORD data type as an extension under the -vfc compiler option. RECORD type declaration statements have the following form:

```
RECORD /structure-name/ record-namelist
```

where:

structure-name

is the name of a previously declared structure.

record-namelist

is a list of variable names, array names, or array declarations, separated by commas.

Structures and records are discussed in detail in the "Structure declaration" section of Appendix E, "VAX FORTRAN compatibility."

POINTER statement

The POINTER statement allows you to declare one or more pointer variables, and to define their pointees. A pointer is a four-byte variable that contains the address of another variable or array, which is referred to as the pointee.

The *POINTER* statement has the following form:

```
POINTER (p, s)
```

where

p

is the pointer being declared; it will hold the address of *s*. *p* must be a variable, and cannot be declared as any other data type. Constants, arrays, statement functions and external functions cannot be used.

s

is the pointee corresponding to *p*. *s* contains the variable whose address is contained in *p*. *s* cannot be a pointer. *s* cannot be associated with any other known piece of named and referenced storage except through assignments to *p* or by associating two or more pointees with one pointer. *s* cannot be of type *CHARACTER*.

Note

Associating two pointees with one pointer can inhibit optimization.

For purposes of arithmetic and data type conversions, the same rules that apply to variables of type *INTEGER*4* apply to pointers. Converting pointers to non-*INTEGER* variables should be avoided. Performing division and multiplication on pointers should also be avoided.

A pointer can be passed into a subprogram as long as it is declared as a pointer in the subprogram.

The *LOC* function

The *LOC* function can be used to assign the address of a variable into a pointer, as shown in the following code segment.

Example:

```
POINTER (IPX, X)  
IPX = LOC (Y)
```

In this example, *X* and *Y* can be used to reference the same value. For more information on the *LOC* function, see the *LOC (3F)* man page.

Dynamic memory allocation

Pointers provide a means by which CONVEX FORTRAN programs can call external routines that dynamically allocate memory. Other more automated dynamic memory allocation methods are available when the `-f90` flag is specified; refer to Appendix C, "Fortran 90 compatibility," for more information.

The following example shows a CONVEX FORTRAN program calling a C routine that dynamically allocates a block of memory of a size specified by the user.

FORTRAN main program:

```
PROGRAM DYNARRAY
REAL A(1)           ! MUST DEFINE NUMBER OF
                   ! DIMENSIONS
POINTER(IPA, A)     ! DECLARE POINTER
INTEGER FALLOC      ! DECLARE C FUNCTION
PRINT*, "Enter array size:"
READ*, ISIZE        ! GET DESIRED ARRAY SIZE
IPA = FALLOC(ISIZE*4) ! ALLOCATE STORAGE
C ISIZE*4 BECAUSE DEFAULT REALS ARE 4 BYTES
.
.
.
```

C routine:

```
int falloc_(int *size) /* trailing underscore
                       required for FORTRAN */
{
    int block;
    block = malloc(*size); /* allocate desired
                           memory w/malloc */
    return block;         /* return starting
                           address */
}
```

The array `A` is declared in `DYNARRAY` as a single-element, one-dimensional array because CONVEX FORTRAN needs to know the number of dimensions of an array, as well as the number of elements in each dimension except the last, at compile time. Since this is a one-dimensional array, the declared length of one element need not be static; it is superseded by the storage allocated when `FALLOC` is called.

For arrays declared with multiple dimensions, only the last dimension can be superseded by dynamically-allocated storage. However, dynamic multidimensional arrays can be created by dynamically allocating the total number of elements required as a one-dimensional array and then passing dimensioning information, along with the array, to a subroutine which accesses the array elements using multidimensional notation. This is illustrated in the following example, in which *IROW* and *JCOL*, the size of each dimension of a two-dimensional array, have been obtained from the user. This example calls *FALLOC*, which is defined in the preceding example.

Example

```
REAL A(1)
POINTER (IPA, A)
INTEGER FALLOC
.
.
.
IPA = FALLOC (4*IROW*JCOL)
CALL MAKEARRAY (A, IROW, JCOL)
.
.
.

SUBROUTINE MAKEARRAY (A, IROW, JCOL)
REAL A (IROW, JCOL)
.
.
.
```

Note

Fortran 90 automatic arrays, accessible under the -F90 option, provide a simpler way to allocate dynamic arrays local to a subroutine. Refer to Appendix C, "Fortran 90 compatibility," for more information on the dynamic memory allocation features supported under the -F90 option.

The POINTER statement is also supported in Cray mode, along with many Cray functions that are designed to allocate and manipulate dynamic memory. Refer to Appendix D, "Cray FORTRAN compatibility," for more information.

For more information on calling non-FORTRAN routines from FORTRAN, refer to the CONVEX Interlanguage Programming Guide.

DIMENSION statement

The DIMENSION statement names arrays and specifies the bounds of the array. It has the following form:

```
DIMENSION arrname ([lb:] ub [, ...]) [, arrname (...)]
```

where:

arrname

is the name of the array you are dimensioning. The array is typed according to a type declaration statement (that can either precede or follow the DIMENSION statement) or implicitly. Multiple arrays can be dimensioned in a single DIMENSION statement. *If the array is to be automatic, arrname must be local to a subroutine.*

lb

is the lower dimension bound. This must be an integer value or parameter and must be smaller than *ub*. *lb* defaults to 1.

ub

is the upper dimension bound. This must be an integer value or parameter and must be greater than *lb* if *lb* exists.

For static arrays, ub must be a constant. For allocatable arrays, ub can be a constant, variable or expression.

If you do not specify a lower bound, the default value is 1, and the upper bound is the number of elements given for the dimension. To use a lower bound that is not 1, you must specify both bounds. The bounds values can be positive, negative, or zero. Separate lower- and upper-bound values with a colon.

For automatic arrays, ub and/or lb for at least one dimension must either consist of or be derived from an integer argument passed to the subroutine in which the automatic array resides.

Different dimensions of a multidimensional automatic array can be declared to be different sizes based on one or more dimension arguments passed to the subroutine.

The type of array and the product of the subscripts determine the number of storage units allocated to each array named in the statement. FORTRAN arrays can have from one to seven dimensions.

Example:

```
DIMENSION MYRAY (4, 5)
```

The preceding example specifies an implicitly typed INTEGER two-dimensional array with 20 elements. The product of the dimension declarators, (4, 5), determines the total number of storage elements assigned to the array. For more information on declaring and dimensioning arrays, refer to Chapter 3, "Arrays and substrings."

ALLOCATABLE **statement**

Allocatable arrays are dynamic arrays of predetermined rank. Storage for allocatable arrays is allocated on the heap at runtime when an ALLOCATE statement is executed, and must be deallocated through use of the DEALLOCATE statement when the array is no longer needed. Refer to Chapter 3, "Arrays and substrings," for more information on the ALLOCATE and DEALLOCATE statements.

The ALLOCATABLE statement is used to declare allocatable arrays. You must supply either the -f90 or -cfc compiler options if you declare allocatable arrays in your program. Allocatable array declarations have the following form:

```
[type arrexp ]  
ALLOCATABLE (arrexp [, ...])
```

where

type

is an optional type definition for the array. The form for this is identical to the form for a standard array type definition, except instead of specifying constant dimension parameters, you must either supply a rank definition or provide no parameters.

arrexp

is the array name or rank definition if it was not given in a preceding type statement. The array name may occur in both the type and ALLOCATABLE statements; the rank definition must occur in exactly one, or the compiler flags an error.

An allocatable array's rank must be supplied either in the ALLOCATABLE statement, in a DIMENSION statement, or in the array's type declaration, which, if used, precedes the ALLOCATABLE statement. Rank definitions are discussed in detail in the "Allocatable array declarations" section of Chapter 3, "Arrays and substrings." For a more detailed discussion of allocatable arrays, refer to Appendix C, "Fortran 90 compatibility."

EQUIVALENCE statement

The EQUIVALENCE statement causes two or more entities within a program unit to refer to the same storage area. Thus, the same storage unit can be referenced by more than one name. Each statement contains two or more variables, array elements, array names, or substring names, separated by a comma. An array must be dimensioned with a DIMENSION, type, or COMMON statement before it or any of its elements can be equivalenced. All elements contained in the same set of parentheses are allotted storage in the same location.

Example:

```
DIMENSION MYARRAY (10), COM (12)
EQUIVALENCE (F,G,H), (MYARRAY(9),COM(10)), (L,M,N)
```

In this example, variables F, G, and H share the same location; the 9th element in array MYARRAY and the 10th element in array COM share the same location; the variables L, M, and N share the same location.

If different data types are equivalenced, the EQUIVALENCE statement does not imply mathematical equivalence or type conversion. Type is associated with the name used to reference a location; the name determines how data is stored or read from the location. Names of dummy arguments of an external procedure in a subprogram, or a variable name that is a function name cannot appear in an EQUIVALENCE statement.

Note

Use of the EQUIVALENCE statement can interfere with program optimization.

Equivalencing arrays

Making one array element equivalent to an element of another array also defines the relative locations of the other array elements.

Example:

```
DIMENSION A(5), B(4,3)
EQUIVALENCE (A(2),B(2,2))
```

The entire array A shares part of the storage space allotted to array B. The EQUIVALENCE statements:

```
EQUIVALENCE (A,B(1,2))
```

or

```
EQUIVALENCE (A(5),B(1,3))
```

also align the two arrays in the same manner as `EQUIVALENCE (A(2),B(2,2))`. Table 8 shows how these statements align the arrays.

Table 8
Array locations

Array A		Array B	
Elements	Location number	Elements	Location number
B(1,1)	1		
B(2,1)	2		
B(3,1)	3		
B(4,1)	4		
B(1,2)	5	A(1)	1
B(2,2)	6	A(2)	2
B(3,2)	7	A(3)	3
B(4,2)	8	A(4)	4
B(1,3)	9	A(5)	5
B(2,3)	10		
B(3,3)	11		

Two or more elements of the same array cannot share the same storage location. In the following example, the statements are invalid because they allocate the same storage for `A(1)` and `A(3)`.

Example:

```
DIMENSION A(5), B(4,3)
EQUIVALENCE (A(1),B(3,3)), (A(3),B(3,3))
```

When making arrays equivalent for storage, you can identify an array element with a single subscript (the linear element number), even if the array was defined as a multidimensional array.

Example:

```
REAL A(10,10), Z(200)
EQUIVALENCE (A(100), Z(150))
```

associates element A(10, 10) with element Z(150).

Equivalencing substrings

When making character substrings equivalent for storage, the EQUIVALENCE statement also defines storage locations for the other corresponding characters in the strings.

Example:

```
CHARACTER PROD*12, N*8
EQUIVALENCE (PROD(6:10), N(4:8))
```

specifies that PROD(6) and N(4), PROD(7) and N(5), and so on through PROD(10) and N(8) share storage locations. Similarly, N(1) now shares storage with PROD(3), N(2) with PROD(4), and so forth.

Equivalencing two or more character substrings that begin at different character positions within the same character variable or array is prohibited. You cannot use EQUIVALENCE statements to indicate that contiguous storage units are to be noncontiguous.

Using EQUIVALENCE in common blocks

You can extend a common block of storage with the EQUIVALENCE statement if you extend locations beyond the last element and do not add to the beginning of the common block.

Example:

```
DIMENSION A(5), B(2,3)
COMMON B
EQUIVALENCE (A(1), B(1,2))
```

The preceding example extends the common block beyond the last element. The existing common block includes B(1, 1) through B(2, 3) and A(1) through A(4); A(5) is the extended portion being added beyond the last element B(2, 3) of the existing common block.

If you change COMMON B in the previous example to COMMON A:

```
DIMENSION A(5), B(2,3)
COMMON A
EQUIVALENCE (A(1), B(1,2))
```

the extension is invalid. The common block now includes A(1) through A(4), and B(1,2) through B(2,3); B(1,1) and B(2,1) comprise the extended portion preceding the common block, which is invalid.

NAMELIST **statement**

The NAMELIST statement associates a single unique name with a list of variables or array names. This name defines a list of entities that can be modified (read) or transferred (written). Thus, you can use this unique name in namelist-directed I/O statements in place of the entities list. The NAMELIST statement has the form:

```
NAMELIST /nlgrpname/varlist [ [, ]/nlgrpname/varlist ] . . .
```

where

nlgrpname

is a symbolic name representing the list of entities to be read or written.

varlist

is a list of variable or array names (separated by commas) to be associated with the nlgrpname. A variable or array name can occur in more than one varlist. An entity can be a dummy argument. These entities can be typed explicitly or implicitly to any data type.

An entity can be of type INTEGER, REAL, LOGICAL, COMPLEX, or CHARACTER. If the entity and the constant value assigned to it are not of the same type, the compiler performs the arithmetic assignment conversion. You cannot, however, convert between numeric and character data types.

The following example shows the format for namelist-directed input:

```
$CONTROL
TESTCASE='40004.00',
CONDITION=.FALSE.,
BEGIN=100,
REPEAT=10,
$END
```

The following statement illustrates the use of `NAMELIST`:

```
NAMELIST /EXAM1/ TESTA,TESTB,TESTC /EXAM2/ TOTTEST
```

In this example, the `NAMELIST` statement defines two group names—`EXAM1` and `EXAM2`. The first represents three entities (`TESTA`, `TESTB`, and `TESTC`), while the second represents one entity (`TOTTEST`). The order in which you list the entities in the `varlist` determines the order in which the values are output; however, the order of input values is immaterial. Also, you do not need to define every entity in the corresponding `varlist` during input. For instance, using the previous example, you could input values for only `TESTA` and `TESTB`. The value of `TESTC` would remain unchanged.

Although you cannot use array elements and character substrings in a namelist, you can use namelist-directed I/O to assign values to elements of arrays or substrings of character variables that occur in the namelist. You can also use a variable or an array name in several namelists. (Refer to Chapter 9 and Chapter 10 for more information on namelist-directed I/O.)

EXTERNAL statement

An `EXTERNAL` statement identifies a symbolic name as representing an externally-defined procedure or dummy procedure. It indicates that a given name is the name of a subprogram instead of a variable or array name. An `EXTERNAL` statement must be used for a subprogram or dummy procedure name that appears as an actual argument in a function reference or in a `CALL` statement. The `EXTERNAL` statement has the following form:

```
EXTERNAL n [,n]...
```

where *n* is the symbolic name of a user-supplied subprogram, block data subprogram, or dummy procedure.

If an `EXTERNAL` statement declares an intrinsic name as an external procedure, all references to the intrinsic name are treated as references to an external procedure rather than as calls to an intrinsic function. For example, if you declare `COS` in the `EXTERNAL` statement (`EXTERNAL COS`), all subsequent references are to an external procedure `COS`, not the intrinsic function `COS`.

INTRINSIC statement

The `INTRINSIC` statement permits a name that specifies an intrinsic function to be used as an actual argument. The `INTRINSIC` statement has the following form:

```
INTRINSIC n [,n] . . .
```

where *n* is one of the FORTRAN intrinsic functions.

If the name of an intrinsic function is to be used as an actual argument in a program unit, it must appear in an `INTRINSIC` statement in that program unit. Not all intrinsic functions can be used as actual arguments. Intrinsic functions for type conversion, maximum and minimum value, and lexical comparison functions (for example, `INT`, `IFIX`, `IDINT`, `REAL`, `FLOAT`, `SINGL`, `DBLE`, `CMPLX`, `ICHAR`, `CHAR`, `LGE`, `LGT`, `LLE`, `LLT`, `MAX`, `MAX0`, `AMAX1`, `AMAX0`, `MAX1`, `MIN`, `MIN0`, `AMIN1`, `DMIN1`, `AMIN0`, and `MIN1`) cannot be used as actual arguments.

SAVE statement

The `SAVE` statement retains the values of designated variables and arrays in a subroutine or function when a `RETURN` or `END` statement is executed. Thus, items specified in a `SAVE` statement do not become undefined when the subroutine or function completes execution. In the next call to the subroutine or function, a saved item has the same value it had on return from the preceding call.

The `SAVE` statement has the following form:

```
SAVE [n [,n] . . . ]
```

where *n* is a variable name, statically-sized array name, or a named common block contained between slashes (for example, `/NCOM/`). Dummy argument names, subprogram or function names, automatic or allocatable array names, or common block entity names are not allowed. When a common block name appears in a `SAVE` statement, all the variables and arrays in the common block are saved.

If the *SAVE* statement contains no arguments, the values of all allowable entities are retained.

DATA statement

The DATA statement establishes initial values for arrays, array elements, substrings, and variables. Values are initialized when the program unit is compiled and can be changed during program execution.

The DATA statement is nonexecutable. All entities initialized with a DATA statement are defined when program execution begins; all entities not initialized with a DATA statement are undefined when program execution begins. Undefined entities must be defined before they can be referenced.

DATA statement form

The DATA statement has the following form:

```
DATA nlist/clist / [ [, ]nlist/clist / ] ...
```

where

nlist

is a list of one or more array names, array element names, character substring names, implied-DO lists or variable names. Dummy argument names or function names cannot appear in the *nlist*.

clist

is a list of constants (numeric, character, logical, or *Hollerith*) or symbolic names of constants (defined with a PARAMETER statement). Items in *clist* are consecutively assigned to the entities in *nlist*; the first item in *nlist* corresponds to the first item in *clist*. Constants can be repeatedly associated with entities in the *nlist*.

The number of names in the *nlist* must equal the number of constants in the *clist*. The following statement is invalid because two values are associated with one variable.

```
DATA MYVAR/5,9/
```

You can repeat the same value for more than one element by placing a nonzero, unsigned integer constant indicating the number of repetitions and an asterisk (*) before the value.

Example:

```
DATA C,LOW,MYEX(2),MYEX(3)/"NAME",.FALSE.,2*3/
```

This statement initializes a character value of NAME for C, logical value .FALSE. for LOW, and 3 for MYEXAMPLE (2) and MYEXAMPLE (3). As long as you retain the correct name and value association, the order and grouping is immaterial.

The previous example can also be represented as:

```
DATA C/"NAME"/,LOW/.FALSE./,MYEX(2),MYEX(3)/2*3/
```

When a character entity is longer than its corresponding character constant, blanks are added on the right. If a character entity is shorter, extra characters on the right are ignored. For example, the following statements initialize DUR to TEMPORARY^^^ and LOC to HOLDING^CELL.

Example:

```
CHARACTER*12 DUR,LOC  
DATA DUR,LOC/"TEMPORARY","HOLDING^CELL"/
```

The character entity is the same length as the constant in LOC, so no blanks are added or ignored. Three blanks are automatically added to DUR, however, because the character entity is longer than its corresponding character constant.

You can use a DATA statement to initialize all or part of an array.

Example:

```
DIMENSION MYRAY(4,3)  
DATA MYRAY /12*5/
```

The above example indicates that the value of all MYRAY elements are initialized to 5. Elements of the array are initialized in the order of subscript progression.

Implied-DO

Implied-DO lists can occur in DATA statements in the form:

$$(dlist, i=m_1, m_2, m_3)$$

where

dlist

is a list of array element names and implied-DO lists.

i

is the name of an integer variable termed the implied-DO variable.

m₁, m₂, m₃

are integer constant expressions that can contain implied-DO variables of other implied-DO lists. The constants specify the initial value, terminal value, and increment, respectively, for the integer variable. If you omit the comma and the value of *m₃*, the increment value defaults to 1. The increment count must be positive.

Examples using implied-DO

The statements:

```
REAL C(8), D(12)
DATA E, (C(I), I=2, 6, 2), F, (D(J), J=1, 3)/4*0, 4*1/
```

initialize E, C(2), C(4), C(6) to 0.0 and F, D(1), D(2), D(3) to 1.0.

The statements:

```
DIMENSION B(10, 10)
DATA ( (B(I, J), I=1, 5), J=1, 5)/25*2/
```

initialize the 25 elements of the submatrix to 2.0; the submatrix is located at the upper-left corner of B.

The following statements initialize 10 elements of the matrix B to 5:

```
INTEGER B(4,4)
DATA ( (B(I,J),J=1,I),I=1,4)/10*5/
```

The matrix has elements B(1,1...1,4), B(2,2...2,4), B(3,3), B(3,4) and B(4,4).

DATA statement extensions

If, in the DATA statement, the constant value in `clist` and the entity in `rlist` have numeric data types, you can determine the data-type conversion in addition to the standard as follows:

- *If an octal or hexadecimal constant is assigned to a variable or array element, the data type of the variable or array element determines the number of digits that can be assigned. Where the constant has fewer digits than the variable or array element, the constant is extended on the left with zeros. If the constant has more digits than can be stored, the constant is truncated on the left.*
- *If a Hollerith or character constant is assigned to a numeric variable or numeric array element, the number of characters that can be assigned depends on the data type of the variable or array element. Where the Hollerith or character constant has fewer characters than the variables or array element, spaces are added to the right of the constant. If the constant has more characters than can be stored, excess characters on the right of the constant are eliminated.*

The constant value in `clist` can be of the numeric data type and the entity in `rlist` of the CHARACTER data type. When this is the case, the entity must conform to the following:

- *The character entity must have a length of one character.*
- *The constant must be an integer, octal, or hexadecimal constant and must have a value in the range 0 through 255.*

Following these restrictions permits the entity to be initialized with the character that has the ASCII code specified by the constant, which, in turn, allows a character entity to be initialized to any 8-bit ASCII code.

The next example initializes the real array *R* to all zeros, the real variable *PI*, the character variable *C*4* to "TEST", and the character variable *CR*1* to the ASCII character code 'OD'X.

Example:

```
DATA R, PI /20*0.0, 3.14159265/  
DATA C, CR /7HTESTING, 'D'X/
```

Arrays and variables in COMMON can be initialized in DATA statements in program units other than BLOCK DATA subprograms. Each variable or array element can only be initialized once.

An assignment statement evaluates an expression and assigns the value to a variable, a substring, or an array element. The ASSIGN statement, which is discussed later in this chapter, assigns a statement label value to a variable.

The assignment statement has the form:

$$v = ex$$

where v is a variable, array element, or substring, and ex is an expression.

If the type of the variable on the left side of the equal sign is the same as that of the expression on the right, the value is assigned directly. If the data types differ, the value of the expression is converted to the type of the left side entity before the value is assigned. For example, in the statement

$$I = (L + M) / K \quad !\text{where } L = 9, M = 6, K = 3$$

both the variable I and the expression $(L+M)/K$ are of type INTEGER, so the value 5 is assigned directly. If the variable is INTEGER and the expression is not, the expression is converted to INTEGER and assigned to the variable. Thus, the statement

$$I = (R + S) / T \quad !\text{where } R = 8.0, S = 9.0, T = 3.0$$

assigns the INTEGER value of 5 ($17/3$ truncated) to I . In this example, $(R+S)/T$ is truncated and converted to INTEGER.

Character conversions

The character assignment statement has the following form:

$$v = ce$$

where v is a character variable, array element, or substring, and ce is a character expression.

The assigned entity and the expression can have different lengths. If the entity (v) is of greater length than the length of the character expression (ce), CONVEX FORTRAN inserts blanks after the characters until the length is equal to v . If the length of v is less than the length of ce , extra expression characters on the right are truncated until v and ce are equal in length. For example, the following statements assign CONVEX^^^^^^ to NAME and supercompute (12 characters only) to PROD.

Example:

```
CHARACTER*12 NAME, PROD
NAME = 'CONVEX'
PROD = 'supercomputer'
```

The same character positions defined in v cannot be referenced in ce within the same statement. When you assign a value to a character substring, the character positions in the character variable or array element not included in the substring are not affected. If a value was previously assigned, the value remains the same; if the value was undefined, it remains undefined. Using a differing substring within the same variable, such as $A1(1:3) = A1(4:6)$, is acceptable.

Fortran 90 array assignments

CONVEX FORTRAN supports both the use of array-valued expressions in assignment statements and masked array assignments (*WHERE* statements and constructs) under the *-f90* compiler option. Both these methods of manipulating array data are discussed in this section.

Array-valued expressions

The ability to use array-valued expressions in assignment statements is supported by CONVEX FORTRAN under the `-f90` flag. This feature allows you to assign a value or expression to an entire array or array section with one assignment, using the following form:

$$arr = ex$$

where *arr* is the name of an array or an array section description and *ex* is an expression.

As with assignment to a scalar variable, mismatched expression types are converted to the type of the argument on the left before assignment.

Example:

```
INTEGER IX(10)
REAL X
      .
      .
      .
IX = IX * X
```

Here, each element of *IX* is multiplied by the *REAL* variable *X* and truncated to an *INTEGER*. The result then replaces the original element in the array.

Note

Fortran 90 array assignments are translated into loops by the compiler. Any optimization options specified at compile time are then applied to the generated loop; for instance, the loop would be vectorized at optimization options `-o2` and higher.

Array sections can be similarly assigned. Refer to the section "Array sections" in Appendix C for more information.

For more general information on the Fortran 90 features available in CONVEX FORTRAN, refer to Appendix C, "Fortran 90 compatibility."

Masked array assignment

When the `-f90` compiler option is specified, CONVEX FORTRAN allows assignment of values to an array under a mask specified via the *WHERE* statement or *WHERE* construct. The *WHERE* statement evaluates a logical expression to determine to

which elements the assignment is being applied. The *WHERE* construct works similarly, but is terminated with the *ENDWHERE* statement. It can contain several assignments and an *ELSEWHERE* statement, which allows alternate assignments to be applied for the complement of *mask-expression*.

The *WHERE* statement has the following form:

```
WHERE (mask-expr) assignment-stmt
```

The *WHERE* construct has the following form:

```
WHERE (mask-expr)  
  [assignment-stmt]  
  [...]  
  [ELSEWHERE  
    [assignment-stmt]  
    [...]]  
ENDWHERE
```

where

mask-expr

is a logical expression of the same shape as the array(s) being manipulated in the *assignment-stmt*(s).

assignment-stmt

is an array assignment. The array must be the same shape as the array in *mask-expr*

On execution, the *mask-expr* is evaluated. Any following assignments are executed only on array elements corresponding to those elements for which *mask-expr* evaluated to *.TRUE.* If an *ELSEWHERE* statement is present, its assignments are applied to array elements corresponding to those elements for which *mask-expr* evaluated to *.FALSE.*

Example 1:

```
REAL DATA(1000), OFFSET(10000), ADJUSTED(1000), LMT  
LOGICAL NORMAL(1000)  
LMT = 130.0  
.  
.  
.  
WHERE(DATA .LE. LMT) NORMAL = .TRUE.
```

In this example, all elements of the logical array `NORMAL` that have the same index as elements in the real array `DATA` that are less than or equal to `LMT` are set to `.TRUE`.

The following example assumes the same variable declarations as Example 1.

Example 2:

```
WHERE (DATA .GT. LMT)
  ADJUSTED = FIX (DATA)
  NORMAL = .FALSE.
ELSEWHERE
  ADJUSTED = 0.0
  NORMAL = .TRUE.
ENDWHERE
```

In this example, elements of `ADJUSTED` corresponding to elements of `DATA` greater than 130.0 are set to `FIX (DATA)`, and all other elements of `ADJUSTED` are set to 0.0. Similarly, all elements of `NORMAL` corresponding to `DATA` elements greater than 130.0 are set to `.FALSE.` and all other elements of `NORMAL` are set to `.TRUE.`

Array sections and subscript expressions can be used with the `WHERE` construct to change the correspondence of elements between arrays.

Example 3:

```
WHERE (DATA .LE. LMT)
  OFFSET (ISET:ISET+1000) = DATA
ELSEWHERE
  OFFSET (ISET:ISET+1000) = 0.0
ENDWHERE
```

In this example, elements of `OFFSET` starting at the value `ISET` and going through `ISET+1000` are set to correspond to elements of `DATA` where the `DATA` elements are less than or equal to `LMT`, and are set to 0.0 where the `DATA` elements are greater than `LMT`.

Note

It is possible to assign values to the same array element in both the `.TRUE.` and `.FALSE.` branches of the `WHERE` construct when using array sections or subscript expressions that differ in each branch of the construct. This action can inhibit vectorization and parallelization of the `WHERE` construct, and should therefore be avoided.

Remember, the *WHERE* construct differs from the block-*IF* in that both clauses can and often do execute. It is therefore possible for assignments that execute in the *ELSEWHERE* clause of the *WHERE* construct to change values assigned in the *WHERE* clause when array section notation is used. The following example, which illustrates this, assumes the same variable declarations as Example 1.

Example 4:

```
WHERE (DATA .LE. LIMIT)
  OFFSET (ISET:ISET+999:2) = DATA(1:1000:2)
ELSEWHERE
  OFFSET (ISET+1:ISET+1000:2) =
  OFFSET (ISET:ISET+999:2)
ENDWHERE
```

In this example, every other element from *OFFSET (ISET)* through *OFFSET (ISET+999)* is set in the *WHERE* clause. Then in the *ELSEWHERE* clause, because of the sectioning notation used, each value assigned in the *WHERE* clause is copied into the element immediately following it in *OFFSET*. The end result is that *OFFSET (ISET) = OFFSET (ISET+1) = DATA (1)*, *OFFSET (ISET+2) = OFFSET (ISET+3) = DATA (3)*, and so on for every pair of elements of *OFFSET* starting at *ISET*.

Data conversion rules

Table 9 summarizes the data conversion rules for assignment statements. These rules apply for both arrays and scalar variables of the indicated type.

Table 9
Conversion of expressions

Type of variable (V)	Type of expression (E)	Value assigned
INTEGER/ LOGICAL	INTEGER/ LOGICAL	Direct assignment of E to V.
	REAL	Truncate E to INTEGER and assign to V.
	REAL*8	<i>Truncate E to INTEGER and assign to V.</i>
	REAL*16	<i>Truncate E to INTEGER and assign to V.</i>
	COMPLEX	Truncate real part of E to INTEGER and assign to V. Imaginary part not used.
	COMPLEX*16	<i>Truncate real part of E to INTEGER and assign to V. Imaginary part not used.</i>
REAL	INTEGER/ LOGICAL	Convert to REAL and assign to V.
	REAL	Direct assignment of E to V.
	REAL*8	<i>Assign most significant digits of E to V; least significant digits of E rounded.</i>
	REAL*16	<i>Assign most significant digits of E to V; least significant digits of E rounded.</i>
	COMPLEX	Assign real part of E to V. Imaginary part not used.
	COMPLEX*16	<i>Assign most significant digits of real part of E to V; least significant digits rounded. Imaginary part not used.</i>
REAL*8 (DOUBLE PRECISION)	INTEGER/ LOGICAL	Convert to REAL and assign to V.
	REAL	Assign E to most significant portion of V. Least significant portion of V is assigned 0.
	REAL*8	<i>Direct assignment of E to V.</i>
	REAL*16	<i>Assign most significant digits of E to V; least significant digits of E rounded.</i>
	COMPLEX	Assign real part of E to most significant portion of V; assign 0 to least significant portion of V. Imaginary part not used.
	COMPLEX*16	<i>Assign real part of E to V. Imaginary part not used.</i>

Table 9 (continued)
Conversion of expressions

Type of variable (V)	Type of expression (E)	Value assigned
REAL*16	INTEGER/ LOGICAL REAL REAL*8 REAL*16 COMPLEX COMPLEX*16	<p>Convert to REAL and assign to V.</p> <p>Assign E to most significant portion of V. Least significant portion of V is assigned 0.</p> <p>Assign E to most significant portion of V. Least significant portion of V is assigned 0.</p> <p>Direct assignment of E to V.</p> <p>Assign real part of E to most significant of V; assign 0 to least significant of V. Imaginary part not used.</p> <p>Assign real part of E to most significant of V; assign 0 to least significant of V. Imaginary part not used.</p>
COMPLEX	INTEGER/ LOGICAL REAL REAL*8 REAL*16 COMPLEX COMPLEX*16	<p>Convert to REAL and assign to real part of V. Assign 0.0 to imaginary part of V.</p> <p>Assign E to real part of V. Assign 0.0 to imaginary part of V.</p> <p>Assign most significant digits of E to real part of V; least significant portion of E is rounded. Assign 0.0 to imaginary part of V.</p> <p>Assign most significant digits of E to real part of V; least significant portion of E is rounded. Assign 0.0 to imaginary part of V.</p> <p>Direct assignment of E to V.</p> <p>Assign most significant portion of E to real part of V; least significant portion of real part of E is rounded. Assign most significant imaginary part of E to imaginary part of V; least significant portion of imaginary E is rounded.</p>
COMPLEX*16	INTEGER/ LOGICAL REAL REAL*8 REAL*16 COMPLEX COMPLEX*16	<p>Convert to REAL and assign to V. Assign 0.0 to imaginary part of V.</p> <p>Assign E to most significant portion of real part of V. Assign 0.0 to imaginary part of V.</p> <p>Assign E to real part of V. Assign 0.0 to imaginary part of V.</p> <p>Assign most significant digits of E to real part of V; least significant portion of E is rounded. Assign 0.0 to imaginary part of V.</p> <p>Assign real part of E to most significant of real V; assign 0 to least significant portion of real part. Assign imaginary part of E to most significant of imaginary V; assign 0 to least significant portion of imaginary part.</p> <p>Direct assignment of E to V.</p>
RECORD	RECORD	Must be a RECORD of the same type.

ASSIGN statement

The ASSIGN statement allows you to assign a statement label to an INTEGER variable. It has the following form:

```
ASSIGN s TO i
```

where

s

is the label of an executable statement or FORMAT statement in the current program unit and

i

is an INTEGER variable name.

Execution of an ASSIGN statement causes the statement label (number) to be assigned to the integer variable (*i*). The variable is now defined for use only as a statement label reference; it is undefined as an integer variable. This statement label is required for referencing in an assigned GOTO statement or as a format identifier in an input/output statement. The variable cannot be referenced in any other way while defined with a statement label value.

Example:

```
ASSIGN 75 TO M
      .
      .
      .
GOTO M
```

The above example transfers control to a statement with the label 75. Do not use the ASSIGN statement for arithmetic purposes. For example, ASSIGN 75 TO M is not equivalent to $M = 75$. Likewise, you an arithmetically assigned variable cannot be used as a statement label. If you define the integer variable with a statement label, you can redefine it with the same or a different statement label value or integer value. For example, the statement

```
M = 50
```

returns M to the status of an integer variable; it cannot be used in a GOTO statement.

CONVEY FORTRAN usually executes statements in the order in which they are written. Control statements provide a means of altering the normal program execution sequence. Control statements include:

- GOTO—Unconditional, computed and assigned
- IF—Arithmetic, logical and block, including ELSE IF, ELSE and END IF
- DO—Indexed DO and DO WHILE
- END DO
- CONTINUE
- CALL
- RETURN
- STOP
- PAUSE
- END

GOTO statements

You can use GOTO statements to change the flow of a program by transferring control to a specified statement. The three types of GOTO statements are unconditional GOTO, computed GOTO, and assigned GOTO.

Unconditional GOTO statement

The unconditional GOTO statement has the form:

`GOTO sl`

where *sl* is the label of an executable statement that appears within the current program unit.

During statement execution, control transfers to the statement identified by the statement label. The identified statement then executes.

Example:

```
                GOTO 50
20      A = 5 * D
        .
        .
        .
50      T = T + 1
```

In this example, control is transferred to statement 50. To execute statement 20 and those statements immediately following it, control must transfer at some point to statement 20.

Computed GOTO statement

The computed GOTO statement specifies the next executable statement from a list of several statements. The computed GOTO statement has the following form:

```
GOTO (slist) [, ]e
```

where

slist

is a list of labels of executable statements within the current program unit separated by commas (*l₁, l₂, ...*).

e

is an arithmetic expression.

The computed GOTO statement evaluates the expression and transfers control to the labeled statement whose position in *slist* corresponds to *e*. If the value of *e* is less than 1 or greater than the number of statement labels in *slist*, control passes to the next statement in the program unit. For example, the statement:

```
GOTO (10,15,20,15,30) I
```

transfers control based on the value of variable I. If I is 2, control passes to statement 15; if I is 5, control goes to statement 30. If the value of I is less than 1 or greater than 5, control passes to the first executable statement immediately following the computed GOTO.

As a CONVEX extension, the arithmetic expression (e) is converted to an integer data type when necessary.

Example:

```
GOTO (10,15,20,15,30)X
```

Here X is converted to type INTEGER before it is compared to (10, 15, 20, 15, 30) to determine the branch destination.

Assigned GOTO statement

The assigned GOTO statement has the following form:

```
GOTO v [ [, ] (slist) ]
```

where

v

is an integer variable name defined by an ASSIGN statement with the value of an executable statement label.

slist

is a list of one or more executable statement labels separated by commas (l_1, l_2, \dots) within the program unit.

When an assigned GOTO is executed, *v* must have the value of a label attached to some executable statement in the same program unit. This label should not be attached to a statement that exists within a block IF statement or DO statement block. If several statement label values are present for *slist*, the label assigned to *v* must be a member of that list.

Example:

```
          ASSIGN 30 TO IFUN
10      GOTO IFUN (20,30,50)
          .
          .
          .
30      FUN = 25.0 * 4.0
```

Statement 30 is executed immediately after statement 10. Control transfers to the statement label last assigned to v by the execution of a prior ASSIGN statement.

IF statements

During program execution, IF statements permit transfer of control based on the value of an arithmetic or logical expression. The three types of IF statements are arithmetic, logical, and block.

Arithmetic IF statement

The arithmetic IF statement evaluates an expression and transfers control based on the value of the expression. The arithmetic IF statement has the following form:

$$\text{IF } (e) \ l_1, l_2, l_3$$

where

e
is an arithmetic expression.

l_1, l_2, l_3
are labels of executable statements contained within the current program unit. All three labels must be included. Do not use labels of statements within DO statement blocks or IF statement blocks.

During program execution, the arithmetic expression is evaluated. If the value of the expression is less than zero, control transfers to the statement with the label l_1 . If the value is equal to zero, control transfers to the statement with the label l_2 . If the value is greater than zero, control transfers to the statement with the label l_3 . For example, the following statement:

$$\text{IF } (IA*IB) \ 40, \ 20, \ 50$$

transfers control to the statement with label 40 if the product of IA and IB is less than zero; to statement 20 if the product is zero; to statement 50 if the product is greater than zero. You can repeat the same statement label in the arithmetic IF statement.

For example, the following statement transfers control to statement 200 if MYEXAM is zero or greater than zero; if MYEXAM is less than zero, control transfers to statement 100.

Example:

```
IF (MYEXAM) 100, 200, 200
```

Logical IF statement

The logical IF statement has the form:

```
IF (le) es
```

where

le

is a logical expression.

es

is an executable statement other than a DO, *END DO*, *END*, logical IF, or block IF statement. Do not use the statement to transfer control to any executable statement within a block IF statement or DO statement block.

The logical IF statement evaluates the value of the logical expression. If the value is true, the statement (*es*) is executed. After the statement is executed, control transfers to the next executable statement unless control is directed elsewhere by the statement (*es*). If the value evaluates to false, the next executable statement is executed.

Consider the following example.

Example:

```
IF (Y .AND. Q) Z = 7  
IF (Y .LT. Q) GOTO 50  
IF (Y .LE. Q) CALL SUB1
```

In the first statement, if Y and Q are true, the value of Z is replaced by 7; otherwise, the value of Z remains unchanged. In the second statement, if the value of Y is less than Q, control transfers to the executable statement at 50; if Y is greater than Q, control transfers to the next executable statement. In the third statement, if the value of Y is less than or equal to Q, the subroutine SUB1 is called. If Y is greater than Q, control passes to the next executable statement, and SUB1 is not called.

Block IF statement

The block IF statement permits one statement or a block of statements to be executed depending on the value of the logical expression. The block begins with an IF THEN statement, followed by the statement block, and ends with an END IF statement. The ELSE statement and the ELSE IF THEN statement can be included in the block IF statement.

There are several variations of the block IF statement. The basic form is

```
IF (le) THEN
.
.
.
END IF
```

Where *le* is a logical expression. If *le* is true, all the lines (the block) between IF THEN and END IF are executed sequentially; otherwise, control transfers to the first executable statement following END IF. You can include one or more statements in the block.

Consider the following example:

Example:

```
IF (H .LE. 40) THEN
    P = H * PR
END IF
```

If H is less than or equal to 40, $P = H * PR$ is executed. After execution of the block, control transfers to the first executable statement following END IF.

Another variation of the block IF statement has the following form:

```
IF (le) THEN
    .
    .           !EXECUTABLE STATEMENTS FOR TRUE VALUE
    .
ELSE
    .
    .           !EXECUTABLE STATEMENTS FOR FALSE VALUE
    .
END IF
```

If the logical expression (*le*) is true, the first block of statements is executed and the block following the ELSE statement is ignored. Control then passes to the next executable statement by means of the END IF statement. However, if *le* is false, the IF THEN block is skipped and control passes to the ELSE statement. Thus, the ELSE statement (block) is executed only if the logical expression (*le*) of the IF statement is false. For example, in the block

```
IF (H .LE. 40) THEN
    P = H * PR
ELSE
    O = (H - 40) * PR * 1.5
    P = H * PR + O
END IF
```

control transfers to the ELSE statement if H is greater than 40. The ELSE block executes and, unless the ELSE transfers control out of the block, control passes to the END IF, which transfers control to the next executable statement. However, if H is less than or equal to 40, the IF THEN block executes; the ELSE block is skipped.

A more complex block IF statement has the form:

```
IF (le1) THEN
  .
  .
ELSE IF (le2) THEN
  .
  .
ELSE IF (len) THEN
  .
  .
ELSE                               !OPTIONAL
  .
  .
END IF
```

This IF block allows any number of additional logical expressions (*le*) to be specified. If the value of the first logical expression (*le*₁) is false, the ELSE IF expressions are evaluated until the value of the expression is true. If *le*₂ through *le*_{*n*} are false, those block sequences are skipped and control transfers to the optional ending ELSE statement, or, if none is present, to the END IF statement. The next example contains an IF block with ELSE IF THEN and ELSE statements.

Example:

```
IF (N .LE. J) THEN
  K = M
ELSE IF (N .GT. J/3) THEN
  K = I
ELSE IF (N .EQ. J/3) THEN
  K = -M
ELSE
  K = L
END IF
```

The IF block is evaluated sequentially. Evaluation of each ELSE IF THEN statement continues until a true value is determined. Then the statements associated with that ELSE IF THEN statement execute and control transfers to the END IF statement, which transfers control to the next executable statement. If all ELSE IF THEN statements evaluate to false, the ELSE statement block, if there is one, executes.

Nested block IF statements

The initial block IF can contain nested block IF statements as long as the nested block IF is completely contained within a statement block. Each block begins with an IF THEN statement and ends with an END IF statement.

Example:

```

IF (T .GT. 40) THEN
  Y = X * 1.5
  IF (AT .GT. 60) THEN
    B = 25
  ELSE
    B = 0
  END IF
ELSE
  NP = H * P
END IF

```

If T is greater than 40, the block executes ($Y=X*1.5$). Then the nested IF THEN is evaluated and executed according to the value of AT. If AT is greater than 60, the block executes. If AT is less than or equal to 60, control transfers to the ELSE statement. (The nested block IF must have an END IF.) If T is less than or equal to 40, the nested IF block does not execute. Control transfers to the outer ELSE statement.

Short-circuit evaluation of conditionals

Short-circuiting the evaluation of conditionals increases the efficiency of IF statements by skipping irrelevant tests when logical operators are involved in the conditional. CONVEX FORTRAN short-circuits evaluation of IF statements that contain .AND. and .OR. operators which have logical operands and are used in a logical context. Take, for example, the following IF statement:

```
IF ((A .EQ. B) .OR. F(G)) THEN
```

Assuming $A = B$, the compiler evaluates $(A .EQ. B)$ and once it has determined that this condition is true, skips the evaluation of $F(G)$ and evaluates the THEN portion of the statement. Similarly, given the code

```
IF ((A .EQ. B) .AND. F(G)) THEN
```

if (A .EQ. B) evaluates to false, the compiler skips the evaluation of F(G) and proceeds past the THEN portion of the statement.

Short-circuit evaluation works with all types of IF statements (arithmetic, logical, and block). Performing arithmetic (+, -, *, /) on or applying non-logical operands or functions to a logical expression disables short circuit evaluation within that expression. Logical valued expressions used as arguments to function calls within an IF statement's conditional expression are not short-circuited. Note that the binary operators .EQ., .NE., .LE., .GT., and .GE. always produce a logical result.

The compiler short-circuits evaluation of conditionals by default. You can disable short circuiting by specifying the `-nosc` option on the compiler command line.

DO statement

A DO statement specifies a DO loop. A group of statements must follow the DO statement, be located within the program unit, and end with a terminal statement. You cannot transfer control into the range of a DO loop from elsewhere in the program unit, but you can terminate execution of a DO loop by transferring control outside the loop. The DO statement has the following form:

```
DO [sl [, ] ] v = ex1, ex2[, ex3]
```

where

sl

is the label of an executable statement (followed by an optional comma) in the current program unit. *If you do not specify a label, the DO loop must end with an END DO statement. That is, a statement of the form DO I=1, 100, 2 must end with END DO.* Nested DO loops cannot share an unlabeled END DO statement, but they can share a labeled terminal statement.

The label identifies the last statement (terminal statement) of the DO loop and the label must textually follow the DO statement. You cannot use an unconditional GOTO, assigned GOTO, arithmetic IF, block IF, ELSE IF, ELSE, END IF, RETURN, END, or DO statement as the terminal statement.

v

is the name of the integer, real, double-precision, or quad-precision variable called the DO-variable.

ex_1 is an integer, real, double-precision, or quad-precision expression that specifies the beginning value of v on the initial execution of the DO loop. (Each succeeding value is determined by $ex_1 + ex_3$, $ex_1 + 2 \times ex_3$, $ex_1 + 3 \times ex_3$).

ex_2 is an integer, real, double-precision, or quad-precision expression that specifies the ending value of v .

ex_3 is an integer, real, double-precision, or quad-precision expression that indicates the increment for v after each time the body of the DO loop is executed. If you omit the increment value, it defaults to 1.

The DO statement executes a loop that begins at the DO and ends with the terminal statement. During execution, CONVEX FORTRAN evaluates ex_1 , ex_2 , and ex_3 to determine the beginning value of v , to determine the final value of v , and to set the step value. Each time the DO loop executes, the value of the increment expression (ex_3) is added to the DO variable, and the iteration count is reduced by 1. The DO loop executes until the value of v exceeds ex_2 .

Transfer of control outside the loop terminates loop execution. The DO-variable v retains its value at loop termination. For example, the following statement:

```
DO 20 MYEXAM = 2, 6, 2
```

executes the DO loop with MYEXAM taking the values 2, 4, and 6. The loop terminates when MYEXAM becomes 8, and this value remains in MYEXAM after termination.

The iteration count is determined by $\text{INT}((ex_2 - ex_1 + ex_3) / ex_3)$. The loop executes until the iteration equals zero. Normally a negative or zero value indicates that the DO loop is not executed. *If you specify the -F66 compiler option, the body of the loop executes once even when the iteration count is zero or negative.* Because internal representation of real numbers is not exact, using a real number for the DO-variable can produce an unexpected count. You cannot redefine the DO-variable (v) within the range of the DO loop. You can, however, alter the initial, terminal, and increment parameters within the loop without affecting the iteration count. The actual terminal and increment values used in the loop are not affected either.

After execution of the DO loop (provided the loop is not nested), control transfers to the first executable statement after the terminal statement. If the loop is nested and the loops share a terminal statement, control transfers outward to the next most enclosing DO loop.

Nested DO loops

You can nest DO loops if each inner DO loop is entirely within the range of the outer DO loop. The DO loop can be entered only through the DO statement. During execution, control can be transferred out of the range before execution is completed and then returned within the range of the DO loop. You cannot, however, transfer control from an outer loop to an inner loop. *DO loops can share terminal statements but not unlabeled END DO statements.* If DO loops share terminal statements, a transfer to that statement can be made only from within the range of the innermost DO.

Example:

```
DO 10 I = 1,100
  .
  .
  DO 10 X = 1,100
    IF (X .GT. XMAX) GOTO 10
    .
    .
10 CONTINUE
```

In this example, the DO loops share a common terminal statement at line 10.

Extended-range DO loops

To maintain compatibility with older FORTRAN implementations, CONVEX FORTRAN allows extended-range DO loops. An extended-range DO loop is a loop in which control transfers outside the body of the DO loop and then back into the loop. The statements in the extended range are logically in the body of the loop.

The following rules apply to extended-range DO statement control transfers:

- You can transfer into the range of a DO statement only if you make the transfer from the extended range of the same DO statement; the transfer is invalid otherwise.
- The extended range of a DO statement must not change its control variable.

The following example illustrates valid control transfer into the range of a DO statement from its extended range.

Example:

```

5   DO 10 I = 1,10
      IF(A(I).LE.0)GOTO 20      !DO LOOP
30  A(I) = -1
10  CONTINUE
      RETURN
20  A(I) = F(X(I))              !EXTENDED RANGE
      GOTO 30

```

The following example illustrates invalid transfer into the range of a DO statement.

Example:

```

      IF (FLAG)GOTO 5
      I = 3
      GOTO 30 !INVALID CONTROL TRANSFER TO DO LOOP
5   DO 10 I = 1,10
      IF (A(I).LE.0)GOTO 20
30  A(I) = -1
10  CONTINUE
20  RETURN

```

DO WHILE statement

The DO WHILE statement allows continual execution of a DO loop as long as a logical expression contained in the statement remains true. This statement has the following form:

```
DO [s [, ] ] WHILE (e)
```

where

s

is the label of an executable statement that must physically follow in the same program unit.

e

is a logical expression.

The *DO WHILE* statement checks the value of the logical expression at the beginning of each iteration of the loop, starting with the first. When the value is true, execution of the statements in the loop follows; when the value is false, control transfers to the statement following the loop. If you do not include a label in the *DO WHILE* statement, you must terminate the *DO WHILE* loop with an *END DO* statement.

Example:

```
DO WHILE (ARRAY (I,J) .GT. 1.0)
    ARRAY (I,J) = ARRAY(I,J)/2.0
    I = I+1
    J = J-1
END DO
```

The condition is only tested at the top of the loop. If the condition becomes true during the execution of the loop, the loop does not terminate until control passes back to the top of the loop.

***END DO* statement**

The *END DO* statement ends the *DO* and *DO WHILE* statements. You must include the *END DO* statement at the end of a *DO* loop if the *DO* or *DO WHILE* statement defining the loop does not contain a terminal-statement label. You can also use the *END DO* statement as a labeled terminal statement if the *DO* or *DO WHILE* statement does contain a terminal-statement label.

Example:

```
REAL A(101)
DO 10 I = 1,100
A(I) = SIN(A(I))/COS(A(I+1))
10 END DO
```

```
REAL A(101)
DO I=1,100
A(I)=SIN(A(I))/COS(A(I+1))
END DO
```

CONTINUE statement

Because the execution of a CONTINUE statement has no effect, you can place it anywhere in a program that an executable statement is allowed. When used as a terminal statement in a DO loop, a CONTINUE statement must be labeled; otherwise, no label is required. This statement has the following form:

```
[s] CONTINUE
```

where *s* is an optional statement label.

CALL statement

The CALL statement transfers program control to a subroutine. It has the following form:

```
CALL sub([a [,a] ... ])
```

where *sub* is the name of the subroutine and *a* is an actual argument or argument list. Refer to Chapter 11, "Subprograms," for more information.

RETURN statement

The RETURN statement returns control from a subroutine to the calling program unit. It has the following form:

```
RETURN [e]
```

where *e* is an optional integer expression that specifies an alternate statement in the calling program that is to receive control. Refer to Chapter 11, "Subprograms," for a complete description.

STOP statement

The STOP statement terminates program execution and has the following form:

```
STOP [s]
```

where *s* is a string of five or fewer digits, or a character constant to be displayed when the STOP statement executes.

Example:

```
STOP '-JOB FINISHED'
```

PAUSE statement

The PAUSE statement suspends program execution until the operator orders execution to resume. It has the following form:

```
PAUSE [s]
```

where

s

is a string of five or fewer digits or a character constant to be printed.

Example:

```
PAUSE 'LOAD TAPE NUMBER 1'
```

When the preceding statement is executed, the system displays the following information:

```
PAUSE:  LOAD TAPE NUMBER 1
To resume execution, type: go
Any other input will terminate the program.
```

To continue, type **go** and press **RETURN**. The system displays:

```
Execution resumed after PAUSE.
```

END statement

The `END` statement ends a main program without displaying a message. In a function or subroutine subprogram, it returns control to the calling program and performs the same function as a `RETURN` statement in a subprogram. It has the following form:

```
END
```

The `END` statement must end every program unit and can appear only in columns 7 through 72 of an initial line.

Input/output statements

9

Input/output (I/O) statements provide a method for transferring data between internal storage and external media or between internal storage and internal files. CONVEX FORTRAN supports *READ*, *ACCEPT*, and *DECODE* statements for input, and *WRITE*, *TYPE*, *PRINT*, and *ENCODE* statements for output. Table 10 lists supported I/O statements by category.

Table 10
Data transfer I/O statements

Statement category	Statement name						
	READ	WRITE	ACCEPT	TYPE	PRINT	DECODE	ENCODE
Sequential/ External:							
Formatted	Yes	Yes	Yes	Yes	Yes	No	No
Unformatted	Yes	Yes	No	No	No	No	No
List-directed	Yes	Yes	Yes	Yes	Yes	No	No
<i>Namelist-directed</i>	Yes	Yes	Yes	Yes	Yes	No	No
Sequential/ Internal:							
Formatted	Yes	Yes	No	No	No	Yes	Yes
<i>List-directed</i>	Yes	Yes	No	No	No	No	No
Direct/ External:							
Formatted	Yes	Yes	No	No	No	No	No
Unformatted	Yes	Yes	No	No	No	No	No

Auxiliary statements control the connection of files to external devices, position files, or retrieve information about a file or unit. These statements are *OPEN*, *CLOSE*, *REWIND*, *INQUIRE*, *BACKSPACE*, *ENDFILE*, and *FIND*.

Note

Unformatted internal I/O statements, direct list-directed, and direct namelist-directed I/O statements are not allowed. All other variations are allowed.

Records

A sequence of characters or values processed as a unit constitutes a record; I/O statements transfer all data as records. Formatted records contain characters; unformatted records (those written without format specification) consist of bytes that represent binary values and have header and trailer fields containing the record length. Each unformatted I/O statement transfers one record. Formatted, list-directed, and namelist-directed I/O statements transfer as many records as required by the I/O data list. Each read or write starts a new record.

Formatted records

A formatted record contains a sequence of characters (letters, numbers, and special symbols). Formatted records can be read or written only by formatted input/output statements. You cannot use formatted I/O on files connected for unformatted access.

With formatted input, if the input statement requires more characters than are available, characters are read as spaces. If the input statement does not require all the characters in the record, unneeded characters are ignored.

The processor reads or writes the current record and possibly additional records during data transfer. The length of the record is measured in characters and depends on the number of characters written to the record. The length can be zero. Any record values left unfilled during data transfer to fixed-length records are written as spaces. When the size of the data is greater than the record length and when an output statement writes to a fixed-length record, an error condition occurs.

Unformatted records

An unformatted record is a sequence of zero or more bytes, surrounded by a four-byte header and a four-byte trailer, each containing the record length. You cannot use unformatted I/O on files connected for formatted I/O. For each unformatted I/O statement, the processor reads or writes one record.

The number of bytes written determines the length of the unformatted record; the length can be zero. On input, if the data list requires more bytes than are available, an error condition occurs. For fixed-length records, the data list in the output statement must not specify more values than the record can hold. Any record bytes left unfilled during data transfer to fixed-length records become zeros.

ENDFILE record

The ENDFILE statement writes the ENDFILE record that ends the file. An ENDFILE record is also written when a file opened for writing is closed, either through the CLOSE statement, through a REWIND statement, or implicitly through program termination. The ENDFILE record appears only as the last record of a file. When such a file is closed with a REWIND statement, the ENDFILE record is written at the current position before rewinding. You cannot use an ENDFILE statement on a file connected for direct access.

Files

A sequence of records that are input to or output from a program constitutes a file. A file is either internal (array or variable) or external (located on a peripheral device). There are two methods of accessing files: sequential and direct.

ConvexOS allows you to create and manipulate files up to one terabyte minus 512 bytes. This can be done transparently with Version 7.0 of CONVEX FORTRAN and all later versions. For complete details and limitations, refer to the *ConvexOS Extensions User's Guide*. File size limits due to non-operating system factors are noted where appropriate.

Internal files

An internal file is a character variable, array, array element, or substring into which records are read or written. If the file consists of a character variable, array element, or substring, it

constitutes a single record. When the file consists of an array, each element constitutes a record. Internal files provide transfer and conversion of data from internal storage to internal storage.

A record in the internal file can be read only if the record is defined. When the processor writes the record, the record of the internal file becomes defined. Also, you can use character assignment statements to define a record.

You can specify an internal file only in READ, WRITE, ENCODE, and DECODE statements.

Units

Before you can access an external file, you must associate (connect) it with a unit. Executing the OPEN statement accomplishes the connection by assigning a logical number to the external file. This number is the unit designator, which provides a means for referencing the file. Internal files are not connected or opened but are referenced by variable, array, or substring name. Connection also can be accomplished implicitly by the system. You cannot connect a file to more than one unit at a time. You can, however, connect a unit to a file that does not exist, that is, a new file that has not been written.

The following statements illustrate various ways to open a file. For instance, the statement

```
OPEN (7)
```

opens the file `fort.7`; this is the file associated with unit 7 by default. The following statement:

```
OPEN (8, FILE='TEST.DAT')
```

connects unit 8 to the file `TEST.DAT`.

The following statement:

```
OPEN (9, STATUS='SCRATCH')
```

opens a scratch (temporary) file associated with unit 9. When the file is closed or the program ends, the file is deleted.

When an OPEN statement is executed for an unopened unit, the program environment is searched for a shell variable associated with the unit. This variable is named `FORnnnOPEN`, where `nnn` is a three-digit number representing the unit (for example, from 000 to 999, leading zeroes required). This shell variable can contain

attributes that override the attributes specified in the `OPEN` statement for that unit. Refer to the "Conversion using a shell variable" section of this chapter for more information and examples.

In the absence of an associated shell variable, a unit takes its attributes from the list of attributes specified with the `OPEN` statement. Attributes not specified in the shell variable are also taken from the `OPEN` statement.

Note

When the data format attribute (see the "Binary datafile format conversions" section of this chapter) is specified, either in an `OPEN` statement or a shell variable, the source code must be compiled with `CONVEX FORTRAN V6.0` or higher and linked with `CONVEX FORTRAN V6.0` or higher libraries. If a version of `CONVEX FORTRAN` prior to `V6.0` is used, unpredictable behavior, including program aborts and errors, can result.

To reassign the unit, terminate the connection. A `CLOSE` statement (or an `OPEN` statement for another file) terminates the connection. The connection is terminated implicitly when the program ends.

Accessing files

You can use either the sequential or the direct method for accessing records of a file. Connection of a file to a unit, typically accomplished with an `OPEN` statement, determines the method of access.

Sequential access

To connect a file for sequential access, use the `OPEN` statement.

Examples:

```
OPEN (10, FILE='MYEXAM', ACCESS='SEQUENTIAL')
OPEN (10, FILE='MYEXAM') ! SEQUENTIAL ACCESS
                           ! BY DEFAULT
```

If you do not specify the `ACCESS` keyword, the access mode defaults to `SEQUENTIAL`. To change the access mode, close the file and reopen it specifying the new mode.

A file connected for sequential access cannot be read or written with direct access I/O statements. A data-transfer statement causes the next record to be read or written when a file is

connected for sequential access; the records are accessed in order of placement in the file. The last record must be an endfile record.

Direct access

Connecting a file for direct access allows the records to be written or read in any order. The record number specified in the I/O transfer statement determines the order of processing. You cannot use sequential access I/O statements on files connected for direct access.

To establish a direct-access file, open a unit for direct access.

Example:

```
OPEN(10, FILE='MYEXAM', ACCESS='DIRECT', RECL=1024)
```

All records of a direct-access file have the same length. The record size is specified in bytes when the file is opened. Every time you read or write a record, you must specify a record number to indicate the record to be read or written.

I/O statement format

The general format of an I/O transfer statement is as follows:

```
READ (clist) iolist  
WRITE (clist) iolist
```

where

clist

is the control information list that controls the data transfer.

iolist

is the I/O list that specifies the data to be transferred.

If invalid data is encountered in a READ statement, execution stops at that point and the remaining variables in the *iolist* are ignored.

Input/output lists

The I/O lists (*iolist*) identify the entities whose values are transferred by I/O data-transfer statements. An *iolist* entity can be:

- Character substring name (CHAR (6 : 10))
- Variable name (L)
- Array name (MYEXAM)
- Array element name (M(3))
- Implied-DO list (J, K, L, M, I=1, 4)
- An expression (K + L or 'JKL'); used for output only. The expressions cannot contain function references with I/O statements in them.

When an array name without a subscript appears in an *iolist*, the elements are processed in the order in which they are stored, for example, M(1, 1), M(2, 1), and so on.

Implied-DO lists

An implied-DO list is used for specifying repetition of part of an I/O list, transferring part of any array, and transferring array elements in an order that is not the same as the order in which they are stored. The implied-DO loop has the form:

$$(dlist, v=ae_1, ae_2 [, ae_3])$$

where

dlist
is an I/O list.

v
is an integer or real variable.

ae₁, *ae₂*, *ae₃*
are arithmetic expressions.

The variable and arithmetic expressions have the same forms and functions as those in the standard DO statements. The loop begins with the value of *ae₁* and increments by the value of *ae₂* until it equals or exceeds the value of *ae₃*. The loop then exits. Elements in *dlist* can reference *v*, but cannot change the value of *v*. The implied-DO loop can be nested.

The following statements illustrate uses of the implied-DO loop.

Example:

```
WRITE (7) (A,B,I=1,10)
C   WRITES THE PAIR A,B 10 TIMES

READ (7) (A(I),I=5,10)
C   READS ELEMENTS 5 THROUGH 10 OF ARRAY A

WRITE (7) ((A(I,J),J=1,N),I=1,N)
C   WRITES THE ARRAY A BY ROWS
```

Specifiers

There are seven specifiers for use in the control information list to provide information on various aspects of data transfer. Each specifier includes a keyword, an equal sign, and a parameter for the specifier. The specifiers are:

- Unit
- Format
- Record
- Status
- Error
- End of file
- *Namelist*

Unit specifier

The unit specifier identifies the external or internal unit being accessed. It has the following form:

[UNIT=] *u*

where *u* is an internal or external identifier. As an external file identifier, *u* is an integer in the range 0 to 255 or *, which defaults to a preassigned input or output unit. As an internal file identifier, *u* is the name of a character variable, array, array element, or substring.

The keyword UNIT= is optional if the unit specifier is the first item in a list of specifiers.

Format specifier

You must include a format specifier in each data transfer statement to or from a formatted file. A format specifier is the label of a `FORMAT` statement, a character expression within the transfer statement, or an asterisk indicating list-directed formatting. A format specifier has the form:

[`FMT =`] *f*

or

[`FMT =`] *

where

f

is a character expression (character constant or name of a character variable, array element, or substring) that contains a runtime format, a statement label of a `FORMAT` statement, or an integer variable with an assigned `FORMAT` statement label. The `FORMAT` statement must be in the current program unit.

*

indicates list-directed formatting that uses default formatting based on the I/O list data types.

If the first item of the control information list is the unit specifier (without the keyword `UNIT=`) and the second item is the format specifier, you can omit `FMT=` from the format specifier. If no format specifier is included, the I/O statement is unformatted.

Record specifier

The record specifier, when used in a data-transfer statement, indicates which record is to be read or written in a file connected for direct access. You can not use the record specifier for sequentially accessed files.

A record specifier has the following forms:

`REC = r`

or

'*r*'

where r is a numeric expression with a positive value that specifies the position of the record to be accessed for I/O. While r can be any size INTEGER, the value of r must fit in an INTEGER*4 storage location. *If the second form is used, the unit specifier cannot use the UNIT keyword and the value for r must appear immediately after the unit specifier with no intervening comma, for example, WRITE (5'10).* The second form is valid only if the `-vfc` compiler option is specified.

Note

Because the record number for I/O must be accessible through an INTEGER*4 variable assigned via the REC specifier in an I/O statement, the number of records in the file cannot exceed $2^{31}-1$ records. This may override the normal one terabyte-512 bytes file size limit under ConvexOS. Refer to the *ConvexOS Extensions User's Guide* for more information.

Status specifier

The status specifier provides a means for determining an error or end-of-file condition. A status specifier has the form:

IOSTAT = *ios*

where *ios* is an integer variable or array element.

After the I/O statement containing the status specifier executes, the status variable contains one of the following:

- A positive integer, which indicates an error condition exists; this integer is the error number.
- 0, which indicates normal execution; no error or end-of-file condition exists.
- -1, which indicates end-of-file condition.

If you indicate only the status specifier (no END or ERR specifier) in the statement and an error condition exists or an end-of-file condition occurs during program execution, program execution continues at the next executable statement.

Error specifier

An error specifier designates a statement to receive control if an error occurs during program execution. An error specifier has the following form:

ERR = *s*

where *s* is the label of an executable statement in the same program unit as the error specifier.

When the processor detects an error during program execution, the I/O statement terminates immediately. The value of the status specifier (if included) becomes a positive integer and control transfers to the statement whose label appears in the error specifier.

End-of-file specifier

You can use the end-of-file specifier in a statement to transfer control to a specific statement on an end-of-file condition. An end-of-file specifier has the form:

END = *s*

where *s* is the statement label of an executable statement in the same program unit as the end-of-file specifier.

An end-of-file condition exists when the end-of-file record is read in an external file opened for sequential access, or when an attempt is made to read a record beyond the range of an internal file. When an end-of-file condition is detected during program execution, the READ statement terminates, the value of the status specifier (if included) becomes -1, and control transfers to the statement whose label appears in the end-of-file specifier.

Namelist specifier

The namelist specifier specifies that namelist-directed I/O is to be used and specifies the group name for the entities that are modified during input or written on output. The namelist specifier has the following form:

NML = nlgrpname

where nlgrpname is the symbolic name that has been defined for the entities in a NAMELIST statement.

If the first item of the control information list is the unit specifier without the keyword UNIT =, you can omit the keyword NML = from the namelist specifier, but you must place the namelist specifier (nlgrpname) as the second item in the control information list. Otherwise, you must use the keyword NML =. You cannot use the namelist specifier in a statement containing a format specifier.

READ statement

READ is an input statement that assigns values from a record to the *iolist* variables. Execution of the READ statement with an external file causes input data to be transferred from the external file into internal storage or memory. Execution of the READ statement with an internal file causes data to be transferred between internal storage locations.

The statement has the following form:

```
READ (clist) [iolist]
```

or

```
READ f [,iolist]
```

where

clist

is a control information list (described in the "Specifiers" section of this chapter). You must include a unit specifier in the READ statement *clist*, and if the record is formatted, a format specifier. The record specifier must be included for direct-access files. You must include the namelist specifier for namelist-directed I/O. The status, error, and end-of-file specifiers are optional.

iolist

is the I/O list that identifies the data to be transferred. The entities include variables, array elements, substrings, implied-DO lists, or array names.

f

is the format specifier. The specifier is a character array name, character expression, character constant, FORMAT statement label, or an integer variable assigned the label of the FORMAT statement. An asterisk (*) indicates list-directed formatting.

When the READ statement executes, at least one record consisting of values from the I/O list is read. The file is then positioned at the beginning of the next record.

By default, CONVEX FORTRAN reads a sequential, formatted file unless an OPEN statement contains FORM='UNFORMATTED' or ACCESS='DIRECT'. Unformatted reads from internal files are not permitted.

The values transfer as one record unless you include a slash, which indicates a new record, in the format specification.

External sequential-access READ statements

There are four classes of external sequential-access READ statements: formatted, unformatted, list-directed, and namelist-directed. The use of IOSTAT, ERR, and END status specifiers is optional in all four classes of statements.

Formatted

The formatted sequential READ statement, which requires a unit (*u*) and a format specifier (*f*), has the form:

```
READ (u,f [ , IOSTAT,ERR,END] ) [iolist]
```

Example:

```
READ (UNIT=50, FMT=10, IOSTAT=IOERR, ERR=120) D, E, F
READ (50, 10) D, E, F
```

Both statements sequentially read values into D, E, and F according to the format specified by statement 10. The first READ statement returns any error codes in the variable IOERR and transfers control to statement 120 on an error condition.

Unformatted

An unformatted sequential READ statement, which requires a unit specifier (*u*), has the form:

```
READ (u [ , IOSTAT,ERR,END] ) [iolist]
```

Example:

```
READ (UNIT=*, FMT=*, END=260) D, E
READ (50)
```

The first statement sequentially reads from the implicit input unit values into the variables D and E without any conversion. The second statement skips the next record in the file connected to unit 50.

List-directed

A list-directed sequential-access statement, which must contain an asterisk (*) to indicate list-directed formatting, has the following forms:

```
READ (u, * [, IOSTAT, ERR, END] ) [iolist]
READ * [, iolist]
```

Examples:

```
READ (UNIT=50, FMT=*) D, E, F
READ (50, *, IOSTAT=IOERR)
READ *, D, E, F
```

The first statement assigns values to D, E, and F from the current record of the file connected to unit 50. Conversion from ASCII to internal format is done according to the rules for list-directed formatting. The second statement skips the current record of the file connected to unit 50. The last statement reads from the implicit input unit into the variables D, E, and F under list-directed formatting.

Namelist-directed

The namelist-directed sequential *READ*, with a unit specifier (*u*) and a namelist specifier (*nl*) in the control information list, has the form:

```
READ (u, nl [, IOSTAT, ERR, END] )
```

When the namelist-directed *READ* is used without specifying a control information list, it has the following form:

```
READ nlgrpname
```

where *nlgrpname* represents the name associated with a list of entities.

When you use the namelist-directed *READ* statement, you must have a *NAMELIST* statement in the program segment.

Example:

```
NAMELIST /SAM/ NAME, EXAM1, EXAM2, EXAM3
CHARACTER*5 NAME
READ (UNIT=50, NML=SAM) !or READ SAM
```

The first statement associates the name (*SAM*) with the four entities. The second statement defines *NAME* to be a *CHARACTER*5* variable; *EXAM1*, *EXAM2*, and *EXAM3* are implicitly typed. The third statement reads input data and assigns values to the namelist entities—*NAME*, *EXAM1*, *EXAM2*, and *EXAM3*. The *READ* statement reads data until it finds the specified name (*SAM*). Then it translates the data from external to internal form, using the data type of the entities and the form of the input. Then the translated data is assigned to the specified entities (*NAME*, *EXAM1*, *EXAM2*, and *EXAM3*) in the order they appear in the input records. (See Chapter 10, "Format specifications," for detailed information on inputting values.)

External direct-access READ statements

There are two classes of external direct-access *READ* statements: formatted and unformatted.

Formatted

A formatted direct-access statement must contain a unit specifier (*u*), record number specifier (*m*), and format specifier (*f*). This statement has the form:

```
READ (u,f [, IOSTAT,ERR,END], m) [iolist]
```

Example:

```
READ (119, 100, REC=25) D, E, F
```

This statement uses the format given at line 100 to read record number 25 of the file connected to unit 119 and assigns values to variables *D*, *E*, and *F* from this record.

The *REC* and *END* keywords are mutually exclusive. The following statement is invalid:

```
READ (10, 20, REC=1, END=30) A      ! Wrong!
```

Unformatted

An unformatted direct-access *READ* statement, which must contain a unit (*u*) and record number (*m*) specifier, has the following form:

```
READ (u,m [, IOSTAT,ERR,END]) [iolist]
```

Example:

```
READ (50, REC=1) D, E, F
```

In this case, the statement reads the first record of the file connected to unit 50 and assigns values from it without translation to the variables D, E, and F.

Internal READ statements

The internal READ statement transfers and converts information from internal storage. In the internal READ statement, the name of the character variable, array, array element, or substring (*iu*) is used in place of the external identifier. The use of IOSTAT, ERR, and END status specifiers is optional.

There are two types of internal READ statements: sequential-access and direct-access.

Sequential access

The internal sequential-access READ statement is always formatted and has the following form:

```
READ (iu,f [, IOSTAT, ERR, END]) [iolist]
```

The following statement transfers values from MYEXAM to A and B, converting them from ASCII to internal form according to the format at line 25.

```
READ (MYEXAM,25) A, B
```

The following statement uses list-directed formatting:

```
READ (MYEXAM,*) A, B
```

Direct-access

The internal direct-access READ statement has the following form:

```
READ (u,f,m, [IOSTAT, ERR, END]) [iolist]
```

Example:

```
READ (ARR, 10, REC=2) A
```

This statement converts the second element of the array *ARR* from ASCII to internal form and stores the result in *A*. The logical record length is the length of the array element. Thus, a character variable array is similar to a fixed-length, direct-access file, and follows the same rules.

ACCEPT statement

The *ACCEPT* statement sequentially reads data from the standard input unit and has the following formats:

ACCEPT f [,iolist]

or

ACCEPT [,iolist]*

or

ACCEPT nigrpname

where

f
is the non-keyword form of a format specifier.

specifies list-directed formatting.

iolist
is an I/O list.

nigrpname
is the non-keyword form of the namelist specifier.

The *ACCEPT* statement is similar to the *READ* formatted or list-directed, sequential, external statement except that reading is always done from the standard input unit.

Example:

```
ACCEPT 100, I, J
100 FORMAT (2I2)
```

As shown, the *ACCEPT* statement reads integer data from the standard input unit and assigns values to the integer variables *I* and *J*.

WRITE statement

The WRITE statement transfers data from internal storage to external devices or from internal storage to internal files. The WRITE statement has the form:

```
WRITE (clist, f) [iolist]
```

where

clist

is a control information list (described in the "Specifiers" section of this chapter) that must include a unit specifier. A formatted record must include a format specifier. A record specifier must be included for direct-access output to a file. Status and error specifiers are optional. (An end-of-file specifier is not allowed in WRITE statements.)

iolist

is the I/O list that identifies the data to be transferred. The entities can include variables, array elements, substrings, implied DO lists or array names, and expressions.

f

is the format specifier and is a character array name, a character expression, a character constant, statement label of a FORMAT statement, or an integer variable assigned the label of the FORMAT statement. An asterisk (*) indicates list-directed formatting.

The WRITE statement writes at least one record consisting of values from the I/O entities. The file is then positioned at the beginning of the next record. In the absence of an OPEN statement, the type of file written by CONVEX FORTRAN is dependent on the form of the WRITE statement used. If an OPEN statement is present and specifies FORM='UNFORMATTED' or ACCESS='DIRECT', an unformatted file is written. Unformatted writes to internal files are not permitted. Direct-access, internal I/O is permitted. The logical record length is the length of the array element. A character variable array is similar to a fixed-length, direct-access file, and follows the same rules as formatted I/O.

A WRITE statement of the following form normally writes to stdout (for example, to the terminal):

```
WRITE (*, *) [iolist]
```

In a FORTRAN 77 subroutine called from a C program, however, this statement writes to the file `fort.6` by default. To change this behavior, you must either initialize FORTRAN I/O by calling `f_init` from the C program, or define the environment variable `FOR006` as `SYSS$OUTPUT`.

The values transfer as one record unless you include a slash, which indicates a new record, in the format specification. If the `WRITE` statement specifies I/O to a nonexistent file, the file is created unless an error condition occurs.

Sequential-access `WRITE` statements

There are four classes of sequential `WRITE` statements: formatted, unformatted, list-directed, and namelist-directed. The use of `IOSTAT` and `ERR` specifiers is allowed in all four classes of statements.

Formatted

The formatted sequential-access `WRITE` statement, which requires a unit (*u*) and a format (*f*) specifier, has the form:

```
WRITE (u,f [,IOSTAT,ERR]) [iolist]
```

Examples:

```
WRITE (UNIT=50, FMT=10, IOSTAT=IOERR, ERR=120) D, E, F
WRITE (50, 10) D, E, F
```

Both statements transfer formatted values from the variables `D`, `E`, and `F` to the file connected to unit 50. The first statement, however, allows for transfer of control if an error condition exists. In this example, if an error condition exists, the error number is assigned to `IOERR` and control transfers to statement 120.

Unformatted

An unformatted sequential `WRITE` statement, which requires a unit (*u*) specifier, has the form:

```
WRITE (u [,IOSTAT,ERR]) [iolist]
```

Examples:

```
WRITE (UNIT=50) D, E
WRITE (50)
```

The first statement writes two unformatted values to unit 50.
The second writes an empty record to unit 50.

List-directed

A list-directed sequential-access `WRITE` statement, which must contain an asterisk (*) to indicate list-directed formatting and a unit (*u*) specifier, has the following form:

```
WRITE (u,* [, IOSTAT,ERR]) [iolist]
```

Example:

```
WRITE (UNIT=50, FMT=*) D, E, F
```

writes D, E, and F according to the default format used for list-directed I/O.

Namelist-directed

The namelist-directed `WRITE` statement, which requires a unit (*u*) and a namelist (*nl*) specifier, has the form:

```
WRITE (u, nl [, IOSTAT,ERR])
```

Example:

```
WRITE (UNIT=50, NML=SAMPLE)
```

In this example, the statement transfers data from the variables specified by the namelist specifier `SAMPLE` to the file connected to unit 50.

External direct-access `WRITE` statements

There are two classes of external direct-access `WRITE` statements: formatted and unformatted. The use of `IOSTAT` and `ERR` status specifiers is optional.

Formatted

A formatted direct-access statement, which must contain a unit (*u*) specifier, record number (*rn*) specifier, and format (*f*) specifier has the form:

```
WRITE (u,f,rn [, IOSTAT,ERR]) [iolist]
```

Examples:

```
WRITE (50, 100, REC=25) D, E, F
WRITE (50, 100, REC=25, ERR=100) D, E, F
```

Both statements write variables D, E, and F to record number 25 of unit 50 according to the format specified in statement 100. The second statement also transfers control to statement 100 if an error condition exists.

Unformatted

An unformatted direct-access statement, which must contain a unit (*u*) and record number (*m*) specifier, has the form:

```
WRITE (u,m [, IOSTAT,ERR] ) [iolist]
```

Examples:

```
WRITE (50, REC=25) D, E, F
WRITE (50, REC=25, ERR=250) D, E, F
```

Both statements write variables D, E, and F to record number 25; no data formatting occurs. The second statement transfers control to statement number 250 if an error condition exists.

Internal WRITE statements

The internal WRITE statement converts data from one location in memory to another. In the internal WRITE statement, the name of the character variable, array, array element, or substring (*iu*) is used in place of the external unit identifier. The use of IOSTAT and ERR status specifiers is optional.

There are two classes of internal WRITE statements: sequential access and direct-access.

Sequential access

The internal sequential-access WRITE statement is always formatted and has the following form:

```
WRITE (iu,f [, IOSTAT,ERR] ) [iolist]
```

Examples:

```
WRITE (MYEXAM,25) A, B
WRITE (MYEXAM,*) A, B
```

These statements transfer values from A and B to MYEXAM, converting them from internal form to ASCII. The first example uses a format at statement 25. The second statement uses list-directed formatting.

Direct access

The internal direct-access WRITE statement has the form:

```
WRITE (iu,f,m, [IOSTAT,ERR,END]) [iolist]
```

Example:

```
WRITE (ARR, 10, REC=2) A
```

This statement transfers the values from A to the second element of ARR, converting the values from internal form to ASCII according to the format specified at statement 10.

PRINT and TYPE statements

You can use either the PRINT statement or the TYPE statement to transfer formatted records to the standard output device. These statements use the sequential mode of access and have the following forms:

```
PRINT f [, iolist] or TYPE f [, iolist]  
PRINT * [, iolist] or TYPE * [, iolist]  
PRINT nlgrpname or TYPE nlgrpname
```

where

f
is the format specifier.

specifies list-directed formatting.

iolist
is an I/O list.

nlgrpname
is the non-keyword form of the namelist specifier.

Example:

```
CHARACTER*16 CLASS, RANK
TYPE 400, CLASS, RANK
400 FORMAT ('CLASS=',A,'RANK=',A)
```

The *TYPE* statement writes one record to the standard output device; the record contains four fields of character data.

**Special
input/output
statements**

The *ENCODE*, *DECODE*, and *FIND* statements are extensions to the ANSI standard and have been included to allow for compatibility with other versions of FORTRAN and for ease in transporting older FORTRAN programs to CONVEX machines.

***ENCODE* statement**

The *ENCODE* statement is equivalent to the internal formatted sequential-access *WRITE* statement. *ENCODE* transfers data between arrays or variables in internal storage and translates the data from internal to character form.

The *ENCODE* statement has the following form:

```
ENCODE (c,f,b [,IOSTAT=ios] [,ERR=s]) [iolist]
```

where

- c* is an integer expression (the number of characters (bytes) to be translated to character form).
- f* identifies the format (an error results if you specify more than one record).
- b* is an array, array element, variable, or character substring reference, any of which receives the characters after translation to external form.
- ios* is either an integer array element or an integer variable that is defined as a positive integer if an error occurs and as a zero if no error occurs.

s

is the label of an executable statement to which control transfers if an error occurs during I/O transfer.

iolist

is an I/O list that contains the data to be translated to character form.

The `ENCODE` statement translates the elements in the I/O list to character form, as specified by the format identifier, and stores the characters in *b*. If the number of characters transferred is less than *c*, the remaining positions are padded with blanks. If *b* is an array, its elements are processed in the order of subscript progression.

The data type of *b* in any given statement determines the number of characters that the `ENCODE` statement processes. An array of `LOGICAL*2`, for example, can contain two characters per element, so that the maximum number of characters is twice the number of elements in that array; a character array can contain characters equal in number to the length of each element multiplied by the number of elements; a character variable or character array element can contain characters equal in number to its length.

The interaction between the format specifier and the I/O list is the same as that of a formatted I/O statement.

Example:

```
CHARACTER*8 A
I=1000
J=9
ENCODE (8,100,A) I,J
100 FORMAT(2I4)
```

Result:

```
A='1000^^^9'
```

In this example, the character string *A* gets the contents of the integers *I* and *J*.

DECODE statement

The *DECODE* statement is equivalent to the internal sequential-access *READ* statement. *DECODE* transfers data between arrays or variables in internal storage and translates the data from character to internal form. The *DECODE* statement has the following form:

```
DECODE (c,f,b [,IOSTAT=ios] [,ERR=s] ) [iolist]
```

where

c
is an integer expression (the number of characters (bytes) to be translated to internal form).

f
identifies the format (an error results if more than one record is specified).

b
is an array, array element, variable, or character substring reference that contains the characters to be translated to internal form.

ios
is either an integer array element or an integer variable that is defined as a positive integer if an error occurs and as a zero if no error occurs.

s
is the label of an executable statement.

iolist
is an I/O list that receives the data after translation to internal form.

The *DECODE* statement translates the character data in *b* to internal (binary) form according to the format specifier and stores the elements in the list. If *b* is an array, its elements are processed in the order of subscript progression.

The data type of *b* in any statement determines the number of characters that the *DECODE* statement processes. An array of *LOGICAL*2*, for example, can contain two characters per element, so that the maximum number of characters is twice the number of elements in that array. A character array can contain characters equal in number to the length of each element

multiplied by the number of elements. A character variable or character array element can contain characters equal in number to its length.

The interaction between the format specifier and the I/O list is the same as that of a formatted I/O statement.

Example:

```
CHARACTER*8 A
DATA A/'1000^^^9'/
DECODE(8,100,A) I,J
100 FORMAT (2I4)
```

Result:

```
I=1000
J=9
```

In this example, the contents of the character string A are transferred to the integers I and J.

***FIN*D statement**

The *FIN*D statement positions a direct-access file to a particular record. It also sets the associated variable of the file to that record number. No transfer of data occurs. The *FIN*D statement is represented as:

```
FIN
```

D ([UNIT=] *u*, REC=*r* [, ERR=*s*] [, IOSTAT=*ios*])

where

u

is a logical unit number; it must refer to a direct-access file.

r

is the direct-access record number; it cannot be less than 1 or greater than the number of records defined for the file.

s

is the label of the executable statement to which control transfers if an error occurs.

ios

is an integer variable or integer array element that is defined as a positive integer if an error occurs and as a zero if no error occurs.

Example 1:

```
FIND (2,REC=1)
```

This statement positions unit 2 to the first record of the file; the associated file variable is set to 1.

Example 2:

```
FIND (4,REC=INDX)
```

*This statement positions the file to the record identified by the content of *INDX*; the associated file variable is set to the value of *INDX*.*

Auxiliary input/output statements

Auxiliary statements control the connection of files to external devices, position files, or retrieve information about files or units. Auxiliary statements include the following:

- OPEN
- CLOSE
- INQUIRE
- REWIND
- BACKSPACE
- ENDFILE

OPEN statement

The OPEN statement connects an existing external file to the specified unit, changes the attributes of a connected file, or creates a new file and connects it to the specified unit. The statement has the following form:

```
OPEN (specifier [,specifier] . . . )
```

Where *specifier* is an expression normally including a keyword and its value. You must include a unit number in the OPEN statement; all other specifiers are optional. The specifiers can be listed in any order except that the unit number must be first when it is given without the UNIT= keyword.

Example:

```
OPEN (7, FILE='TEST.DAT', RECORDTYPE='FIXED', RECL=80)
```

connects the file TEST.DAT to unit 7 and defines the file to be a sequentially accessed formatted file with fixed-length records of 80 characters.

Note

The **OPEN** statement operates somewhat differently under COVUEshell. Please refer to the *CONVEX COVUEshell Reference Manual* for details.

The following sections describe the OPEN statement keywords, which are summarized in Table 11. In the descriptions, the term "numeric expression" can be any integer or real expression. The value of the expression is converted to the INTEGER data type before it is used in the OPEN statement.

Table 11
OPEN statement keywords

Keyword	Values	Function	Default
ACCESS	'SEQUENTIAL' 'DIRECT' 'APPEND'	Access mode	'SEQUENTIAL'
ASSOCIATEVARIABLE	V	Next direct-access record	NA
BLANK	'NULL' 'ZERO'	Interpretation of blanks	'NULL'
BLOCKSIZE	E	Physical block size	System default
CARRIAGECONTROL	'FORTRAN' 'LIST' 'NONE'	Print control	'LIST' (Formatted) 'NONE' (Unformatted)
DEFAULTFILE	C	Default file specification	NA

Table 11 (continued)
OPEN statement keywords

Keyword	Values	Function	Default
<i>DISPOSE or DISP</i>	'KEEP' or 'SAVE' or 'DELETE'	<i>File disposition at close</i>	<i>Depends on STATUS keyword</i>
ERR	S	Error transfer label	NA
<i>FILE or NAME</i>	C	File-name specification	fort .n, where n is the unit number
FORM	'FORMATTED' 'PRINT' 'UNFORMATTED' 'UNFORMATTED/ dataformat'	Format type	'FORMATTED' for sequential access; 'UNFORMATTED' for direct access
<i>IOSTAT</i>	V	<i>I/O status</i>	NA
<i>MAXREC</i>	E	<i>Direct-access record limit</i>	Unlimited
<i>NOSPANBLOCKS</i>	None allowed	<i>Ignored—for VAX compatibility only</i>	NA
<i>READONLY</i>	NA	<i>Write protection</i>	<i>Depends on file access rights</i>
<i>RECL or RECORDSIZE</i>	E	Record length	As specified at file creation
<i>RECORDTYPE or RT</i>	'FIXED' 'VARIABLE'	<i>Record structure</i>	'VARIABLE' for sequential access; 'FIXED' for direct access
<i>STATUS or TYPE</i>	'OLD' 'NEW' 'SCRATCH' 'UNKNOWN'	File status at open	'UNKNOWN'
UNIT	E	Logical unit number	NA

Key:

E is a numeric expression.

C is a character expression, *numeric array name*, *numeric variable name*, or *numeric array element name*.

V is an integer variable name.

S is a statement label.

ACCESS keyword

The ACCESS keyword indicates the method of file access—direct or sequential. 'APPEND' implies sequential access with positioning after the last record in the file. You must include the record length, RECL, in the list when ACCESS='DIRECT'. The keyword has the form:

ACCESS = *cex*

where *cex* is a character expression 'DIRECT', 'SEQUENTIAL', or 'APPEND'. The default is 'SEQUENTIAL'.

The following statement opens the file `tst` for sequential access with positioning after the last record in the file.

Example:

```
OPEN (UNIT=10, FILE='tst', ACCESS='APPEND')
```

ASSOCIATEVARIABLE keyword

The ASSOCIATEVARIABLE keyword specifies an integer variable to be updated after each direct access I/O operation. The keyword has the form:

ASSOCIATEVARIABLE = *asv*

where *asv* is an integer variable.

After each direct-access I/O operation, *asv* is set to the number of the next sequential record in the file. This identifier is valid for direct-access mode only; it is ignored for other access modes.

Because the ASSOCIATEVARIABLE is modified by direct-access I/O operations in any routine called by the program in which it is associated, passing it as a parameter to the called routine or declaring it in a COMMON memory area can create a hidden alias. Avoid passing the ASSOCIATEVARIABLE or declaring it in a COMMON memory area.

Note

Optimizing routines that reference the ASSOCIATEVARIABLE can cause unexpected results. Always compile these routines at optimization level `-no`.

BLANK keyword

The **BLANK** keyword determines the interpretation of blank characters in numeric formatted input fields. The keyword has the form:

```
BLANK = blnk
```

where *blnk* specifies the character expression 'NULL' or 'ZERO'. The default is 'NULL' unless the *-F66 compiler option is specified, in which case the default is 'ZERO'*.

If you specify **BLANK='NULL'**, all blanks are ignored. When **BLANK='ZERO'**, all blanks except leading blanks are read as zeros.

BLOCKSIZE keyword

The **BLOCKSIZE** keyword specifies the physical transfer size (in bytes) for the file. The keyword has the form:

```
BLOCKSIZE = bls
```

where *bls* is a numeric expression.

The default is the system default for the device. If you specify **BLOCKSIZE**, the physical record for block devices is set to the value of *bls*, with a maximum of 64 kbytes. For other devices (such as raw tape), the **BLOCKSIZE** value is rounded up to a multiple of the file system block size. An **INQUIRE** statement with a **BLOCKSIZE** keyword returns the rounded-up value.

The following statements write one physical record of 200 bytes to the block-mode tape device */dev/mt12*. The physical record contains two logical records, each 100 bytes long.

Example:

```
CHARACTER*1 A(100), B(100)
.
.
.
OPEN (7, FILE='/dev/mt12', BLOCKSIZE=200,
^ RECORDTYPE='FIXED', RECL=100)
.
.
.
```

```
WRITE (7, '(100A1)') (A(I), I=1,100)
WRITE (7, '(100A1)') (B(I), I=1,100)
.
.
.
```

CARRIAGECONTROL keyword

The *CARRIAGECONTROL* keyword determines the carriage control processing to be used for printing a file. The keyword has the form:

```
CARRIAGECONTROL = cc
```

where *cc* is a character expression having a value equal to 'FORTRAN', 'LIST', or 'NONE'.

The default is 'LIST' for formatted files and 'NONE' for unformatted files. 'LIST' transmits the first character of each formatted output record unchanged. FORTRAN replaces the first character of each formatted output record with the control characters required to interpret it on an ASCII output device. These control characters are **CTRL-L** for start of page, newline for double spacing, and null (no character) for single spacing.

Files created with *CARRIAGECONTROL*= 'LIST' can be printed with the `fpr` utility.

DEFAULTFILE keyword

The *DEFAULTFILE* keyword defines part of a default file specification. The keyword has the form:

```
DEFAULTFILE = c
```

where *c* is a character expression.

The *DEFAULTFILE* keyword is used to fill in missing parts of the file name specified with the *FILE*= or *NAME*= keyword.

Example 1:

```
OPEN(FILE='info',DEFAULTFILE='.dat')
C OPEN FILE INFO.DAT
```

This statement opens the file `info.dat`.

Example 2:

```
OPEN(FILE='info.dat',DEFAULTFILE='/mnt/user1/')  
C OPEN /MNT/USER1/INFO.DAT
```

This statement opens the file `/mnt/user1/info.dat`. Under the COVUEshell, the equivalent statement would appear as in the following example.

Example 3:

```
OPEN(FILE='info.dat',DEFAULTFILE='mnt:[user1]')  
C OPEN MNT:[USER1]INFO.DAT
```

The `DEFAULTFILE` keyword is used mainly for interactively requesting a file name, especially to fill in a part of the file name that is a default, such as a directory name or extension. Any components in the `FILE=` keyword override those in the `DEFAULTFILE=` keyword.

DISPOSE keyword

The `DISPOSE` keyword allows you to keep, save, or delete files connected to the unit when the unit is closed. The keyword has the form:

```
DISPOSE = dis
```

or

```
DISP = dis
```

where `dis` is a character expression having a value equal to `'KEEP'`, `'SAVE'`, or `'DELETE'`.

Specifying `'KEEP'` or `'SAVE'` retains the file after the unit is closed; specifying `'DELETE'` deletes the file. For scratch files, the default is `'DELETE'`. For all other files, the default is `'KEEP'`. The ANSI standard method of deleting a file is to use the `STATUS='DELETE'` keyword in the `CLOSE` statement.

The following example causes the file associated with unit 10 to be deleted when closed.

```
OPEN (UNIT=10, DISP='DELETE')
```

ERR keyword

The ERR keyword specifies a statement number to which control is passed if an error occurs during execution of the OPEN statement. The keyword has the form

ERR = *sl*

where *sl* specifies the statement label of an executable statement that appears in the same program unit as the error specifier.

FILE keyword

The FILE (*or NAME*) keyword specifies the name of the file being connected to the unit. The keyword has the following form:

FILE = *fn*

or

NAME = *fn*

where *fn* represents a character expression. If the FILE specification does not appear in an OPEN statement, the unit is connected to a predefined file.

The following statement opens the file `tst.in` for sequential access and connects it to unit 1.

```
OPEN (UNIT=1, FILE='tst.in')
```

FORM keyword

The FORM keyword indicates either formatted or unformatted I/O. The keyword has the form:

FORM = *f*

where *f* represents the character expression with the value 'FORMATTED', 'UNFORMATTED', or 'PRINT'.

If you do not specify a format specifier in the OPEN statement, formatted I/O is assumed for sequentially accessed files. Unformatted I/O is assumed for direct-access files. *If FORM is specified as unformatted, then it can be followed by an optional data format qualifier:*

FORM = UNFORMATTED[/*dataformat*]

The *dataformat* qualifier specifies the format of the data file. Legal values are *NATIVE* (*CONVEX-NATIVE* or *CONVEX_NATIVE*), *IEEE* (*CONVEX-IEEE* or *CONVEX_IEEE*), *VAX-D* (or *VAX_D*), *VAX-G* (or *VAX_G*), *CRAY*, *CRAYUB* and *USER-DEFINED* (or *USER_DEFINED*). For more information on data file conversions, refer to the "Binary data file format conversions" section at the end of this chapter.

Data file conversion using the *dataformat* qualifier is most useful for programs that run only once or infrequently, programs that read or write small amounts of binary data, and programs whose data is irregular (the layout of the records within a file changes from one record to the next). If the amount of data that a program reads or writes is large and the layout of the data file is regular, consider writing a custom data conversion program instead.

Note

When the *dataformat* attribute is specified, either in an *OPEN* statement or a shell variable, the source code must be compiled with *CONVEX FORTRAN V6.0* or higher and linked with the *CONVEX FORTRAN V6.0* or higher libraries. If a version of *CONVEX FORTRAN* prior to *V6.0* is used, unpredictable behavior, including program aborts and errors, can result.

Specifying *FORM='PRINT'* implies "formatted" and *CARRIAGECONTROL='FORTRAN'* enables vertical format control for that unit. Vertical format control is interpreted at runtime only on sequential formatted writes to a *PRINT* file.

As an alternative to specifying *FORM='PRINT'*, use the *fpr* utility before printing the file to interpret vertical format control.

IOSTAT keyword

The *IOSTAT* keyword provides a variable that is set to indicate the status of an *OPEN* operation. The keyword has the form:

$$\text{IOSTAT} = \text{ios}$$

where *ios* is an integer variable or integer array element. A nonzero value returned in *ios* indicates an error condition.

MAXREC keyword

The *MAXREC* keyword determines the total number of records allowed in a direct-access file. The keyword has the form:

$$\text{MAXREC} = \text{mr}$$

where *mr* is a numeric expression.

The *MAXREC* keyword applies only to direct-access files. The default is an unlimited number of records.

The following statement opens the file associated with unit 1 for direct access. Records past record number 100 cannot be accessed.

Example:

```
OPEN (1, ACCESS='DIRECT', MAXREC=100)
```

NOSPANBLOCKS keyword

The *NOSPANBLOCKS* keyword specifies that records must not cross disk block boundaries. This keyword is provided for VAX compatibility only and is ignored by CONVEX FORTRAN. The keyword has the form:

```
NOSPANBLOCKS
```

READONLY keyword

The *READONLY* keyword specifies that an existing file can be read but not written. The keyword has the form:

```
READONLY
```

Using *READONLY* in an *OPEN* statement does not prevent the file from being removed when it is closed.

RECL keyword

The *RECL* (or *RECORDSIZE*) keyword specifies the record size for fixed-length records and the maximum record size for variable-length files. The keyword has the form:

```
RECL = ie
```

or

```
RECORDSIZE = ie
```

where *ie* is an integer expression with a positive value.

You must specify *RECL* when the file is opened with *RECORDTYPE='FIXED'*. The record size is measured in bytes; the default is 80 bytes. For fixed-length records, *RECL* is the size of the logical record buffer. For variable-length records, *RECL* is an initial approximation of the logical record size and the buffer is incremented in multiples of *RECL* bytes.

The following statement opens the direct-access unformatted file connected to unit 10. Each record in the file is 20 bytes long.

```
OPEN (UNIT=10, RECL=20, ACCESS='DIRECT', FORM='UNFORMATTED')
```

RECORDTYPE keyword

The *RECORDTYPE* keyword specifies fixed- or variable-length records for a file. The keyword has the form:

```
RECORDTYPE = typ
```

where *typ* is a character expression whose value is equal to 'FIXED' or 'VARIABLE'. The defaults are shown in Table 12.

Table 12
RECORDTYPE defaults

File access	Default RECORDTYPE
<i>Direct</i>	<i>FIXED</i>
<i>Sequential</i>	<i>VARIABLE</i>
<i>List-directed</i>	<i>VARIABLE</i>

If you specify *RECORDTYPE='FIXED'*, you must also specify *RECL*. *RECORDTYPE='VARIABLE'* is not allowed for direct-access files.

The following statement specifies sequential-access, fixed-length records, each of which is 10 bytes long. The file is connected to logical unit 10.

Example:

```
OPEN (10, RECORDTYPE='FIXED', RECL=10)
```

SHARED keyword

The *SHARED* keyword is provided for compatibility with FORTRAN source written for non-UNIX operating systems only. ConvexOS, like UNIX, provides for shared files by default. The compiler ignores the *SHARED* keyword.

STATUS keyword

The *STATUS* (or *TYPE*) keyword determines the status of the file to be opened; the default is 'UNKNOWN'. The keyword has the form:

```
STATUS = sta
```

or

TYPE = *sta*

where *sta* is a character expression with the value of 'OLD', 'NEW', 'SCRATCH', or 'UNKNOWN'.

If you specify *STATUS*='OLD', the file must exist. To create a new file, specify *STATUS*='NEW' in the OPEN statement. When 'SCRATCH' is designated as the status, the unit is connected to a predefined file and must not be named. When the CLOSE statement is executed, the 'SCRATCH' file is deleted. When *STATUS*='UNKNOWN', CONVEX FORTRAN searches to see if the file exists. If it does, status becomes 'OLD'; if it does not exist, status becomes 'NEW'.

If the *STATUS* (or *TYPE*) keyword is not specified, by default, scratch files are deleted and all other files are retained.

UNIT keyword

The UNIT keyword specifies the logical unit to which a file is to be connected. The keyword has the form:

[UNIT=] *u*

where *u* is a numeric expression. Valid logical unit numbers are 0 through 255.

If the unit number appears as the first parameter of the OPEN statement, the UNIT keyword can be omitted; otherwise, the UNIT keyword is required.

If a unit is connected to a file but the *FILE*= specifier does not appear in the OPEN statement, the file to which the unit is currently connected is opened. In this case, the *BLANK*= and *FORM*= specifiers are the only specifiers that can have a value different from the one currently in effect. When the OPEN statement executes, the new value of the *BLANK*= specifier becomes effective. The position of the file is unaffected.

If the file to be opened is not the file to which the unit is connected, the effect is the same as executing a CLOSE statement immediately prior to executing the OPEN statement.

If a file is connected to a unit, you cannot reopen the file with a different unit number.

CLOSE statement

Use the CLOSE statement to disconnect a file from a unit. The CLOSE statement has the following forms:

```
CLOSE ([UNIT=u] [, STATUS=p] [, ERR=s] [, IOSTAT=ios])
```

or

```
CLOSE ([UNIT=u] [, DISPOSE=p] [, ERR=s] [, IOSTAT=ios])
```

or

```
CLOSE ([UNIT=u] [, DISP=p] [, ERR=s] [, IOSTAT=ios])
```

where

u

is a logical unit number that must be an integer expression.

p

is a character expression that determines the disposition of the file. Its values are 'KEEP', 'SAVE', or 'DELETE'.

s

is the label of an executable statement.

ios

is an integer variable or integer array element.

The following statement

```
CLOSE (7, STATUS='DELETE')
```

disconnects the file opened to unit 7 and deletes it. Specifying either 'SAVE' or 'KEEP' retains the file after you close the unit. If the unit is not connected to a file, the CLOSE statement has no effect.

The status specification supersedes the disposition specified in the OPEN statement. For scratch files, the default is 'DELETE'. For all other files, it is 'KEEP'. If you disconnect a unit or file by the CLOSE statement, either can be connected again within the same executable program to the same file or unit.

INQUIRE statement

The INQUIRE statement determines specific information about a file or unit, such as the access mode or block size. This statement has the following forms:

```
INQUIRE (FILE=fi, list)
```

or

```
INQUIRE ( [UNIT=u], list)
```

where

fi

is a character expression, *numeric array name*, *numeric variable name*, or *numeric array element name* whose value is the name of the file being queried.

Note

The *FILE* name you specify when using the INQUIRE statement under the COVUEshell must be the absolute ConvexOS path name. VMS path names are not recognized.

list

is a list of specifiers that indicate the information to be determined for the file or unit. Each specifier appears in the list only once. Table 13 describes the valid specifiers.

u

is the external unit specifier (number) that identifies the unit to be queried. The unit need not exist nor need it be connected to a file. If the unit is connected to a file, the inquiry includes the connection and the file.

Although you can position FILE=*fi* and UNIT=*u* any place in the list that specifies properties, if you omit the UNIT keyword, *u* must be the first item in the list.

The following statement returns the access mode of the file connected to unit 99 in the character variable ACC.

Example:

```
INQUIRE (99, ACCESS=ACC)
```

The following statement returns the form of the file, 'FORMATTED' or 'UNFORMATTED', in the character variable FM.

Example:

```
INQUIRE (FILE='TEST.IN', FORM=FM)
```

Table 13 enumerates and describes INQUIRE specifiers.

Note

Because the next record number must be accessible through an INTEGER*4 variable assigned via the NEXTREC specifier in an INQUIRE statement, the number of records in the file cannot exceed $2^{31}-1$ records. This may override the normal one terabyte-512 bytes file size limit under ConvexOS. Refer to the *ConvexOS Extensions User's Guide* for more information.

Table 13
INQUIRE specifiers

Specifier/form	Specifier variable values
ACCESS = <i>character</i> *	'DIRECT' or 'SEQUENTIAL' if connected; 'UNKNOWN' if no connection.
BLANK = <i>character</i> *	'NULL' or 'ZERO' if connected and formatted I/O; 'UNKNOWN' if no connection of unformatted I/O.
BLOCKSIZE = <i>integer</i> *	0 if not connected. Block size set on OPEN; system default if not set on OPEN.
CARRIAGECONTROL = <i>character</i> *	'FORTRAN' if FORTRAN specified on OPEN; 'LIST' if specified on OPEN; 'NONE' if specified on OPEN; 'UNKNOWN' if not connected.
DIRECT = <i>character</i> *	'YES' if direct access permitted; 'NO' if direct access not permitted; 'UNKNOWN' if not connected.
ERR = <i>statement label</i>	Control transfers to statement if error condition.
EXIST = <i>logical</i> *	.TRUE. if by file and exists; .TRUE. if by unit and unit is in allowed set of unit numbers; .FALSE. otherwise.
FORM = <i>character</i> *	'FORMATTED' if connected for formatted; 'UNFORMATTED' if connected for unformatted.
FORMATTED = <i>character</i> *	'YES' if formatted I/O permitted; 'NO' if formatted I/O not permitted; 'UNKNOWN' if file not connected.
IOSTAT = <i>integer</i> *	0 if no error condition; positive integer if error condition.
NAME = <i>character</i> *	<i>file name</i> if file has name; blank if no file name. If the file is not currently open, the absolute pathname is returned.
NAMED = <i>logical</i> *	.TRUE. if file has a name; .FALSE. if no name.
NEXTREC = <i>integer</i> *	<i>next record number</i> if record length specified on OPEN.
NUMBER = <i>integer</i> *	Unit number of file connected; -1 if no unit connected.
OPENED = <i>logical</i> *	.TRUE. if file/unit connected; .FALSE. if file/unit not connected.
RECL = <i>integer</i> *	<i>record length</i> set on OPEN if connected for direct access; 0 otherwise.
RECORDTYPE = <i>character</i> *	'FIXED' if fixed-length record; 'VARIABLE' if variable-length record; 'UNKNOWN' if not connected.
SEQUENTIAL = <i>character</i> *	'YES' if sequential access permitted; 'NO' if sequential access is not permitted; 'UNKNOWN' if not connected.
UNFORMATTED = <i>character</i> *	'YES' if unformatted records permitted; 'NO' if unformatted records not permitted; 'UNKNOWN' if undetermined.

*The specifier variable can be either a variable or array element of the stated type.

File-positioning statements

File-positioning statements allow manipulation of external files. You cannot use these statements with internal files. The positioning statements are:

- **REWIND**—repositions before the first record.
- **BACKSPACE**—repositions to beginning of preceding record.
- **ENDFILE**—writes an endfile record.

File-positioning statements have the following form:

```
REWIND ( [UNIT=]u [,ERR=s] [,IOSTAT=ios] )
```

or

```
REWIND u
```

```
BACKSPACE ( [UNIT=]u [,ERR=s] [,IOSTAT=ios] )
```

or

```
BACKSPACE u
```

```
ENDFILE ( [UNIT=]u [,ERR=s] [,IOSTAT=ios] )
```

or

```
ENDFILE u
```

where

u

is the unit specifier. If the unit specifier is the first argument, you can omit *UNIT=keyword*.

s

is the statement label to which control transfers if an error condition exists. (If *IOSTAT* and *ERR* are omitted, the program terminates on an error.)

ios

is an integer variable or integer array element that is set to either a zero if no error condition exists, or a positive integer error code if an error occurs during program execution. (If *IOSTAT*, without *ERR*, is included in the statement, execution continues at the next statement on an error.)

REWIND statement

The REWIND statement positions a file at its initial point. If the file is already at its starting point, REWIND takes no action. If the unit is not connected to a file, REWIND has no effect. The following statements reposition the file MYEXAM to its beginning.

Example:

```
.  
. .  
. .  
OPEN (10, FILE='MYEXAM', STATUS='OLD')  
READ (10, END=200) A, B, C  
. .  
. .  
200 REWIND 10  
. .  
. .  
. .
```

BACKSPACE statement

The BACKSPACE statement positions the file connected to the specified unit to the preceding record. If the file is already at the first record, no action is taken. If the file is positioned after the endfile record, BACKSPACE positions the file before the endfile record. You cannot backspace a file that does not exist. Do not attempt to BACKSPACE in an unformatted sequential access Cray pure data file.

The following statement repositions the file connected to unit 10 to the beginning of the preceding record.

Example:

```
BACKSPACE 10
```

The following statements assign A and B the same value from the file connected to unit 8.

Example:

```
READ (8, *) A  
BACKSPACE (8)  
READ (8, *) B
```

ENDFILE statement

The ENDFILE statement writes an endfile record to the file connected to the specified unit and positions the file after the endfile record. After ENDFILE writes the endfile record, no additional records can be read or written without using BACKSPACE or REWIND to reposition the file for data-transfer operations.

The following statements write endfile records to the files connected to units 101 and 4, respectively.

Example:

```
ENDFILE (UNIT=101)
ENDFILE (4)
```

Binary data file format conversions

The binary data file format conversion feature of CONVEX FORTRAN allows programs to read from and write to unformatted, binary data files whose data format is different from the program's mode. Mode refers to the format in which the program is processing data, either NATIVE or IEEE. For example, a CONVEX FORTRAN program in native mode can read from and write to an unformatted binary file whose format is vax-g.

Note

*The REAL*16 data type is only supported in native floating point mode; programs using IEEE mode can neither read nor write REAL*16 data.*

CONVEX FORTRAN allows conversions of the data formats listed in Table 14.

Table 14
Data format conversion
routine names

<i>Format</i>	<i>Data format names</i>
<i>CONVEX native</i>	<i>convex_native,</i> <i>convex-native, or native</i>
<i>CONVEX IEEE*</i>	<i>convex_ieee, convex-ieee,</i> <i>or ieee</i>
<i>VAX D_floating</i>	<i>vax_d or vax-d</i>
<i>VAX E_floating</i>	<i>vax_d, vax-d, vax_g or vax-g</i>
<i>VAX G_floating</i>	<i>vax_g or vax-g</i>
<i>VAX H_floating</i>	<i>vax_d, vax-d, vax_g or vax-g</i>
<i>Cray†</i>	<i>CRAY</i>
<i>Cray unformatted un- blocked sequential access</i>	<i>CRAYUB</i>
<i>User-Defined</i>	<i>user_defined or user-defined</i>

**REAL*16* is supported in native floating point mode only.

† See Appendix D for details on converting Cray files.

Note

You must use the `cvbin` utility to convert VAX files into a format recognizable by CONVEX FORTRAN before attempting to read these files. `cvbin` can be used to copy and convert files from a DECnet or COVUEnet node, or it can be used to convert local files. Refer to the `cvbin(1)` man page or to the CONVEX COVUEbinary *User's Guide* for more information.

You convert data by specifying one of the data format names listed above in an `OPEN` statement or by using a shell variable. Both methods are described in this section. When you specify a data format that is different from the program's current mode, a data conversion routine converts the data when `READ` or `WRITE` statements execute.

For more information on reading and writing Cray files, refer to Appendix D, "Cray FORTRAN compatibility." For more information on reading and writing VAX FORTRAN files, refer to Appendix E, "VAX FORTRAN compatibility."

When to use the conversion feature

Specify conversion only in programs that match one of the following criteria:

- "One-time" programs or programs run infrequently.

- Programs that read or write small amounts of binary data.
- Programs whose data files contain record layouts that vary from record to record if writing a conversion program is not practical or cost-effective.

Do not use the conversion feature on programs that repeatedly read unchanging data in the same data file over and over again. A custom conversion program converts data permanently, but the conversion feature must convert data each time it is read.

Programs that read very large files can use a standalone conversion program that is optimized for its particular data file format. This method should be more efficient if data values are known to fall in certain ranges and if issues such as overflow and underflow can be ignored.

-dFc option

If you are using the `-vFc` compiler option and plan to use the format conversion feature, then you must also use the `-dFc` compiler option. The `-dFc` option helps convert VAX data by instructing the compiler to decompose the VAX record I/O element by element.

Conversion using OPEN statement

To convert data using the `OPEN` statement, use the `FORM` keyword to specify unformatted data and a data format type:

```
OPEN (... , FORM='UNFORMATTED/format' , ...)
```

where *format* is one of the data format names listed in Table 14 and indicates the current data format of the file you want to convert. You can specify the data format in uppercase, lowercase, or mixed case.

When binary data is read from the file specified in the `OPEN` statement, the conversion routines convert the data from the specified data format to the data format implied by the program's mode, either `NATIVE` or `IEEE`.

When binary data is written to the file specified in the `OPEN` statement, the conversion routines convert the data from the data format implied by the program's mode to the data format specified in the `OPEN` statement.

Restrictions on conversions

Observe the following restrictions when using data file format conversions:

- If a program which uses *EQUIVALENCE* statements writes or reads a file specifying any type of conversion and the data type of the actual value in memory does not match the data type of the variable specified in the I/O statement, the value usually becomes corrupted.
- For unions in records, the compiler cannot determine the data type of the actual value in memory. Therefore, when a record containing a union is specified in an I/O statement and data format conversion (*-dfc* option) is specified, the compiler issues a diagnostic message. You can modify the program to read or write each structure's elements rather than the entire structure.
This restriction prevents you from reading or writing data that is probably corrupted.
- All conversions performed expect *REAL* data types to contain numeric data. Using *REAL* or *COMPLEX* data types for Hollerith data does NOT work if a floating-point conversion takes place.
- The VAX does not support *INTEGER*8* or *LOGICAL*8* data types, so attempting to read or write these data types when the data format is *vax-d* or *vax-g* causes a diagnostic message to be issued at runtime (unless an *ERR=* or *IOSTAT=* clause is specified).
- *REAL*16* is supported only in native mode. Therefore, native mode programs that attempt to read or write a *REAL*16* value from or to a file opened with data format *IEEE* print a diagnostic message during execution and halt (unless an *ERR=* or *IOSTAT=* expression is specified in the I/O statement).

Note

Because *REAL*16* variables cannot be specified in *IEEE* mode programs, there is no way to read or write *vax-h* or *NATIVE* format *REAL*16* values in *IEEE* mode programs.

- The Cray FORTRAN compiler does not support arithmetic items less than eight bytes long. So attempting to read or write these data types when the data format is *CRAY* or *CRAYUB* causes a diagnostic message to be issued at runtime (unless an *ERR=* or *IOSTAT=* clause is specified).

- *CONVEX FORTRAN's internal representation of floating point data is different from Cray's; the CONVEX representation allows for greater precision through a slightly smaller range of numbers. If you attempt to read a floating point number written on a Cray and the number is larger than the largest number representable under CONVEX FORTRAN, the read will fail with a floating point exception.*
- *When the FORM = UNFORMATTED/CRAY data format conversion is specified, only unformatted, direct access, unblocked ("pure") files can be read. You can convert sequential access, unformatted, blocked files into readable unblocked format with the fcUnblock utility. Refer to Appendix D, "Cray FORTRAN compatibility," for details on reading various unformatted Cray files. Refer to the fcUnblock(3f) man page for more information on fcUnblock.*
- *Optimizing code containing type conversions (including implicit type conversions) can cause some code to vectorize poorly or not at all.*

Error handling using data format conversions

*When you specify ERR= or IOSTAT= in an I/O statement, most errors reading REAL*16 values in IEEE mode do NOT cause a diagnostic message to be issued (except for calling stub routines for user-defined I/O). Instead, the program branches to the user-defined label, or if only an IOSTAT= expression was specified, it branches to the statement immediately following the active I/O statement.*

If you do not specify ERR= and IOSTAT= in the active I/O statement, then all errors except overflow and underflow cause a diagnostic message to be issued, and execution terminates.

When an overflow or underflow occurs, then an appropriate value is stored in the user's variable for a READ or in the I/O buffer for a WRITE. The library then activates an overflow or underflow. If you ignore underflow and overflow, program execution continues. If not ignored, these conditions cause the program to receive a signal (SIGFPE) that terminates the program unless trapped via the signal(3) routine. You can call errtrap to change the default handling of overflows and underflows. By default, programs ignore underflow (zero is substituted), and overflow causes a diagnostic message to be issued and terminates execution.

Reading or writing a reserved operand value (ROP) or Not a Number value (NaN) has no effect on the program's execution. The appropriate value for the target data format is generated. No traps or signals occur.

User-defined conversions

You can create your own conversion routines to convert a data file from a format that does not have a supplied conversion routine. To use your own routine, specify `user_defined` as the data format of the file in the `OPEN` statement or in a shell variable.

You must supply two functions for each FORTRAN data type: `REAL*4`, `INTEGER*2`, and so on. One function converts the data from the user-defined format to the appropriate CONVEX FORTRAN data type. The other function converts a CONVEX FORTRAN data type to the user-defined data type.

Table 15 shows a set of stub routines for user-defined data types. You only need to supply the routines that you actually use. If you forget to write a routine that is called, the stub routine prints an appropriate diagnostic message and halts execution of the program. This error message ignores `ERR=` and `IOSTAT=` clauses.

Each user-defined conversion routine is called as follows:

`NAME (source, target, bytecount)`

where `source` and `target` are pointers to a block of storage, and `bytecount` is a pointer to an integer value. The routine is expected to pick up the values through the `source` pointer and store the converted data values through the `target` pointer. `bytecount` is the number of data values the routine converts multiplied by the size of one item in bytes. For example, a call of the form

`CVT_INT4_TO_UD (SOURCE , TARGET , 40)`

converts 10 `INTEGER*4` values (40/4) from CONVEX binary integer format to the user-defined format. `SOURCE` points to the 10 values before conversion, and the user-supplied routine stores the values through `TARGET`.

The function must return 0 to indicate success and -1 to indicate failure. The function must also provide support for reserved values, infinity, overflows, and underflows as needed.

Sample conversion routine

The FORTRAN source code shown in Figure 2 converts INTEGER*4 binary values into one's complement binary format. No provision is made in this routine for the maximum negative integer that cannot be represented in one's complement format.

Figure 2
FORTRAN example
conversion routine

```
INTEGER FUNCTION CVT_INTEGER4_TO_UD (SOURCE, TARGET, BYTECOUNT)
INTEGER SOURCE(*), TARGET(*), BYTECOUNT
INTEGER N
N = BYTECOUNT / 4           ! COMPUTE AN ITEM COUNT

DO 10 I=1,N                  ! CONVERT N ITEMS
  TARGET(I) = SOURCE(I)      ! IF VALUE < 0, SUBTRACT 1
  IF (TARGET(I) .LT. 0) TARGET(I) = TARGET(I) - 1
CONTINUE

CVT_INTEGER4_TO_UD = 0       ! RETURN SUCCESS
RETURN
END
```

The C source code to accomplish the same task is shown in Figure 3.

Figure 3
C example conversion routine

```
int cvt_int_to_ud_ (source, target, byteCount)
  int * source, * target, * byteCount;
{
  int n;
  n = ( *byteCount) / 4; /* Compute an item count          */
  while ( n-- > 0) { /* Convert n items                    */
    *target = *source; /* Move the value                  */
    if (*target < 0) /* If the value is negative,                */
      (*target)-- /* Subtract one.                      */
    source++; /* Adjust the source & target pointers */
    target++; /* for the next value/destination          */
  }
  return (0); /* Return success.                          */
}
```

Note

A trailing underscore is required on the routine name if it is written in C.

Because some of the FORTRAN data types are not directly supported in C, it can be useful to declare the pointer's source and target as something else, such as a char or a struct pointer.*

Note

Only one user-defined data format is supported; therefore, a given program is restricted to one set of user-supplied conversion routines.

You can support several different user-defined data formats concurrently by modifying the program to save the current user-defined data format in a variable in COMMON, which can be queried by each conversion routine.

User-supplied conversion routine names

Table 15 lists user-defined conversion routine names.

Table 15
User-supplied conversion routine names

FORTRAN type	Routine names	
	FORTRAN-to-user-defined	User-defined-to-FORTRAN
INTEGER*1	cvt_integer1_to_ud	cvt_ud_to_integer1
INTEGER*2	cvt_integer2_to_ud	cvt_ud_to_integer2
INTEGER*4	cvt_integer4_to_ud	cvt_ud_to_integer4
INTEGER*8	cvt_integer8_to_ud	cvt_ud_to_integer8
REAL*4	cvt_real4_native_to_ud cvt_real4_ieee_to_ud	cvt_ud_to_real4_native cvt_ud_to_real4_ieee
REAL*8	cvt_real8_native_to_ud cvt_real8_ieee_to_ud	cvt_ud_to_real8_native cvt_ud_to_real8_ieee
REAL*16 [†]	cvt_real16_native_to_ud cvt_real16_ieee_to_ud	cvt_ud_to_real16_native cvt_ud_to_real16_ieee
COMPLEX*8	cvt_complex8_native_to_ud cvt_complex8_ieee_to_ud	cvt_ud_to_complex8_native cvt_ud_to_complex8_ieee
COMPLEX*16	cvt_complex16_native_to_ud cvt_complex16_ieee_to_ud	cvt_ud_to_complex16_native cvt_ud_to_complex16_ieee
LOGICAL*1	cvt_logical1_to_ud	cvt_ud_to_logical1
LOGICAL*2	cvt_logical2_to_ud	cvt_ud_to_logical2
LOGICAL*4	cvt_logical4_to_ud	cvt_ud_to_logical4
LOGICAL*8	cvt_logical8_to_ud	cvt_ud_to_logical8
CHARACTER	cvt_character_to_ud	cvt_ud_to_character

[†] REAL*16 is not supported in IEEE mode programs.

For all floating point data types (REAL, COMPLEX, and so on), the program's mode (NATIVE or IEEE) has been included as part of the name. This allows the runtime library to trap an unintended conversion.

*COMPLEX data types are really two adjacent REAL values. For COMPLEX*8, there are two adjacent REAL*4 values. Because a byte count is passed to the conversion routines, the conversion routines for COMPLEX*8 can usually be identical to the REAL*4 routines.*

Conversion using a shell variable

Each time an *OPEN* statement is executed for a unit number not previously opened, the program's environment is searched for the shell variable named *FOR_{mmm}OPEN*, where *mmm* is the unit number. *mmm* must be 3 digits long with leading zeros as needed. If this shell variable exists, the attributes it contains are applied to the file in question. If *FOR_{mmm}OPEN* is not found, the file attributes specified in the *OPEN* statement are used.

Any attributes not specified in the shell variable are taken from the *OPEN* statement. In the absence of an associated shell variable, the attributes specified with the *OPEN* statement are used.

All *OPEN* operations for that particular unit use those attributes specified in the shell variable.

Note

The data format attribute (*dataformat=xxxx*) is only valid if the file is opened for unformatted I/O. Otherwise, it is ignored.

Table 16 lists the attributes that you can specify in the shell variable.

Table 16
Shell variable attributes

Keywords & abbreviations	Legal values
BLANK BLNK	null, zero
BLOCKSIZE BLKSZ	<number>
CARRIAGECONTROL	FORTRAN, LIST, NONE
MAXREC	<number>
RECL	<number>
RECORDTYPE RT	FIXED, VARIABLE
DISPOSE DISP	KEEP, DELETE
POSITION POS	as is, REWIND, APPEND
DATAFORMAT DF	NATIVE, IEEE, vax-d, vax-g, CRAY, CRAYUB, user-defined

Note

When the data format attribute is specified, either in an `OPEN` statement or a shell variable, the source code must be compiled with CONVEX FORTRAN V6.0 or higher and linked with the CONVEX FORTRAN V6.0 or higher libraries. If a version of CONVEX FORTRAN prior to V6.0 is used, unpredictable behavior, including program aborts and errors, can result.

Use the following `cs`h command to specify a `FOR mmm OPEN` value, where mmm is a three-digit number from 000 to 255.

```
% SETENV FOR $mmm$ OPEN "RECL=256, DATAFORMAT=VAX-G"
```

The corresponding `sh` commands are as follows:

```
$ FOR $mmm$ OPEN="RECL=256, DATAFORMAT=VAX-G"
```

```
$ EXPORT FOR $mmm$ OPEN
```

The corresponding `COVUEshell` command is:

```
$SET COVUE ENVIRON : FOR $mmm$ OPEN = "RECL=256, DATAFORMAT=VAX-G"
```

You can use keyword abbreviations to reduce the total length of the shell variable's value. Keywords and their abbreviations can be given in either uppercase or lowercase. The shell variable name must be in uppercase.

Format specifications describe the format of data to be read or written and define any editing that is required. You can use any of the following format specification methods with formatted input/output (I/O) statements:

- The label of a FORMAT statement that contains the format, for example:

```
WRITE (6,50) A, B
50 FORMAT (I4)
```

- An integer variable assigned the label of a FORMAT statement, for example:

```
ASSIGN 50 TO L
WRITE (2,L) A1, A2
```

- A character array, character variable, or other character expression that specifies the format, for example:

```
READ (10, '(I4,I6)') L, M
```

- An asterisk that indicates list-directed I/O, for example:

```
WRITE (10, *) K, L, M
```

FORMAT statement

The nonexecutable FORMAT statement provides information necessary to produce the desired format for I/O statements. The FORMAT statement has the form:

```
sl FORMAT (flist)
```

where *sl* is a required statement label and *flist* is a nonempty format list.

Each item in the *flist* is of the form:

```
[r]ed ned [r]fs
```

where

r
represents the repeat count.

ed
is a repeatable edit descriptor. Repeatable descriptors indicate the type and layout of the next data value in the file. The repeatable descriptors are: I, O, Z, D, F, E, G, A, L, and Q.

ned
is a nonrepeatable descriptor. Nonrepeatable descriptors specify format characteristics such as spacing and skipping data that are not required. These descriptors are: H, X, P, T, TL, TR, SP, SS, S, BN, BZ, B, SU, R, slash (/), colon (:), *dollar sign* (\$), apostrophe (') and *asterisk* (*) descriptors.

fs
is a nonempty *flist*.

A repeatable edit descriptor has one of the following forms:

$[r]c[r]cw [r]cw.m [r]cw.d[Ee]$

$[r]cw.d[De] [r]cw.d.e$

where

r
is a repeat specification (unsigned integer constant) that indicates repetition of the descriptor *r* times in the format specification. The repeat specification *r* cannot be used with all descriptors. If you omit the repeat specification, the count defaults to 1. You must include at least one repeatable descriptor in the format specification for I/O statements that have one or more items in the I/O list.

c
is a format descriptor that may or may not be repeatable.

w
is an unsigned integer constant that indicates the field width in characters. A field containing only blank characters represents the value of zero. Leading blanks are not significant; other blanks are ignored or represented as zero depending on the value of the BLANK keyword when the file was connected.

m

is an unsigned integer constant that indicates the minimum number of characters, including leading zeros, that must appear within the field.

d

is an unsigned integer constant that indicates the number of characters to the right of the decimal point for real values.

E or D

identifies the exponent field.

e

is an unsigned, integer constant that indicates the number of characters to output as the exponent.

Not all of the previously identified terms are required for formatting. For instance, *e* can be used for formatting real values but is invalid for use with integer format descriptors, for example, *I*, *O*, *Z*. Do not use PARAMETER constants for the terms *r*, *w*, *m*, *d*, or *e*.

FORMAT control

When data transfer occurs, format control depends on information provided by the next format descriptor and the next item in the I/O list, if any. Generally, the format specification is interpreted from left to right, and elements in the I/O list are correlated with the corresponding repeatable edit descriptors. There are no corresponding list elements for the nonrepeatable descriptors. The I/O statement terminates if, during execution of the data transfer statement, a repeatable edit descriptor is encountered but there is no corresponding item in the I/O list. For example, the statement:

```
READ (*, '(I4,5F6.2)' ) K, X, Y
```

causes three values, not six, to be read using the descriptors *I4* and *F6.2*. The additional three *F6.2* descriptors are not used. If there is another item in the I/O list but no repeatable descriptor, however, control reverts to the beginning of the format specification, and a new record is started.

Example:

```
READ (5, '(I4,I6)' ) K, J, I, N
```

This statement causes values to be read in from character positions 1 to 4 and 5 to 10 of the current record and assigned to K and J, respectively. Control then reverts to the beginning of the format specification; values from character position 1 to 4 and 5 to 10 of the next record are read and assigned to I and N, respectively. Reversion to the beginning of the format list causes multiple records to be transferred. The end-of-file condition is flagged if there are insufficient records in the file to satisfy the execution of the input statement.

You can transfer data entirely from the descriptors to the external records. In this case, there is no corresponding item in the I/O list for the descriptors so format control communicates information directly to the record. You can use H and character constant descriptors to transfer data directly to the external record from the format specification. For example, the following statement outputs the characters 'CONVEX FORTRAN' to the file connected to unit 2:

```
WRITE (2, ' ("CONVEX FORTRAN" )')
```

There is no output list in the above WRITE statement. Because the format identifier is a character constant containing the format specification, the apostrophes in the format specification must be represented by two consecutive apostrophes in the format identifier.

Usually, a new I/O statement positions the file at the next record. *The use of \$ while writing causes suppression of a newline at the end of the current record.* The slash descriptor (/) terminates processing of the current record; the next record is used for the remaining descriptors.

Processing of repeatable edit descriptors positions the file after the last character transferred. This is also true of the H and apostrophe edit descriptors. Positioning left or right within the current record is accomplished by the X, T, TL, and TR descriptors.

Repeat count

You can use the descriptors A, O, Z, F, E, D, G, L, and I in a repetitive sequence by preceding the descriptor with an unsigned, integer constant that specifies the number of repetitions. For example, the following two statements are equivalent:

```
30 FORMAT (F6.0,F6.0,8X,F10.3,F10.3,F10.3)
30 FORMAT (2F6.0,8X,3F10.3)
```

You can also repeat a group of descriptors by enclosing the descriptors in parentheses and preceding them with an unsigned, integer constant that specifies the number of repetitions. The repeat count defaults to 1 when you do not specify the count. For example, the following two statements are equivalent:

```
30  FORMAT (F6.0,F6.0,8X,F10.3,E12.4,5X,F10.3,  
        ^ E12.4,5X,F4.0)  
30  FORMAT (2F6.0,8X,2(F10.3,E12.4,5X),F4.0)
```

Descriptors

Field and edit descriptors in CONVEX FORTRAN are grouped into the following categories:

- Character (A)
- Editing, character constants, and Hollerith constants (apostrophe, *asterisk*, T, TL, TR, P, Q, *dollar sign*, colon, slash, X, H, B, BN, BZ, S, SP, SS, SU, R)
- Integer (I, O, Z)
- Logical (L)
- Real and complex (D, E, F, G)

A descriptor

The A descriptor transfers character or Hollerith values and is represented by

A[w]

In an input statement, the A field descriptor transfers *w* characters from the external record and assigns them to the corresponding I/O list element. If the *w* field is not specified, the size equals the length of the character variable, character substring reference, or character array element. For numeric I/O list elements, the size depends on the data type, as shown in Table 17.

Table 17
 Character assignment for
 numeric I/O list elements

I/O list elements	Maximum no. of characters
<i>LOGICAL*1</i>	<i>1</i>
<i>LOGICAL*2</i>	<i>2</i>
<i>LOGICAL*4</i>	<i>4</i>
<i>LOGICAL*8</i>	<i>8</i>
<i>INTEGER*1</i>	<i>1</i>
<i>INTEGER*2</i>	<i>2</i>
<i>INTEGER*4</i>	<i>4</i>
<i>INTEGER*8</i>	<i>8</i>
<i>REAL*4</i>	<i>4</i>
<i>REAL*8 (DOUBLE PRECISION)</i>	<i>8</i>
<i>REAL*16</i>	<i>16</i>
<i>COMPLEX</i>	<i>8</i>
<i>COMPLEX*16 (DOUBLE COMPLEX)</i>	<i>16</i>

If w is less than the size of the *iolist* item, on input the characters are stored left-justified and padded on the right with blanks. If w is greater than the size of the *iolist* item, on input the rightmost characters are stored in the variable.

On output, the value is right-justified in the field and w characters from the entity are written to the record. If w is greater than the number of characters in the entity, leading blanks are added to right-justify the value. If w is less than the number of characters in the *iolist* item, only the leftmost w characters are written.

The following example illustrates reading into a CHARACTER*5 variable.

Example:

Format code	External field	Internal value
A	CONVEXCOMPUTER	CONVE
A4	CONVEXCOMPUTER	CONV^
A14	CONVEXCOMPUTER	PUTER

The following example illustrates writing from a CHARACTER*10 variable.

Example:

Format code	Internal value	External field
A	MY^EXAMPLE	MY^EXAMPLE
A4	MY^EXAMPLE	MY^E
A14	MY^EXAMPLE	^^^^MY^EXAMPLE

Apostrophe (') descriptor

The apostrophe descriptor has the form of a character constant. Characters that are enclosed within a pair of apostrophes are written to the record. The width of the field equals the number of characters contained within (but not including) the delimiting apostrophes. Use two consecutive apostrophes to produce a single apostrophe. For example, the statements

```
WRITE (6,100)
100 FORMAT ('THE^^'CONVEX''^COMPUTER')
```

produce

```
THE 'CONVEX' COMPUTER.
```

Asterisk (*) descriptor

CONVEX FORTRAN supports the asterisk descriptor under the -cfc option. The asterisk descriptor delimits literal text and behaves identically to the apostrophe descriptor. Use two consecutive asterisks to produce a single asterisk.

H descriptor

The H descriptor writes a literal string to a record. You can use the H descriptor for output editing as an alternative to apostrophe editing.

This descriptor has the following form:

$nHc...c$

The H descriptor writes the n characters immediately following the letter H, including apostrophes and quotation marks. The $c...c$ represents the actual characters to be written. For example, the statements

```
WRITE (6,10)
10 FORMAT (17HENTER 'FILE' NAME)
```

produce

```
ENTER 'FILE' NAME
```

L descriptor

The L descriptor formats logical variables and has the form:

Lw

where w indicates the field width for formatting logical variables.

Optional blanks, optionally followed by a decimal point, a T (t, .T., .t.) for true or an F (f, .F., .f.) for false, constitute the input field. The T or F can be followed by additional characters, for example, .TRUE. or .FALSE., in the field. On input, a field containing only blanks is read as false.

On output, the record contains $w - 1$ blanks, followed by a T or F depending on the value of the corresponding I/O list element.

Input examples:

Format code	External field	Internal value
L2	T60	.TRUE.
L7	^^FALSE	.FALSE.
L7	1234567	Error: invalid

Output examples:

Format code	Internal field	External field
L1	.TRUE.	T
L3	.FALSE.	^^F

I descriptor

The I descriptor provides integer formatting. It has one of the forms:

Iw

or

Iw.m

where

w

is an unsigned, positive integer constant that specifies that the field to be edited is *w* characters wide.

m

is an unsigned, integer constant that specifies the minimum number of digits for output only, including leading zeros if necessary.

During input, the processor transfers *w* characters from the record in integer representation and stores the integer values in the corresponding I/O list elements. Both forms of the I descriptor are treated identically during input. On input, leading blanks are not significant; nonleading blanks are interpreted according to the BLANK specifier in the OPEN statement or the BZ or BN descriptor. If the field contains only blanks, the value is zero. A plus sign (+) or no sign indicates a positive value; a leading minus sign indicates a negative value.

During output, the processor formats the value of the I/O list element and outputs it in a field w characters wide, right-justified. Leading blanks are added, if needed, to fill the field. If the field specified is too small for the value, the field is padded with asterisks (*). If you specify m , the external field contains up to m characters; if necessary, the processor inserts leading zeros to pad to m . The value of m must not be greater than the value of w . If m equals zero and the value of the entity is zero, the output field contains blank characters. The minus sign (-) precedes a negative integer; by default a plus sign (+) does not precede a positive integer. You must include a space for a minus sign for negative integers in the w term.

Input examples:

Format code	External field	Internal value
I3	760	760
I4	^^^^	0
I4	-760	-760
I5	760^^	760
I5	760^^	76000 (blanks interpreted as zeros)
I5	7.60^^	Error: decimal

Output examples:

Format code	Internal value	External field
I4	760	^760
I8.4	760	^^^^0760
I3	-760	***
I4	0	^^^^
I4.0	0	^^^^

***O* descriptor**

*The **O** descriptor transfers unsigned octal values. The descriptor has the form:*

Ow [.m]

where

w

is an unsigned, positive integer constant that specifies the field to be edited is *w* characters wide.

m

is an unsigned, integer constant that specifies the minimum number of digits for output, including leading zeros if necessary.

On input, format code *o* transfers *w* characters from the external field and assigns them as an octal value to the corresponding I/O list element. On output, if *m* is specified and the external field contains fewer digits than *m*, the remaining positions are padded with zeros on the left. You can only use the numerals 0 through 7 in the external field; you cannot use a decimal point (.), a sign (+ or -), or an exponent field.

Input examples:

Format code	External field	Internal octal
<i>o3</i>	523	523
<i>o4</i>	23176	2317
<i>o4</i>	2.317	Error: decimal
<i>o4</i>	-1234	Error: signed

Output examples:

Format code	Internal decimal value	External value
<i>o6</i>	4095	^^7777
<i>o6</i>	-4095	*****
<i>o3</i>	4095	***
<i>o4.3</i>	8	^010

***Z* descriptor**

The *z* descriptor transfers unsigned hexadecimal values and is represented as

Zw [.m]

where

w

is an unsigned, positive integer constant that specifies the field to be edited is *w* characters wide.

m

is an unsigned, integer constant that specifies the minimum number of digits for output, including leading zeros if necessary.

On input, descriptor *Z* transfers *w* characters from the external field and assigns them as a hexadecimal value to the corresponding I/O list element. On output, if *m* is specified and the value has fewer digits than *m*, the remaining positions are padded with zeros on the left. You can use only the numerals 0 through 9 and the letters A (a) through F (f) in the external field; you cannot use a decimal point (.), a sign (+ or -), or an exponent field.

Input examples:

Format code	External field	Internal hex value
Z3	9A1	9A1
Z3	9A1B	9A1
Z3	9A.1	Error: decimal

Output examples:

Format code	Internal decimal value	External value
Z4	4095	^fff
Z5	-1	*****
Z6.4	4095	^^0fff
Z2	4096	**

F descriptor

The F descriptor provides formatting of real numbers. It has the following form:

Fw.d

where

w

is an unsigned, positive integer constant that specifies the field to be edited is *w* characters wide.

d

specifies the number of digits in the fractional (right of the decimal) part of the real number.

During input, the processor transfers *w* characters from the external field and stores the real values in the corresponding I/O list elements. The input field consists of an optional sign followed by a string of digits that can contain a decimal point. If the field has a decimal point, the *d* term has no effect; the location of the explicit decimal overrides the location specified by the field descriptor. If you omit the decimal point and the exponent, the rightmost *d* digits are interpreted as the fractional part of the field with leading zeros assumed if necessary.

On input, leading blanks are not significant; nonleading blanks are interpreted according to the BLANK specifier or the BZ or BN descriptor. If the field contains only blanks, the value is zero. The processor treats a plus sign (+) or no sign as a positive value; a minus sign (-) indicates a negative value.

During output, the processor transfers the value of the I/O list element rounded to *d* decimal positions and outputs it in a field *w* characters wide, right-justified. *w* must include a space for a minus sign when necessary, at least one digit to the left of the decimal point, the decimal point, and *d* digits to the right of the decimal, for example, at least equal to or greater than $d + 3$. Leading spaces are added, if needed, to fill the field.

Input examples:

Format code	External field	Internal value
F8.5	1234567^	12.34567
F8.5	12345.67	12345.67
F8.0	-1.23E-3	-.00123
F8.5	123456789	123.45678

Output examples:

Format code	Internal value	External field
F9.4	123.456789	^123.4568
F5.2	123.456789	*****
F6.3	+1.12	^1.120
F6.3	-1.12	-1.120

If the value is too large for the field, asterisks (*) are output. In native format, if the sign is 1 and the exponent is 0, *Rop* (reserved operand) is output followed by the fraction in hexadecimal. In IEEE format, if the exponent is all ones and the fraction is nonzero, NaN (not a number) is output followed by the fraction in hexadecimal; if the exponent is all ones and the fraction is 0, Inf (infinity) is output.

E and D descriptors

The E and D descriptors are functionally identical. Both transfer real values in exponential form and edit external REAL, DOUBLE PRECISION, or complex data. These descriptors differ only in the exponent symbol they use. The E and D descriptors have the following forms:

Ew.d, *Ew.d.e*, or *Ew.dEe*

Dw.d, *Dw.d.e*, or *Dw.dEe*

where

w

is the width of the field containing the real number. The width is the count of all characters in the field, including sign (if any), decimal point, and exponent.

d

is the fractional part of the field that contains *d* digits.

E or D

identifies the exponent part that contains e digits (no effect on input).

e

indicates the number of digits in the exponent.

On input, the descriptors read w characters from the external field and assign them as a real value to the corresponding I/O list element. The values being read consist of a string of digits with an optional decimal point. When the decimal point is included, the d term has no effect. When the decimal is omitted, however, the least significant d digits of the string, not including the exponent, are considered the fractional part of the value.

On output, the E and D descriptors transfer the value of the I/O list element rounded to d decimal positions and output it to a field w characters wide, right-justified. w must include space for a minus sign (when necessary), the decimal point, d digits to the right of the decimal, a two- or three-digit exponent (depending on whether the I/O list item is REAL*4 or REAL*8), E or D, and the sign of the exponent.

All data values, including REAL*16, can be used with the E and F formats.

In a data value, the exponent can be a signed-integer constant, or E or D followed by zero or more blanks, followed by an optional signed-integer constant. You can use the legal characters—digits 0 through 9, decimal point, plus, minus, E, D, and blank. With the descriptors in the form of $Ew.d$, $Ew.d.e$ or $Ew.dEe$ ($Dw.d$, $Dw.d.e$, $Dw.dEe$), the value of the next item in the list has the following form:

$$[\pm] [0] .x_1, x_2 \dots x_d \text{ exp}$$

where

\pm

signifies a plus or minus sign; the plus sign is optional for a positive value.

0

an optional leading zero.

$x_1, x_2 \dots x_d$

are the d most significant digits of the value after rounding.

exp

is a decimal exponent that is of the form $E \pm z_1 [z_2] \dots z_n$ where z is a digit and there are e digits in the exponent.

Input examples:

Format code	External field	Internal value
E4.3	.625	.625
E6.1	.78-01	.078
D6.3	-.62D4	-6200
D6.3	123456	123.456
E4.3	62567	6.256

Output examples:

Format code	Internal value	External field
D11.4	-6250.	-0.6250D+04
E10.3	625	^0.625E+03
E10.4	0.4568	0.4568E+00
E10.3	0.4568	^0.457E+00
E5.3	24.53	*****
E7.2.1	2.718	0.27E+1

If the value is too large for the field, asterisks (*) are output. The default length for the exponent field for *REAL*16* numbers is three. In native format, if the sign is 1 and the exponent is 0, *Rop* (reserved operand) is output followed by the fraction in hexadecimal. In IEEE format, if the exponent is all ones and the fraction is nonzero, *NaN* (not a number) is output followed by the fraction in hexadecimal; if the exponent is all ones and the fraction is 0, *Inf* (infinity) is output.

G descriptor

The G descriptor edits external single-precision, double-precision, quad-precision, or complex data. It has the following form:

Gw.d, *Gw.d.e* or *Gw.dEe*

where

w

is a nonzero, unsigned, integer constant that indicates the field width in characters.

d

is a nonzero, unsigned, integer constant that indicates the number of significant digits to be printed. On output, when the range of the value to be printed forces E style editing (see below), *d* specifies the number of characters to the right of the decimal point. When the range of an output value allows F editing, *d* specifies the number of significant digits to print.

E

identifies the exponent field.

e

is an unsigned, integer constant that indicates the number of digits in the exponent.

Input editing is identical to F, E, and D editing. You can use the G descriptor when you are not certain that the values you are using can be adequately represented by the F descriptor because of their magnitude—either too large or too small.

On output, the G descriptor uses either the F or E style of editing depending on the magnitude of the value. If the value can be represented using the F format without loss of significant digits, F is chosen; otherwise, E is chosen.

Assume *M* is the magnitude of the data in the field. If *M* is less than 0.1 or greater than or equal to 10^{**d} , the output editing of *Gw.d* or *Gw.dEe* is the same as that of *kPEw.d* and *kPEw.dEe*, respectively, and *k* is the scale factor currently in effect. If *M* is greater than or equal to 0.1 or less than 10^{**d} , however, the F mode of editing is used with output of the four-character exponent field as four blanks after the value. The scale factor has no effect and the value of *M* determines the editing as shown in Table 18.

Table 18
Data conversion based on magnitude

Magnitude of data	Conversion equivalence
$0.1 \leq M < 1.0$	$F(w-n).d,n('')$
$1.0 \leq M$ or < 10.0	$F(w-n).(d-1),n('')$
.	.
.	.
.	.
$10^{** (d-2)} \leq M < 10^{** (d-1)}$	$F(w-n).1,n('')$
$10^{** (d-1)} \leq M < 10^{**d}$	$F(w-n).0,n('')$

The value $n('')$ specifies that four or $e + 2$ spaces are to follow the numeric data representation; n is 4 for $Gw.d$ and $e + 2$ for $Gw.dEe$. Be sure the w term is large enough to include a sign, if necessary, a decimal point, d digits to the right of the decimal and either a 4-character or an $(e + 2)$ -character exponent. Thus, you must make w equal to or greater than $d + 7$ or $d + 5 + e$.

Input examples:

Format code	External field	Internal value
G8.5	^1234567	12.34567
G8.5	12345.67	12345.67
G8.0	-1.234-3	.001234

Output example:

Format code	Internal value	External field
G13.6	-1234	^-1234.00^^^^
G13.6	0.01234	^0.123400E-01
G13.6	1.23456789	^^1.23457^^^^
G10.4	15.65	^15.65^^^^
E10.4	15.65	0.1565E+02
F10.4	15.65	^^^15.6500

If the value is too large for the field, asterisks (*) are output. The default length for the exponent field for $REAL*16$ numbers is three. In native format, if the sign is 1 and the exponent is 0, Rop (reserved operand) is output followed by the fraction in hexadecimal. In IEEE format, if the exponent is all ones and the fraction is nonzero, NaN (not a number) is output followed by the fraction in hexadecimal; if the exponent is all ones and the fraction is 0, Inf (infinity) is output.

B descriptors

The B descriptors operate only during execution of the input statements and affect only the numeric descriptors I, O, Z, F, E, D, and G. The BN and BZ descriptors supersede the default interpretation of blanks while the *B descriptor causes return to the default mode of blank interpretation*. Their forms and meanings are:

B

reverts to default interpretation.

BZ

interprets blanks as zeros.

BN

interprets blanks as nulls.

When execution of a formatted input statement begins, blanks can either be interpreted as zeros or ignored, depending on the value of the BLANK specifier in the OPEN statement. The BLANK specifier can be omitted from the OPEN statement, and the OPEN statement itself can be omitted. Additionally, the -vfc compiler option and use of the COVUEshell can effect the default value of BLANK. Table 19 shows the value of the BLANK specifier under various definition states.

Table 19
BLANK specifier defaults

OPEN/BLANK definition		BLANK specifier value	
OPEN statement	BLANK specifier	Under COVUEshell or -vfc	-vfc and COVUEshell not used
Omitted	NA	BLANK = 'ZERO'	BLANK = 'NULL'
Included	Omitted	BLANK = 'NULL'	BLANK = 'NULL'

The BN descriptor causes the processor to treat all embedded blank characters as nulls in subsequent input fields for the current statement. When the processor encounters the BN specifier, it treats the input field as though the embedded blanks have been moved to the position of leading blanks; the remainder of the field becomes right-justified. The processor assigns the value of zero to a field containing only blanks. If you specify the BZ descriptor, the processor treats all embedded and trailing blanks in subsequent numeric input fields as zeros.

For example, if a file connected to unit 5 contains the record

```
^^^5^-500^^^3^056
```

and the BLANK specifier has a NULL value or the BN descriptor is specified, the statements:

```
READ (5, '(I4, I7, I6)') L, M, N
READ (5, '(BN, I5, I7, I6)') L, M, N
```

assign the value of 5 to L, the value of -500 to M, and the value of 3056 to N. The processor ignores all blanks. If the BZ descriptor is indicated by:

```
READ (5, '(BZ, I5, I7, I6)') L, M, N
```

the values assigned become L = 50; M = -50000; N = 30056. The processor treats all nonleading blanks as zeros. If another input statement refers to unit 5, blank interpretation returns to the default value.

The descriptor B causes a return to the default mode of blank interpretation ('NULL') and is identical to BN. For example, change the previous example to include a B descriptor:

```
READ (5, '(BZ, I5, I7, B, I6)') L, M, N
```

The value of 3056 is assigned to N rather than 30056, because the B descriptor returns blank interpretation to default mode.

P descriptor

The P descriptor specifies a scale factor for real and complex values. The P descriptor can be used on input or output and applies to the F, E, D, and G edit descriptors. The P descriptor has the form:

$$nP$$

where n is an optional integer constant that specifies the number of positions, to the left or right, that the decimal point is to be moved. The value can be signed. Its default value is 0.

If no P descriptor is specified, a scale factor of 0 is assumed. Once specified, a scale factor remains in effect within a FORMAT statement until another P descriptor is encountered. The following example uses a scale factor of 0 for the first format descriptor and a scale factor of 2 for the two remaining descriptors.

Example:

```
PRINT 50, D
50 FORMAT (F8.2, 2PF8.2, F6.2)
```

On input, the scale factor (with the F, E, D, or G descriptors) causes the externally represented number to be multiplied by 10^{*-n} before it is assigned to the I/O list element.

Examples:

Format code	External field	Internal value
3PF7.4	56.789^	.056789
-3PE6.3	56.789	56789.

On output, when you use the scale factor with the F descriptor, the externally represented number equals the internally represented number multiplied by 10^{**n} . When the scale factor is used with E or D, the nonexponent part of the constant is multiplied by 10^{**n} and n is subtracted from the exponent. With G, if the F style of formatting is used, the scale factor is ignored; otherwise, the effect is the same as E editing.

Examples:

Format code	Internal value	External field
-1PF7.3	58.967	^^5.897
2PE10.3	890.11	^89.01E+01

If you use a scale factor when an external field has an explicit exponent, for example, 5.E02, the processor ignores it; in this case, 500 is stored regardless of the scale factor.

S descriptors

The S descriptor can be used to control optional plus (+) characters in numeric output *or to cause integer values to be interpreted as unsigned during output conversion*. If you do not use an S descriptor, positive values do not have leading plus signs. The S, SP, and SS descriptors act only during statement execution and only with I, F, E, and D editing descriptors. *The SU descriptor only affects integer values*. The descriptors have the following forms:

S

reverts to normal interpretation.

SP

adds a plus sign (+).

SS

suppresses plus signs.

SU

outputs integer values as unsigned values.

The *SP* descriptor forces a plus sign during output for all subsequent positive *I*, *F*, *D*, *E*, and *G* values within the format specification. Include space for the plus sign in the numeric fields. When you use the *SS* descriptor, the processor suppresses leading plus sign characters from any position where the plus sign is optional. The *S* descriptor returns the normal plus sign handling option to the processor. For example, if *L* = +5, *M* = 100, *N* = -10, *I* = 50, *J* = 6000, and *K* = -450, the statements

```
WRITE (10,30) L, M, N, I, J, K
30  FORMAT (SS,I2,I5,SP,I4,I4,S,I5,I5)
```

write the record as:

```
^5^^100^-10^+50^6000^-450
```

The SU descriptor causes integer values to be interpreted as unsigned during output conversion. SU remains in effect until another sign-control specifier is encountered or until format interpretation is complete. It has no effect on input. Radix and unsigned specifiers can be used to format a hexadecimal dump as follows:

```
2000 FORMAT (SU, 16R, 8I10.8)
```

R descriptor

The R descriptor changes the radix for integer I/O. This descriptor applies only to integers (I descriptor) and must not be used with other descriptors. The R specifier has the form

[n]R

*where $2 \leq n \leq 36$. The default value is 10. Omitting *n* restores the default decimal radix. The radix specified by *R* remains in effect until another radix is specified or until format interpretation is complete.*

Example:

```
I = 15
WRITE (6,10) I, I, I
10  FORMAT (16R,I4,8R,I4,R,I4)
```

produces

```
^^^F^^17^^15.
```

X descriptor

The X descriptor sets the position in a record and has the form:

nX

where *n* indicates the number of character positions to move forward (skip over) from the current position in the file. The value of *n* must be greater than or equal to 1. The default is 1.

The X descriptor is functionally identical to the TR descriptor. When you use the X descriptor, *n* indicates that the next *n* characters are to be skipped. The character following the number of skipped positions is transmitted. For example, the statements

```
WRITE (*,200) 450, 8921
200 FORMAT (2X,I3,3X,I4)
```

insert two blanks before 450 and three blanks before 8921.

The X format descriptor cannot in itself change the length of a record. For example, the statements

```
WRITE (6,100) I
100 FORMAT (I4,X)
```

produce a four-character record followed by a newline, not a trailing blank.

T descriptors

The T (tab) descriptors control forward and backward positioning within a record for input or output of characters. These descriptors let you skip portions of a record or reread portions of a record. The T descriptors are T, TR, and TL.

The T descriptor has two forms; the first form is

Tn

where n specifies the absolute position within the record. This form indicates transmission of characters at position n . For example, if a file connected to the designated input unit contains the record:

```
^2.5^200^^40
```

then execution of the statements

```
      READ (*,35) A, B
35    FORMAT (T2,F3.0,T11,F3.0)
```

assigns A the value of 2.5 (positions the file at character 2 and reads the next 3 characters according to format specification F3.0) then assigns the value of 40 to B (positions the file at character 11 of the record and reads the next 3 characters).

On output, for example, the statements

```
      PRINT 50
50    FORMAT (T10,'MY',T13,'EXAMPLE')
```

output MY at position 10 and EXAMPLE at position 13.

Another form of the T descriptor is

T or nT

which causes tabbing to the next (or nth) 8-column tab stop. You can, therefore, align columns of alphanumeric without counting. For example, the statements

```
      READ (5,50) K,N
50    FORMAT (T,I4,2T,I3)
```

cause K to be read starting in character position 8 of the current record; the value for N is read starting in position 24 of the current record.

The second of the T-series descriptors has the following form:

TLn

where n , an unsigned, integer constant, indicates that the record is repositioned n characters left (backwards) from the current position in the record. The default is 0. For example, if the external record is 1.2345, the statements:

```
      READ (5,20) X, I
20    FORMAT (F6.0,TL4,I3)
```

produce X = 1.2345 and I = 234.

The third T-series descriptor has the following form:

TR n

where n is an unsigned integer that specifies the number of characters to move right (forward) from the current position in the record. The default is 0. The TR and X field descriptors are identical. For example, assume the external record is as follows:

```
12.345^^^^123
```

The statements

```
      READ (5,20) X, I
20    FORMAT (F6.0,TR4,I3)
```

produce X = 12.345 and I = 123.

The T descriptor cannot itself change the record length. For example, the statements

```
      WRITE (6,10) I
10    FORMAT (I4,TR10)
```

produce a 4-character record with no trailing blanks.

Dollar sign (\$) descriptor

The \$ descriptor suppresses the newline at the end of the current record of a formatted sequential-access write. (In an input statement, the \$ descriptor is ignored.) For terminal I/O, a typed response follows the output on the same line. For example, the statements

```
WRITE (*,'(" enter value for x: ",$)')
READ (*,*) x
```

write ^enter value for x:^ to the output device with the cursor positioned one space to the right of the colon.

Q descriptor

The *Q* descriptor determines the number of unread characters in the current record. It is represented by:

Q

Example:

```
READ(4,100) J,MYEXAM, (ISAM(I), I=1,MYEXAM)
100 FORMAT(I5,Q,80A1)
```

This example reads the first field into variable *J*, stores the number of remaining characters in *MYEXAM*, and causes transfer of that number of characters to the character array, *ISAM*. If you place *Q* first in the format specification, you can determine the actual length of the record.

In an output statement, the descriptor *Q* causes the corresponding I/O list element to be skipped.

Colon (:) descriptor

The colon (:) descriptor ends format control when no items remain in the I/O list. If items remain in the I/O list, the colon descriptor has no effect. The following example:

```
M = 15
WRITE (10,40)M
40 FORMAT (I2, :, ' SAMPLE', I3)
```

writes 15 only, ending format control at the colon. Change the statements slightly, however, and the colon descriptor has no effect.

Example:

```
M = 15
N = 500
WRITE (10,40) M, N
40 FORMAT (I2, :, ' SAMPLE', I4)
```

writes 15 SAMPLE 500; the colon descriptor is ignored because items remain in the I/O list.

Slash (/) descriptor

The slash (/) descriptor indicates the end of data transfer for the current record. For example, the statements:

```
      READ (10,50) L, M, N
50 FORMAT (I2/I4,I3)
```

cause L to be read from the first record, and M and N from the second record.

During input, use sequential slashes to indicate bypassing of records. The first slash indicates the end of input for the current record; subsequent slashes skip records. When you use the slash on a unit connected for sequential access, the remainder of the current record is skipped and the file is positioned at the beginning of the next record. On direct access, 1 is added to the record number and the processor reads that record.

On output, slashes can be used to create empty records. The first slash indicates end of output for the current record; subsequent slashes produce empty records.

Default field descriptor values

If you do not specify a field width value with the field descriptors I, O, Z, L, F, E, D, G, or A, default values for *w*, *d*, and *e* are supplied based on the data type of the I/O list element. These default values are shown in Table 20.

Table 20
Default field descriptors

Field descriptor	List element type	w	d	e
I, O, Z	INTEGER*1, LOGICAL*1	7		
I, O, Z	INTEGER*2, LOGICAL*2	7		
I, O, Z	INTEGER*4, LOGICAL*4	12		
I, O, Z	INTEGER*8, LOGICAL*8	23		
O, Z	REAL*4	12		
O, Z	REAL*8	23		
O, Z	REAL*16	44		
L	LOGICAL	2		
F, E, G, D	REAL, COMPLEX*8	15	7	2
F, E, G, D	REAL*8, COMPLEX*16	24	15	3
F, E, G, D	REAL*16	42	33	3
A	LOGICAL*1, INTEGER*1	1		
A	LOGICAL*2, INTEGER*2	2		
A	LOGICAL*4, INTEGER*4	4		
A	LOGICAL*8, INTEGER*8	8		
A	REAL*4, COMPLEX*8	4		
A	REAL*8, COMPLEX*16	8		
A	REAL*16	16		
A	CHARACTER*n	n		

Comma field separator on input data

A comma between numeric fields overrides the width specified in the field descriptor. Because you can use a comma to end a field, you can avoid padding the input field, which makes entering data from a terminal keyboard easier. A comma field separator can be used with the numeric descriptors (I, O, Z, F, E, D, G, and L).

Example:

```
READ (5,100) I,K
100 FORMAT (2I4)
```

Record:

```
1,2
```

Result:

$I = 1$

$K = 2$

The following constraints apply:

- Two successive commas constitute a null field.
- You cannot use a comma to end a field that is controlled by an A or a character constant field descriptor. If the record reaches its physical end before w characters are read, short-field termination occurs and the characters you input are assigned successfully. Trailing spaces are appended to fill the corresponding I/O list item.

Runtime formats

Format specifications, called runtime formats, can be stored in character variables, character substrings, and character expressions, and in character arrays, *numeric arrays and numeric array elements*. *Numeric arrays and numeric array elements are treated as Hollerith constants.*

You can define or modify a runtime format during program execution. A runtime format is similar to a FORMAT statement but does not have a label or the word FORMAT.

Example:

```
INTEGER*8 IFMT
CHARACTER*8 SFMT
IFMT = 8H(2X, I12)
SFMT = '(2X, I12) '
WRITE (6, IFMT) I
WRITE (6, SFMT) I
```

Cray-style asterisk descriptors can be used in runtime formats in all modes; the `-cfc` option is not required when asterisks appear in runtime formats.

Variable formats

A variable format has an expression, enclosed in angle brackets, that is computed each time it is encountered during format scanning. The expression has the form:

<expression>

The *expression* in the angle brackets can be used in a FORMAT statement wherever you can use an integer, except as the character count of a Hollerith (H) descriptor.

A variable expression in a format statement is subject to the following rules:

- If the expression is not an integer, it is converted to an integer before use.
- Any valid FORTRAN expression can be used, including function calls and dummy argument references.
- The value of the expression must conform to the restrictions on magnitude that apply to its use in a format.
- A variable expression is not allowed in a runtime format.

Do not perform I/O operations within a function call used in a variable format expression; a runtime error will occur.

Example:

```
C TEST OF D AND E DESCRIPTORS WITH REPEAT COUNT
1 FORMAT (<J+2>D10.4,<J/2+1>E10.4)
2 FORMAT (<J+2>D10.4.2,<J/2+1>E10.4.2)
3 FORMAT (1X,'Message = ',A<LEN (STRING)>)
J = 2
READ (5,1) A,B,C,D,E,F
WRITE (6,2) A,B,C,D,E,F
WRITE (*,3) STRING
STOP
END
```

Note that variable formats are *not* allowed in runtime format strings, as shown in the following example.

Example

```
WRITE (*, ' (1X, "Message = ",A<LEN (STRING)>)' ) !
INVALID!!
```

Attempting this will result in a runtime error.

List-directed formatting

List-directed formatting transfers data based on the data type of the entity. A list-directed I/O statement uses an asterisk (*) as the format indicator. For example, the following statement:

```
READ (5, *) J, M, L
```

reads three fields from unit 5 and assigns integer values to the variables J, M, and L.

The list-directed record is a sequence of values and value separators. A value is generally a constant but can also be a null value, or the value can have the form

r^*c or r^*

where

r

is an unsigned, nonzero, integer constant that represents the repeat count.

r^*c

represents successive appearances of the constant c . You can enter a repeat count to assign a value to more than one entity with r^*c .

r^*

repeat count with an empty constant (null value). A null value indicates that the value of the corresponding I/O entity is to remain unchanged.

Separators divide the values in each list-directed record. A value separator is a blank, comma, or a slash optionally enclosed by blanks. Normally, the blanks are considered part of a value separator. In the following cases, the blanks are not considered part of a value separator:

- Leading blanks in the first record, unless followed by a slash or comma
- Blanks embedded in a character constant

List-directed input

You can use list-directed input from any file that allows formatted input. The data type of the constant, which can be logical, integer, real, complex, or character, determines the data type of the value, as well as the translation from external to internal form. A character list element must correspond to a

character constant; likewise, a numeric element must correspond to a numeric constant. If the data type of the external numeric field does not match the data type of the numeric list item, the external value is converted according to the rules for conversion on assignment (refer to Chapter 7, "Assignment statements"). Input fields are separated by blanks, commas, or slashes.

The format of a complex value is left parenthesis followed by a numeric value, a comma, another numeric value (an ordered pair of numeric fields separated by a comma), followed by a right parenthesis. The processor ignores one or more blanks around either parenthesis or the comma. The end of record can occur between the real part and the comma or between the comma and the imaginary part.

Character input

Character constants for list-directed input are usually enclosed in apostrophes. Character constants can span record boundaries.

Embedded blanks, commas, and slashes within a character string are not considered separators. To include an apostrophe as part of a character string, use two consecutive apostrophes without an intervening blank or end of record.

The processor transfers the leftmost characters read, either truncating the constant to fit in the list item or filling it on the right with blanks.

CONVEX FORTRAN allows input of a string not enclosed in quotes. The string must not start with a digit and cannot contain a separator consisting of a right or left parenthesis, or blank (space or tab). A newline ends the string unless escaped with a slash (\). Any string not meeting these restrictions must be enclosed in single or double quotes.

Nulls and slashes

You can specify a null value for a list item with a comma or with r^* in the external record. No characters between successive value separators or no characters preceding the first value separator indicate a null field. When assigning a null for the first value, you can use one comma; for a subsequent null, use two consecutive commas.

A null value does not alter the value of the corresponding input list item.

When the processor encounters a slash on list-directed input, it skips the rest of the I/O list items and ends the READ statement. Those items skipped retain their original values.

Namelist-directed input formatting

To assign input values for a namelist-directed READ, you must delimit the input record (or records) with a dollar sign (\$). Namelist input has the following form:

```
$nlgrpname [ent = value [, ] ] ... $[END]
```

where

\$

indicates the beginning and end of input. You can use the ampersand (&) rather than the \$.

nlgrpname

is the name defined for the entities contained in the namelist.

ent

is a namelist entity. The entity can be a variable, an array name, a subscripted variable, a variable with a substring, or a subscripted variable with a substring.

value

is a constant, a list of constants, or a repetition of constants or null values.

END

is an optional delimiter indicating no more input.

Use constant values for assigned values, array subscripts, and substring specifiers; you cannot use PARAMETER constants.

You can use any data type. Conversion (following rules of arithmetic assignment) is performed if the data type of a namelist entity and its assigned constant value do not match. Conversion between numeric and character data is not allowed.

For the NAMELIST statement

```
NAMELIST/SAM/ NAME, EXAM1, EXAM2, EXAM3
```

the following example shows how to input data to the namelist entities. You can assign the values in any order.

```
$SAM NAME='TESTA', EXAM1=5.2, EXAM2=6.78,  
EXAM3=10.0 $
```

Several acceptable formats for entering input exist. For example, you can also enter the previous input as

```
$SAM^NAME='TESTA'^EXAM1=5.2^EXAM2=6.78^EXAM3=10.0  
$END
```

You can use commas, tabs, and spaces as valid separators in the list of value assignments, and input can begin in any column. You cannot use nonblank control characters in column 1.

The previous example assigns values to all of the namelist entities associated with *SAM*; however, you do not need to assign values to all the defined entities. Only those entities that you assign a value to change; those defined in the namelist but not assigned a value in the input data remain unchanged. Likewise, when you have defined character substrings and array elements in the namelist, only those you specify to receive input are changed. You can change part of a character substring. For example, to change the character variable *NAME* from 'TESTA' to 'TESTB', use the following namelist-directed input:

```
$SAM NAME(5:) = 'B' $END
```

The value for *NAME* is 'TESTB'; the first four positions of the value remain unchanged.

When you assign values to an array name, the first value is associated with the first element, the second value with the second element, and so on. The number of array elements you can assign must be less than or equal to the size of the array.

Assume a program with the following statements:

```
DIMENSION MYRAY(10)  
NAMELIST /SAM2/ MYRAY  
READ SAM2
```

Assume that the input is

```
$SAM2 MYRAY = 10, 8, , 70 $END
```

On execution, the READ statement assigns the following values to the array elements:

MYRAY (1)	10
MYRAY (2)	8
MYRAY (3)	Unchanged
MYRAY (4)	70
MYRAY (5-10)	Unchanged

Values MYRAY (3) and (5-10) remained unchanged because the two consecutive commas in the input indicate not to change the current value, and values for unspecified array elements remain unchanged.

Because values are assigned to the specified array elements and not assigned beginning with the first element, the READ statement can assign new values and not alter unspecified elements. For example, the following line assigns values to MYRAY elements 5-7; the unspecified elements remain unchanged:

```
SSAM2 MYRAY(5) = 9, 85, 60 $END
```

Namelist-directed formatting follows the following rules for list-directed input:

- Do not use spaces or tabs in groupname definitions. In value assignments, the entity name cannot contain spaces or tabs except within a subscript or substring specifier, where they are acceptable within the parentheses.
- The groupname and each entity must be contained within a single record.
- When assigning values, you can precede and follow the equal sign with any number of tabs and spaces.
- Character constants are enclosed in apostrophes. If you want an apostrophe to appear as part of the character string, use two consecutive apostrophes without an intervening blank or end record.
- You cannot use Hollerith, octal, or hexadecimal constants.
- Character constants can span record boundaries. Normally, the end of a record in namelist input is a space character. If the end of record occurs within a

character constant, however, the end of record is ignored; the last character of the previous record is followed by the first character of the next record.

- For fixed record length files, *NAMELIST* reads and writes records of a fixed length.

List-directed output

The format of list-directed output is defined by the data type of the I/O list items except that *r** is not used. Also, neither double nor single quotation marks are output for character constants. Table 21 shows the default output forms that the list-directed *WRITE* statement generates for each data type.

Table 21
List-directed output formats

Data type	Output format
<i>LOGICAL*1</i>	<i>L2</i>
<i>LOGICAL*2</i>	<i>L2</i>
<i>LOGICAL*4</i>	<i>L2</i>
<i>LOGICAL*8</i>	<i>L2</i>
<i>INTEGER*1</i>	<i>I5</i>
<i>INTEGER*2</i>	<i>I7</i>
<i>INTEGER*4</i>	<i>I12</i>
<i>INTEGER*8</i>	<i>I22</i>
REAL	1PG15.7E2
REAL*8	1PG24.15E3
REAL*16	1PG43.33E4
COMPLEX	1X, '(', 1PG14.7E2, ', ', 1PG14.7E2, ')'
COMPLEX*16	1X, '(', 1PG23.15E3, ', ', 1PG23.15E3, ')'
CHARACTER	<i>An</i> , where <i>n</i> represents the character expression length

Namelist-directed output formatting

The format of namelist-directed output is defined by the data type of the list entities in the corresponding *NAMELIST* statement. When you use a namelist-directed *WRITE* statement, the order of data output is specified by the sequence in which namelist entities are defined in the *NAMELIST* statement.

For example, assume a program with the following statements:

```
LOGICAL L4
INTEGER I4
REAL R4
COMPLEX C8
CHARACTER*20 CHAR20

NAMELIST /CONTROL/ L4, I4, R4, C8, CHAR20

READ (5, CONTROL)
WRITE (6, CONTROL)

END
```

with the following input:

```
$CONTROL
  L4 = F, I4 = -123213, C8 = (12,2),
  CHAR20='test case',
  R4=3.14159
$END
```

The `WRITE` statement outputs the following:

```
$CONTROL
L4           = F,
I4           = ^^^^-123213,
R4           = ^3.141590^^^^,
C8           = (^12.00000^^^^, ^2.000000^^^^),
CHAR20       = 'test^case^^^^^^^^^^^^^^'
$END
```

The output for this program segment consists of the current values of all list entities associated with the namelist specifier. You can output a value for an entity that is defined by the `NAMELIST` statement but is not assigned an input value. (The entity can also be undefined or defined elsewhere in the program.) For instance, if you had an entity, `count`, that was defined in the `NAMELIST` statement but received no input, the current value of `count` would be written as well as the values shown in the example.

As the example illustrates, each value begins on a new line for namelist-directed output. Character values are enclosed in apostrophes. As mentioned above, the data types used are the

same as those defined in the *NAMELIST* statement. The format output follows the same form as that of list-directed output. Although you can use the *\$* and *&* characters interchangeably on input, the *\$* character is always used for output.

Carriage-control characters

The first character of a formatted record transfers to the printer as a carriage-control character. Table 22 shows the characters that provide vertical format control.

Table 22
Vertical format control

Character	Interpretation
^	Advances one line; begins output at beginning of next line.
0	Advances two lines; skips one line and begins output.
1	Advances to new page; begins output at the top of a new page.
+	Overwrites; begins output at the beginning of the current line and returns to the left margin.
ASCII NUL	Overwrites with no advance; begins output at beginning of the current line and does not return to left margin.
\$	<i>Prompting; begins output at the beginning of the next line and suppresses carriage return at end of line.</i>

*Specifying FORM='PRINT' in the OPEN statement indicates formatted I/O and implies vertical format control for that unit. You can use the *fpr* utility to interpret the vertical format controls before printing the file.*

Subprograms



Subprograms are program units that can be invoked from other program units. Subprograms usually perform frequently used sequences of operations for the invoking program unit. Arguments (dummy and actual) of the subprogram are used to transfer information between the subprogram and another program unit. The dummy argument appears in the argument list of a subprogram, and the actual argument appears in the argument list of a subprogram reference.

There are two classes of subprograms: `BLOCK DATA` subprograms and procedures. `BLOCK DATA` subprograms provide initial data for common block variables and arrays. Procedure subprograms include functions and subroutines.

`BLOCK DATA` subprogram

A block data subprogram provides initial values for variables and array elements in named common blocks. A block data subprogram must not contain any executable statements.

A block data subprogram begins with a `BLOCK DATA` statement and ends with an `END` statement. You can use only one `BLOCK DATA` statement in a subprogram, but you can use more than one block data subprogram in the program units that constitute the executable program. The statement has the form:

```
BLOCK DATA [name]
```

where *name* is an optional symbolic name for the block data subprogram in which the `BLOCK DATA` statement appears. Do not assign the `BLOCK DATA` the same name as that of an external procedure, main program, common block, or other block data subprogram, or any local name in the block data subprogram.

You can use only these specification statements between the BLOCK DATA and END statements:

- COMMON
- DATA
- DIMENSION
- EQUIVALENCE
- IMPLICIT
- PARAMETER
- SAVE
- any type-declaration statements

Your block data subprogram must contain at least one COMMON statement and one DATA statement.

You must specify all entities having storage units in the common block, but you are not required to initialize all of the values. Be sure to provide specifications to establish the entire block.

Example:

```
BLOCK DATA MYBLOK
COMMON/EX/A, B, C
COMMON/CAT/LIST(100)
DATA A/3.5/, LIST/100*5/
END
```

Procedures

Procedure subprograms include both function subprograms and subroutine subprograms. These can include both user-defined subprograms and those supplied as part of the CONVEX FORTRAN system.

Procedure subprograms can contain executable statements. In general, after its initial defining statement, a procedure subprogram can contain any statement except a PROGRAM, BLOCK DATA, SUBROUTINE, or FUNCTION statement.

Dummy and actual arguments

Dummy arguments are classified as variables, arrays, or dummy procedures. They are used in statement functions, function subprograms, and subroutine subprograms to indicate the number and types of actual arguments to be transferred.

Dummy arguments indicate whether each actual argument is a

single value, array of values, procedure, or statement label. You cannot use a dummy argument name in a DATA, EQUIVALENCE, INTRINSIC, SAVE, or COMMON statement except as a common block name.

Actual arguments, which can be constants, symbolic names of constants, function references, expressions, arrays and array elements, character substrings, alternate return specifiers, or subprogram names, specify the entities that are to be associated with the dummy arguments. The type of each actual argument must agree with the type of its associated dummy argument, except when the actual argument is a subroutine name or an alternate return specifier. Actual arguments must also agree in order and number with the dummy arguments.

A function or subroutine reference establishes an association between the corresponding dummy and actual arguments. The dummy argument holds the value of the actual argument during execution. For example, using the following statement:

```
SUBROUTINE SAMPLE (R,L)
```

specifies R and L as dummy arguments. When the subroutine

```
CALL SAMPLE (B,80)
```

executes, the actual arguments (B,80) replace the dummy arguments (R,L). Thus B replaces R, and 80 replaces L. Any value assigned to R is also assigned to B.

The number of elements of a dummy argument used as an array cannot exceed the number of elements in the actual argument. Also, a type CHARACTER dummy argument length must not be larger than the length of the associated actual argument.

Variables as dummy arguments

To associate a dummy argument variable with an actual argument that is a variable, array element, substring, or expression (including a constant), use the variable, array element, substring, or expression as an actual argument and include a dummy argument of the same data type in the subprogram argument list.

You can define the associated dummy argument within the subprogram if the actual argument is a variable name, array element name, or substring name. If the associated actual argument is a constant or constant name, function reference, or an expression, however, it must not be defined within the subprogram. If you pass a constant to a subroutine as an actual

parameter and that subroutine attempts to modify the corresponding dummy argument, either by a `READ` or `ASSIGNMENT` statement, a segmentation violation occurs because the constants are stored in read-only storage.

Arrays as dummy arguments

If a dummy argument is declared as an array, it can only be associated with an actual argument that is an array or array element of the same type. To pass an array to a subprogram, use the array name as the actual argument. The subprogram must dimension the array to use it. That is, the dummy array must be specified in an array declarator in the subprogram. The declarator has the same format as that for an actual array but with the following differences:

- You cannot use the declarator in a `COMMON` statement. It is, however, permitted in a `DIMENSION` or `type` statement.
- Integer constant expressions and expressions containing integer constants and variables can be used as upper and lower bounds of array dimensions. These dimensions are considered adjustable, as one or both of the dimension-bound expressions is a variable. The array is called an adjustable array.
- You can use an asterisk (*) to specify the upper bound of the last dimension. In this case, the array is known as an assumed-size array.

Common block elements must not be passed as actual arguments if the called routine, or any routine that it calls, accesses that element from the common block.

Adjustable arrays

Adjustable arrays are used to process arrays of different sizes in a single subprogram. The adjustable array dimensions are determined in the reference to the subprogram.

Each dummy argument in the array declarator must be associated with an actual argument when the subprogram is entered. Any variable used in an adjustable dimension or each `COMMON` variable appearing in a dimension-bound expression must have a defined value when the subprogram is entered. The expressions specifying the adjustable dimensions are evaluated when the subprogram is entered. Argument association is not retained through different references to the subprogram. The bound values are determined each time a subprogram is entered.

In the statement

```
DIMENSION E(I, I), G(5, 2*I)
```

E and G are adjustable arrays.

The size of the adjustable array must be less than or equal to the size of the array of the corresponding actual argument.

Assumed-size arrays

An asterisk specifies the upper bound of the last array dimension in an assumed size array declarator. For example:

```
DIMENSION SAM (*)
```

sets the upper bound to assumed-size for a one-dimensional array. If the array has more than one dimension, only the last dimension can be assumed size. For example,

```
DIMENSION SAM (1:N, 1:*)
```

sets the upper bound for a two-dimensional array.

The assumed-size dummy array name cannot appear:

- In an I/O list of a data transfer statement or an implied DO loop without subscripts
- As an internal unit identifier in an I/O statement
- As a runtime format identifier in an I/O statement

The size of the dummy array is the size of the actual argument array when the actual argument is a noncharacter array name. When the actual argument is a noncharacter array element name, however, the size of the array is the array size plus one minus the subscript value.

The size of the dummy array is $\text{INT}(n + 1 - s) / l$ when the actual argument is one of the following:

- A character array name
- Character array element name
- Character array element substring

and:

- s is the character storage unit of an array where the actual argument begins.
- n is the number of character storage units in this array.
- l is the length of an element of the dummy array.

If an assumed size dummy array has n dimensions, the product of the sizes of the first $n-1$ dimensions must not exceed the size of the array.

Character arguments

You can use character values in one or more of the dummy arguments in a subprogram if the actual argument in the calling program unit is of type CHARACTER. Thus, a dummy argument that is a variable name of type CHARACTER can be associated only with an actual argument that is either a character variable, character array element, character substring, or character expression. A dummy argument that is an array name of type CHARACTER can be associated only with an actual argument that is a character array, character array element, or character array element substring.

If the actual argument is a Hollerith constant (for example, 3HSAM), the dummy argument must be of numeric data type. The corresponding dummy argument can have either a numeric or character data type when the actual argument is a character constant (for example, "SAM").

Character argument lengths

The length of the dummy argument must not exceed the length of the actual argument. The subprogram cannot access more characters than are declared for the argument in the calling unit. That is, when the dummy argument is of type CHARACTER, the associated actual argument must be less than or equal to the length of the actual argument. If the length of the dummy argument of type CHARACTER is less than the length of the associated actual argument, only the leftmost characters of the actual argument are associated with the dummy argument.

If you use an assumed-length character argument, it must be a dummy argument. When control transfers to a subprogram, the assumed-length character dummy argument must be associated with a character actual argument. It assumes the length of the corresponding actual argument. Thus, if you specify the dummy argument length as an assumed-length character argument (for example, `* (*)`), the length of the associated actual argument is used.

If the dummy argument is an array, you can specify a length that differs from that of the calling unit. In this case, however, the subprogram cannot access a character beyond the last character reserved by the calling unit for the array. When a dummy argument of type CHARACTER is an array name, the restriction on length is for the entire array and not for each array element.

You can also use a character array dummy argument with an assumed-length. In this case, the length of each element in the dummy argument equals the length of the elements in the actual argument. The assumed length and the array declarator determine the size of the assumed-length character array.

The following example illustrates character argument length specifications.

Example:

```
PROGRAM SAM
CHARACTER A1*2, A2*6, A3*8
.
.
.
END

SUBROUTINE MYEX (A)
CHARACTER A*6
.
.
.
END
```

Assume the following CALL statements occur in the main program SAM:

```
CALL MYEX (A1)
CALL MYEX (A2)
CALL MYEX (A3)
```

The first statement is invalid because the length of the dummy argument exceeds that of the associated argument; the remaining two statements are valid with the six leftmost characters of A3 being associated with A. If, however, the subprogram MYEX is defined as

```
SUBPROGRAM MYEXAM (A)
CHARACTER A* (*)
.
.
.
END
```

all three of the `CALL` statements are valid. The length of the dummy argument `A` is determined by the length of the corresponding actual argument.

Procedures as dummy arguments

A dummy argument is considered a dummy procedure if the dummy procedure name appears in the dummy argument list of a `FUNCTION`, `SUBROUTINE`, or `ENTRY` statement and if:

- It is referenced as a function.
- It appears in a type statement and `EXTERNAL` statement.
- It is referenced as a subroutine.

When you use a dummy argument that is a dummy procedure, associate it only with an actual argument that is an intrinsic function, external function, subroutine, or another dummy procedure.

When you use a dummy argument as an external function, or in a type statement and `EXTERNAL` statement, the associated actual argument must be an intrinsic function, external function, or dummy procedure.

If you use the dummy argument as a procedure name in a function reference and associate it with an intrinsic function, the arguments must agree in order, number, and type with those specified for the intrinsic function.

When you use the dummy argument as a subroutine, the actual argument must be the name of a subroutine or dummy procedure. If a procedure name appears only in a dummy argument list, an `EXTERNAL` statement, and an actual argument list, it is not possible to determine whether the symbolic name becomes associated with a function or subroutine by examination of the subprogram alone.

Alternate return arguments

You can use an asterisk as a dummy argument only in the dummy argument list of a `SUBROUTINE` statement or an `ENTRY` statement in a subroutine subprogram. When you use the asterisk as a dummy argument, the corresponding actual argument must be an alternate return specifier in the `CALL` statement.

Example:

```
SUBROUTINE EXAM(D3 , * , E2 , *)
```

The alternate return argument allows you to return control to any executable labeled statement in the calling program as long as you have included alternate return arguments in the corresponding positions. These actual arguments have the following form:

**label* or *&label*

Functions

A function can be an intrinsic function, a statement function, or an external function (function subprogram), all of which supply a value to the expression. The function is referenced from within another part of the program. When executed, a function has a value and a type. The general form of a function reference is

func ([*a* [*a*,] ...])

where *func* is the symbolic name of the function or dummy procedure being referenced, and *a* is an optional list of actual arguments separated by commas. If you do not include arguments, you must still include the enclosing parentheses.

Intrinsic functions

CONVEX FORTRAN supplies intrinsic functions as a built-in language feature. You can invoke these pre-existing functions by using the function name in any part of a user program; no definition is required.

There are two classes of intrinsic names: generic names and specific names. If you reference a generic intrinsic name, the compiler decides which special intrinsic to invoke based on the type of the actual arguments. When you reference specific names, the arguments to the intrinsic must be of a specific type. For example, the generic intrinsic function, LOG (natural logarithm), can accept arguments of type REAL, DOUBLE PRECISION, REAL*16, and COMPLEX, whereas the specific function, DLOG, can only accept a DOUBLE PRECISION argument.

Using a generic name generally simplifies function referencing because you can use the function name with more than one type of argument. You must use the appropriate specific name whenever the intrinsic function name is used as an actual argument in a subprogram. For either generic or specific functions that require multiple arguments, all arguments must be of the same data type. The compiler does not convert incorrectly typed arguments.

Use of the `IMPLICIT` statement does not alter the type of intrinsic functions.

An intrinsic reference has the following form:

$$\text{inf} (a [,a] [\dots])$$

where *inf* is an intrinsic name and *a* is the argument on which the function operates.

Built-in functions

Built-in functions allow a FORTRAN program to pass arguments to a program that is not written in FORTRAN.

%REF and %VAL functions

Two built-in functions—%REF and %VAL—can be used in the argument list of a CALL statement or function reference to change the form of the argument. Such a change can be necessary if you must call subprograms written in languages other than FORTRAN, because the actual argument may have to be passed in a form different from that used by FORTRAN. These two functions specify how to pass the argument to the subprogram.

Although you can use these functions in the actual argument list of a CALL statement or function reference, you cannot use them in any other context. You need not, however, use these built-in functions when invoking a FORTRAN library procedure or a user-supplied subprogram written in FORTRAN.

There are two built-in argument list functions:

- *%REF(a)*—This function passes the argument by reference.
- *%VAL(a)*—This function passes the argument as a 32-bit immediate value; an argument shorter than 32 bits is sign-extended to a 32-bit value.

a is an actual argument.

Table 23 shows the FORTRAN argument-passing defaults and allowable uses of %REF and %VAL.

Table 23
Built-in functions and
defaults for argument lists

Data type	Default	Functions allowed	
		%REF	%VAL
LOGICAL (*1,2,4)	REF	Yes	Yes [†]
LOGICAL*8	REF	Yes	No
INTEGER (*1,2,4)	REF	Yes	Yes [†]
INTEGER*8	REF	Yes	No
REAL*4	REF	Yes	Yes
REAL*8	REF	Yes	No
REAL*16	REF	Yes	No
COMPLEX	REF	Yes	No
CHARACTER	REF	Yes	No
Hollerith	REF	No	No
Array name			
Numeric	REF	Yes	No
Character	REF	Yes	No
Procedure name			
Numeric	REF	Yes	No
Character	REF	Yes	No

[†]If a logical or integer value occupies less than 32 bits of storage, it is converted to a 32-bit value by sign extension.

%LOC function

The %LOC built-in function computes the internal address of a storage element, as in the following example:

```
IADDR = %LOC(v)
```

where *v* is a variable name, array element name, array name, character substring name, or external procedure name, and IADDR is an INTEGER*4 variable.

The %LOC built-in function produces an INTEGER*4 value that represents the location of its argument. Addresses from user programs are BYTE addresses and are always negative integer values.

The use of %LOC is limited to arithmetic expressions.

Statement functions

The statement function is a nonexecutable, user-defined, single-statement procedure. You can reference a statement function only from the program unit in which it is defined. The form is similar to an arithmetic, logical, or character assignment

statement. *Statement function definitions must precede the use of the statement function.* The statement function returns a single value to the program. The form of a statement function is

$$\text{func ([d [,d] ...]) = exp}$$

where

func

is the name of the statement function. The type is defined by the implicit naming convention or by a prior type statement. Do not use the name to identify any other entity in the current program unit except a common block.

d

represents a variable name called a statement function dummy argument. The dummy argument holds the value of the actual argument during execution. The dummy argument list specifies the order, number, and type of actual argument whose values are used in the function reference. The actual arguments must agree in order, number, and type with the corresponding dummy arguments. (The compiler associates the actual arguments with the dummy arguments of the companion statement definition.)

Each name in the function definition must be unique and must be of the data type of the actual value that replaces it during the function reference. You can use the dummy argument name to identify a variable of the same type, as a dummy argument in a FUNCTION, SUBROUTINE, or ENTRY statement, or as a common block name.

exp

is an expression. Each primary of *exp* must be one of the following:

- A constant or symbolic name of a constant
- A variable reference
- An array element reference
- An intrinsic function reference
- A statement function that has been previously defined in the current program unit
- An external function reference
- A dummy procedure reference

A statement function is referenced using its function reference as a primary in an expression. The following example illustrates a statement function and the statement function reference.

Example:

```
IAVG (IGR1, IGR2, IGR3) = (IGR1 + IGR2 + IGR3) / 3
.
.
.
ISC = IAVG (IOR, IWR, IAP)
```

When the statement function reference executes, all actual arguments that are expressions are evaluated and actual arguments are associated with the corresponding dummy arguments. (The compiler substitutes the values in the actual arguments for the dummy arguments.) Then the processor evaluates the expression (right side of the statement function statement). Conversion occurs, if necessary, of the resulting value to the type of the statement function according to the usual arithmetic assignment rules; or a change occurs, if necessary, in the length of a character expression value according to the usual character assignment rules. The resulting value is available to the expression that contains the function reference.

In a statement function reference, any expressions can be used as actual arguments except array names and character expressions involving concatenation of an operand whose length specification is an asterisk (*) in parentheses, unless the operand is the symbolic name of a constant.

Function subprograms

A function subprogram is a separate program unit that includes a FUNCTION statement followed by a series of statements that define the computing procedure. The calling program unit references it; the statements execute, and through a RETURN or END statement, a single value returns to the function reference in the calling unit. This value is assigned to the function name.

The FUNCTION statement specifies the name of the function, the dummy arguments used by the function, and can indicate the type of the function value. In logical and numeric functions, the FUNCTION statement, including the CONVEX extension *m and optional parentheses, is represented by:

```
[typ] FUNCTION nam [*m] [ ( [d [,d]...] ) ]
```

where

typ

is one of the logical or numeric data-type specifiers.

nam

is the symbolic name of the function subprogram. If you neither specify *typ* nor declare the *nam* in a later type statement, the name implies the data type of the function.

m

is an unsigned, nonzero integer constant specifying the length of the data type. It must be one of the valid length specifiers for the data type given by *typ*.

d

is a dummy argument name that can include variable names, array names, or dummy procedure names. All dummy argument names must match the actual arguments in all references in number, order, and type. The dummy name is local to the program unit and must not appear in a DATA, EQUIVALENCE, INTRINSIC, SAVE, or COMMON statement, except as a common block name.

In character functions, the CHARACTER FUNCTION statement, with the CONVEX extension of **n* and optional parentheses, has the form:

```
CHARACTER[*n] FUNCTION nam [*n] [( [ d [,d] ... ) ]]
```

where *n* is either an unsigned, nonzero integer constant, or a parenthetical asterisk (*) indicating an assumed-length function name. If you specify CHARACTER* (*), the function always assumes the length declared for it in the program unit that invokes it. (An assumed-length character function can have different lengths when invoked by different program units.) If *n* is an integer constant, the value of *n* must agree with the length of the function specified in the program unit that invokes the function. If you do not specify *n*, a length of 1 is assumed. If the length has already been specified after the keyword CHARACTER, you cannot use the optional-length specification following *nam*. Both *nam* and *d* retain the same definition.

You must begin the function subprogram with the FUNCTION statement. You can include any statements except a BLOCK DATA, SUBROUTINE, PROGRAM, or another FUNCTION statement within the function subprogram. End it with an END statement. Between a FUNCTION and END statement, you can use the specified

function name as a variable in an executable statement or in a type statement if you omit *typ*. Include `ENTRY` statements to provide multiple entry points to the subprogram.

A function specified in a subprogram can reference other subprograms but cannot reference itself, directly or indirectly. It must assign a value to its symbolic name at least once.

Subroutine subprograms

A subroutine subprogram (subroutine) is a program unit that performs a specific, user-defined task or subtask for some other program unit of the program. Subroutines are similar to function subprograms; actual and dummy arguments are handled the same for both. The `RETURN` statement returns control to the calling program. They differ in that the subroutine names have no type, and no value is associated with the subroutine name. Also, you use a `CALL` statement, not a function reference, to invoke a subroutine. Within the subroutine, you can specify different points of return to the calling subprogram. The `CALL` statement has the form:

```
CALL sub [ ( [ a [, a] ... ] ) ]
```

where

sub

is the name of the subroutine.

a

is the actual argument which can be a constant, variable, expression, array, array element, character substring, alternate return specifier, intrinsic function name, external procedure name, or dummy procedure name. You can use an `*` or `&` followed by the label of an executable statement to indicate an alternate return. Do not use a character expression involving concatenation of an operand whose length specification is an asterisk in parentheses unless the operand is the symbolic name of a constant. If the actual argument is a Hollerith constant, the dummy argument must be of numeric data type.

Using the `CALL` statement invokes the subroutine. Control passes to the first executable statement using any *a* arguments for the subroutine dummy arguments. After the subroutine returns, control returns to the statement in the calling unit that follows the `CALL` statement unless you specify an alternate `RETURN` in the subroutine.

You must begin a subroutine with a `SUBROUTINE` statement and end it with an `END` statement. It specifies the name of the subroutine and the arguments used by the subroutine. The `SUBROUTINE` statement has the following form:

```
SUBROUTINE name [ ( [d [, d]... ] ) ]
```

where

name

is the symbolic name of the subroutine. Because the subroutine has no data type, you need not apply the naming rules. Because the name is global, do not use it for any other purpose in the program.

d

represents a dummy argument list consisting of a variable name, array name, dummy procedure name, or, if the subroutine uses alternate returns, an asterisk (*). Separate the argument list items with commas. The argument list can be empty. In this case, use of parentheses is optional; for example, either `SUBROUTINE EXAM` or `SUBROUTINE EXAM()` is acceptable.

If you indicate the dummy argument as *, be sure that the corresponding actual argument in the calling unit is also an * or & followed by the label of an executable statement within the calling unit. You can specify an alternate return in the `RETURN` statement by giving the position of this asterisk among other asterisks in the dummy argument.

You must specify an actual argument for each dummy argument in the `SUBROUTINE` statement of the called subroutine. If you use a variable, array element, or array as the actual argument, the data type must match that of the dummy argument. If the argument is the name of a subprogram, you must declare this name in an `EXTERNAL` statement in this program unit.

The `ENTRY` statement can be used to specify multiple entry points for subroutines.

ENTRY statement

You can use the nonexecutable `ENTRY` statement to specify alternative entry points into a function or subroutine subprogram. You can reference an `ENTRY` from any program unit except the subprogram that contains it. Use a function reference for `ENTRY` in a function; use `CALL` for `ENTRY` in a subroutine. You

can place an `ENTRY` anywhere between the initial `FUNCTION` or `SUBROUTINE` statement and the `END` statement, but you must not place it within a block `IF` or the range of a `DO` loop.

The form of the `ENTRY` statement is

```
ENTRY nam [ ( [d [, d] ... ] ) ]
```

where

nam

is the symbolic name of the entry point representing either a subroutine name in a subroutine or an external function name in a function subprogram. When entry is in a function, *nam* has a data type that you can imply or specify. If you use a type statement, it can appear before or after the `ENTRY` statement. For entry in a subroutine, however, there is no data type restriction.

d

is a variable name, array name, or dummy procedure. You can use an asterisk (*) as an alternate return only if the entry is in a subroutine. You can omit the parentheses for an empty argument list in a subroutine entry, but the parentheses must always be included in a function entry and in the entry reference.

You can use dummy arguments in `ENTRY` statements that differ in order, number, type, and name from the dummy arguments you use in the `FUNCTION`, `SUBROUTINE`, or other `ENTRY` statements in the same subprogram. Each reference to a function or subroutine, however, must use an actual argument list that agrees in order, number, and type with the dummy argument list in the corresponding `FUNCTION`, `SUBROUTINE`, or `ENTRY` statement.

If a dummy argument is not currently associated with an actual argument, it is undefined. The association between actual dummy arguments is not retained between references.

An entry name or the name of the function subprogram defines all associated names of the same data type. All entry names within a function subprogram are associated with the name of the function subprogram. You can use function and entry names of different data types.

You must define the variable whose name is used to reference the function before a `RETURN` or `END` statement is executed in the subprogram. You don't need to use associated variables of the

same type unless the function is type CHARACTER; but an associated variable of different type must not become defined within the subprogram.

RETURN statement

The RETURN statement is used to return from a subroutine subprogram to one of several alternate points in the calling program unit. You can use none, one, or more than one RETURN statement in a subroutine. Use the alternate return only with subroutine subprograms, not function subprograms. You can, however, use the RETURN statement without an alternate specifier, in either a function or subroutine subprogram.

The statement has the following form:

```
RETURN [e]
```

where *e* is an optional integer expression that specifies an alternate statement in the calling program is to receive control. *The system converts the value type to integer if necessary.* The *e* represents the number, such as RETURN 2, of the corresponding asterisk, among other asterisks, in the dummy argument list of the subroutine. The alternate return specifier has the form of an asterisk or ampersand followed by the label of an executable statement (for example, *30 or &30).

Example:

```
      .  
      .  
      CALL EXAM(D, *30, E, *40)  
      .  
      .  
30      !RETURN 1 goes here.  
      .  
      .  
40      !RETURN 2 goes here.  
      .  
      .  
      .  
      END
```

```
SUBROUTINE EXAM(D3,*,E2,*)  
.  
.  
.  
RETURN      !Returns after the CALL statement.  
.  
.  
RETURN 1    !Returns to 30.  
.  
.  
RETURN 2    !Returns to 40.  
.  
.  
.  
END
```

In the preceding example, RETURN 1 indicates control transfers to the statement at line 30; RETURN 2 indicates the alternate return transfers control to the statement at line 40. RETURN indicates control transfers to the statement immediately after the CALL statement.

If you do not specify a RETURN, the END statement has the same effect as the RETURN.

Use of the alternate RETURN statement allows you to return control to any labeled statement in the calling program whose label you specify as an alternate return specifier to the subprogram. If e is less than 1 or greater than the total number of asterisks appearing in the dummy argument list, control returns as for a normal RETURN (without specifier).

When a RETURN or END statement executes, the subprogram ends the association between the dummy arguments and the current actual arguments. If the `-re` compiler option is specified, all local data that did not appear in a SAVE statement becomes undefined.

When used within a function, RETURN transfers control to the function reference in the calling unit and returns the function value. In a subroutine, RETURN transfers control to the statement following the CALL statement in the calling program unit.

Intrinsics and commonly used library routines

A

This appendix lists CONVEX FORTRAN's generic and specific intrinsics, as well as its commonly used library routines.

Generic and specific intrinsics

Table 24 lists generic and specific intrinsics. Numbers in the first and second column of Table 24 refer to notes following the table.

Table 24
Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
SQRT 1	SQRT DSQRT QSQRT CSQRT CDSQRT	Square root $a^{1/2}$	1	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16
LOG 2	ALOG DLOG QLOG CLOG CDLOG	Natural log $\log_e a$	1	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16
LOG10 2	ALOG10 DLOG10 QLOG10	Common log $\log_{10} a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16

Table 24 (continued)
 Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
EXP	EXP DEXP QEXP CEXP CDEXP	Exponential e^a	1	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16
SIN 3	SIN DSIN QSIN CSIN CDSIN	Sine $\sin a$	1	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16
SIND 3	SIND DSIND QSIND	Sine (degree) $\sin a^\circ$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
COS 3	COS DCOS QCOS CCOS CDCOS	Cosine $\cos a$	1	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16	REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16
COSD 3	COSD DCOSD QCOSD	Cosine (degree) $\cos a^\circ$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
TAN 3	TAN DTAN QTAN	Tangent $\tan a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
TAND 3	TAND DTAND QTAND	Tangent (degree) $\tan a^\circ$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
ASIN 4,5	ASIN DASIN QASIN	Arcsine $\arcsin a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
ASIND 2,4,5	ASIND DASIND QASIND	Arcsine (degree) $\arcsin a^\circ$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16

Table 24 (continued)
Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
ACOS 4,5	ACOS DACOS QACOS	Arccosine $\arccos a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
ACOSD 2,4,5	ACOSD DACOSD QACOSD	Arccosine (degree) $\arccos a^\circ$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
ATAN 5	ATAN DATAN QATAN	Arctangent $\arctan a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
ATAND 2,5	ATAND DATAND QATAND	Arctangent (degree) $\arctan a^\circ$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
ATAN2 5,6	ATAN2 DATAN2 QATAN2	Arctangent (two arguments) $\arctan a_1/a_2$	2	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
ATAN2D 2,5,7	ATAN2D DATAN2D QATAN2D	Arctangent (two degree arguments) $\arctan a_1^\circ/a_2^\circ$	2	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
SINH	SINH DSINH QSINH	Hyperbolic Sine $\sinh a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
COSH	COSH DCOSH QCOSH	Hyperbolic Cosine $\cosh a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16
TANH	TANH DTANH QTANH	Hyperbolic Tangent $\tanh a$	1	REAL*4 REAL*8 REAL*16	REAL*4 REAL*8 REAL*16

Intrinsics

Table 24 (continued)
 Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
ABS 8	IIABS JIABS KIABS ABS DABS QABS CABS CDABS	Absolute value <i> a </i>	1	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*16 COMPLEX*8 COMPLEX*16	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*16 REAL*4 REAL*8
IABS 8	IIABS JIABS KIABS	Absolute value <i> a </i>	1	INTEGER*2 INTEGER*4 INTEGER*8	INTEGER*2 INTEGER*4 INTEGER*8
INT 9,14 15	IINT JINT KINT IIDINT JIDINT KIDINT IIQINT JIQINT KIQINT	Truncation <i>[a]</i>	1	REAL*4 REAL*4 REAL*4 REAL*8 REAL*8 REAL*8 REAL*16 REAL*16 REAL*16 COMPLEX*8 COMPLEX*8 COMPLEX*8 COMPLEX*16 COMPLEX*16 COMPLEX*16	INTEGER*2 INTEGER*4 INTEGER*8 INTEGER*2 INTEGER*4 INTEGER*8 INTEGER*2 INTEGER*4 INTEGER*8 INTEGER*2 INTEGER*4 INTEGER*8 INTEGER*2 INTEGER*4 INTEGER*8
INT 19				COMPLEX*8 COMPLEX*8 COMPLEX*8 COMPLEX*16 COMPLEX*16 COMPLEX*16	INTEGER*2 INTEGER*4 INTEGER*8 INTEGER*2 INTEGER*4 INTEGER*8
10,14	INT1 INT2 INT4 INT8	Conversion to INTEGER	1	Any numeric Any numeric Any numeric Any numeric	INTEGER*1 INTEGER*2 INTEGER*4 INTEGER*8
IDINT 9,14 15	IIDINT JIDINT KIDINT	Truncation <i>[a]</i>	1	REAL*8 REAL*8 REAL*8	INTEGER*2 INTEGER*4 INTEGER*8

Table 24 (continued)
Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
<i>IQINT</i> 9,14 15	<i>IIQINT</i> <i>JIQINT</i> <i>KIQINT</i>	<i>Truncation</i> <i>[a]</i>	1	<i>REAL*16</i> <i>REAL*16</i> <i>REAL*16</i>	<i>INTEGER*2</i> <i>INTEGER*4</i> <i>INTEGER*8</i>
<i>AINT</i>	<i>AINT</i> <i>DINT</i> <i>QINT</i>	<i>Truncation</i> <i>[a]</i>	1	<i>REAL*4</i> <i>REAL*8</i> <i>REAL*16</i>	<i>REAL*4</i> <i>REAL*8</i> <i>REAL*16</i>
<i>NINT</i> 9,14 15	<i>ININT</i> <i>JNINT</i> <i>KNINT</i> <i>IIDNNT</i> <i>JIDNNT</i> <i>KIDNNT</i> <i>IIQNNT</i> <i>JIQNNT</i> <i>KIQNNT</i>	<i>Nearest integer</i> <i>[a+.5*sign(a)]</i>	1	<i>REAL*4</i> <i>REAL*4</i> <i>REAL*4</i> <i>REAL*8</i> <i>REAL*8</i> <i>REAL*8</i> <i>REAL*16</i> <i>REAL*16</i> <i>REAL*16</i>	<i>INTEGER*2</i> <i>INTEGER*4</i> <i>INTEGER*8</i> <i>INTEGER*2</i> <i>INTEGER*4</i> <i>INTEGER*8</i> <i>INTEGER*2</i> <i>INTEGER*4</i> <i>INTEGER*8</i>
<i>IDNINT</i> 9,14 15	<i>IIDNNT</i> <i>JIDNNT</i> <i>KIDNNT</i>	<i>Nearest integer</i> <i>[a+.5*sign(a)]</i>	1	<i>REAL*8</i> <i>REAL*8</i> <i>REAL*8</i>	<i>INTEGER*2</i> <i>INTEGER*4</i> <i>INTEGER*8</i>
<i>IQNINT</i> 9,14 15	<i>IIQNNT</i> <i>JIQNNT</i> <i>KIQNNT</i>	<i>Nearest integer</i> <i>[a+.5*sign(a)]</i>	1	<i>REAL*16</i> <i>REAL*16</i> <i>REAL*16</i>	<i>INTEGER*2</i> <i>INTEGER*4</i> <i>INTEGER*8</i>
<i>ANINT</i> 9,15	<i>ANINT</i> <i>DNINT</i> <i>QNINT</i>	<i>Nearest integer</i> <i>[a+.5*sign(a)]</i>	1	<i>REAL*4</i> <i>REAL*8</i> <i>REAL*16</i>	<i>REAL*4</i> <i>REAL*8</i> <i>REAL*16</i>

Table 24 (continued)
Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
ZEXT 14,15	<i>IZEXT</i>	Zero-extend <i>a</i>	1	<i>LOGICAL*1</i>	<i>INTEGER*2</i>
	19			<i>LOGICAL*2</i>	<i>INTEGER*2</i>
	19			<i>INTEGER*2</i>	<i>INTEGER*2</i>
	<i>JZEXT</i>			<i>LOGICAL*1</i>	<i>INTEGER*4</i>
	19			<i>LOGICAL*2</i>	<i>INTEGER*4</i>
	19			<i>LOGICAL*4</i>	<i>INTEGER*4</i>
	19			<i>INTEGER*2</i>	<i>INTEGER*4</i>
	19			<i>INTEGER*4</i>	<i>INTEGER*4</i>
	<i>KZEXT</i>			<i>LOGICAL*1</i>	<i>INTEGER*8</i>
	19			<i>LOGICAL*2</i>	<i>INTEGER*8</i>
	19			<i>LOGICAL*4</i>	<i>INTEGER*8</i>
	19			<i>LOGICAL*8</i>	<i>INTEGER*8</i>
	19			<i>INTEGER*2</i>	<i>INTEGER*8</i>
	19			<i>INTEGER*4</i>	<i>INTEGER*8</i>
19	<i>INTEGER*8</i>	<i>INTEGER*8</i>			
REAL 10	<i>FLOATI</i>	Convert <i>a</i> to <i>REAL*4</i>	1	<i>INTEGER*2</i>	<i>REAL*4</i>
	<i>FLOATU</i>			<i>INTEGER*4</i>	<i>REAL*4</i>
	<i>FLOATK</i>			<i>INTEGER*8</i>	<i>REAL*4</i>
	19			<i>REAL*4</i>	<i>REAL*4</i>
	<i>SNGL</i>			<i>REAL*8</i>	<i>REAL*4</i>
	19			<i>COMPLEX*8</i>	<i>REAL*4</i>
	19			<i>COMPLEX*16</i>	<i>REAL*4</i>
	<i>SNGLQ</i>			<i>REAL*16</i>	<i>REAL*4</i>
DBLE 10	19	Convert <i>a</i> to <i>REAL*8</i>	1	<i>INTEGER*2</i>	<i>REAL*8</i>
	19			<i>INTEGER*4</i>	<i>REAL*8</i>
	19			<i>INTEGER*8</i>	<i>REAL*8</i>
	19			<i>REAL*4</i>	<i>REAL*8</i>
	19			<i>REAL*8</i>	<i>REAL*8</i>
	19			<i>COMPLEX*8</i>	<i>REAL*8</i>
	19			<i>COMPLEX*16</i>	<i>REAL*8</i>
	<i>DBLEQ</i>			<i>REAL*16</i>	<i>REAL*8</i>

Table 24 (continued)
Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
QEXT 19		<i>Convert a to REAL*16</i>	1	INTEGER*2	REAL*16
				INTEGER*4	REAL*16
				INTEGER*8	REAL*16
				REAL*4	REAL*16
				REAL*8	REAL*16
	QEXTD 19			REAL*16	REAL*16
	19			COMPLEX*8	REAL*16
	19			COMPLEX*16	REAL*16
IFIX 10,14 15	IIFIX JIFIX KIFIX	<i>Fix a (REAL*4-to-INTEGER conversion)</i>	1	REAL*4	INTEGER*2
				REAL*4	INTEGER*4
				REAL*4	INTEGER*8
FLOAT 10	FLOATI FLOATJ FLOATK	<i>Float a (INTEGER-to-REAL*4 conversion)</i>	1	INTEGER*2	REAL*4
				INTEGER*4	REAL*4
				INTEGER*8	REAL*4
DFLOAT 10	DFLOTI DFLOTJ DFLOTK	<i>Float a (INTEGER-to-REAL*8 conversion)</i>	1	INTEGER*2	REAL*8
				INTEGER*4	REAL*8
				INTEGER*8	REAL*8
QFLOAT 10	QFLOTI QFLOTJ QFLOTK	<i>Float a (INTEGER-to-REAL*16 conversion)</i>	1	INTEGER*2	REAL*16
				INTEGER*4	REAL*16
				INTEGER*8	REAL*16
CMPLX 19		<i>Convert a₁ to COMPLEX*8 or convert a₁ and a₂ to COMPLEX*8</i>	1,2 1,2 1,2 1,2 1,2 1 1	INTEGER*2 INTEGER*2 INTEGER*4 REAL*4 REAL*8 COMPLEX*8 COMPLEX*16	COMPLEX*8 COMPLEX*8 COMPLEX*8 COMPLEX*8 COMPLEX*8 COMPLEX*8 COMPLEX*8
DCMPLX 11,15 19		<i>Convert a₁ to COMPLEX*8 or convert a₁ and a₂ to COMPLEX*8</i>	1,2 1,2 1,2 1,2 1,2 1 1	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 COMPLEX*8 COMPLEX*16	COMPLEX*16 COMPLEX*16 COMPLEX*16 COMPLEX*16 COMPLEX*16 COMPLEX*16 COMPLEX*16

Intrinsics

Table 24 (continued)
 Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
	REAL DREAL	Real part of complex	1	COMPLEX*8 COMPLEX*16	REAL*4 REAL*8
	AIMAG DIMAG	Imaginary part of complex	1	COMPLEX*8 COMPLEX*16	REAL*4 REAL*8
CONJG	CONJG DCONJG	Complex conjugate (if $a = (x, y)$ then $CONJG(a) = (x, -y)$)	1	COMPLEX*8 COMPLEX*16	COMPLEX*8 COMPLEX*16
	DPROD	REAL*8 product of REAL*4 $a_1 a_2$	2	REAL*4	REAL*8
MAX 12,15	IMAX0 JMAX0 KMAX0 AMAX1 DMAX1 IIDMAX1 JIDMAX1 KIDMAX1 QMAX1	Maximum $\max(a_1, a_2, \dots, a_n)$	n	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*8 REAL*8 REAL*8 REAL*16	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 INTEGER*2 INTEGER*4 INTEGER*8 REAL*16
MAX0 12,15	IMAX0 JMAX0 KMAX0	Maximum $\max(a_1, a_2, \dots, a_n)$	n	INTEGER*2 INTEGER*4 INTEGER*8	INTEGER*2 INTEGER*4 INTEGER*8
MAX1 12,14 15	IMAX1 JMAX1 KMAX1	Maximum $\max(a_1, a_2, \dots, a_n)$	n	REAL*4 REAL*4 REAL*4	INTEGER*2 INTEGER*4 INTEGER*8
AMAX0 12,15	AIMAX0 AJMAX0 AKMAX0	Maximum $\max(a_1, a_2, \dots, a_n)$	n	INTEGER*2 INTEGER*4 INTEGER*8	REAL*4 REAL*4 REAL*4

Table 24 (continued)
Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
MIN 13, 15	IMINO JMINO KMINO AMIN1 DMIN1 IIDMIN1 JIDMIN1 KIDMIN1 QMIN1	Minimum $\min(a_1, a_2, \dots, a_n)$	<i>n</i>	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*8 REAL*8 REAL*8 REAL*16	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 INTEGER*2 INTEGER*4 INTEGER*8 REAL*16
MINO 13,15	IMINO JMINO KMINO	Minimum $\min(a_1, a_2, \dots, a_n)$	<i>n</i>	INTEGER*2 INTEGER*4 INTEGER*8	INTEGER*2 INTEGER*4 INTEGER*8
MIN1 13,14 15	IMIN1 JMIN1 KMIN1	Minimum $\min(a_1, a_2, \dots, a_n)$	<i>n</i>	REAL*4 REAL*4 REAL*4	INTEGER*2 INTEGER*4 INTEGER*8
AMINO	AIMINO AJMINO	Minimum $\min(a_1, a_2, \dots, a_n)$	<i>n</i>	INTEGER*2 INTEGER*4 5	REAL*4 REAL*4
DIM	IIDIM JIDIM KIDIM DIM DDIM QDIM	Positive difference $a_1 - (\min(a_1, a_2))$	2	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*16	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*16
IDIM	IIDIM JIDIM KIDIM	Positive difference $a_1 - (\min(a_1, a_2))$	2	INTEGER*2 INTEGER*4 INTEGER*8	INTEGER*2 INTEGER*4 INTEGER*8
MOD	IMOD JMOD KMOD AMOD DMOD QMOD	Remainder $a_1 - a_2 * [a_1 / a_2]$ (Returns the remainder when the first argument is divided by the second)	2	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*16	INTEGER*2 INTEGER*4 INTEGER*8 REAL*4 REAL*8 REAL*16

Intrinsics

Table 24 (continued)
 Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
SIGN	IISIGN	Transfer of sign $ a_1 $ if $a_2 \geq 0$ $- a_1 $ if $a_2 < 0$	2	INTEGER*2	INTEGER*2
	JISIGN			INTEGER*4	INTEGER*4
	KISIGN			INTEGER*8	INTEGER*8
	SIGN			REAL*4	REAL*4
	DSIGN			REAL*8	REAL*8
	QSIGN			REAL*16	REAL*16
ISIGN	IISIGN	Transfer of sign $ a_1 \text{ sign } a_2$	2	INTEGER*2	INTEGER*2
	JISIGN			INTEGER*4	INTEGER*4
	KISIGN			INTEGER*8	INTEGER*8
IAND	IIAND	Bitwise AND (performs a logical AND on corresponding bits)	2	INTEGER*2	INTEGER*2
	JIAND			INTEGER*4	INTEGER*4
	KIAND			INTEGER*8	INTEGER*8
IOR	IIOR	Bitwise OR (performs an inclusive OR on corresponding bits)	2	INTEGER*2	INTEGER*2
	JIOR			INTEGER*4	INTEGER*4
	KIOR			INTEGER*8	INTEGER*8
IEOR	IIEOR	Bitwise XOR (exclusively ORs corresponding bits)	2	INTEGER*2	INTEGER*2
	JIEOR			INTEGER*4	INTEGER*4
	KIEOR			INTEGER*8	INTEGER*8
NOT	INOT	Bitwise complement (complements each bit)	1	INTEGER*2	INTEGER*2
	JNOT			INTEGER*4	INTEGER*4
	KNOT			INTEGER*8	INTEGER*8
ISHFT [†] 16	IISHFT	Bitwise shift (a_1 logically shifted a_2 bits—positive a_2 argument shifts left; negative, right)	2	INTEGER*2	INTEGER*2
	JISHFT			INTEGER*4	INTEGER*4
	KISHFT			INTEGER*8	INTEGER*8
IBITS [†] 17	IIBITS	Bit extraction (extracts bits a_2 through a_2+a_3-1 from a_1)	3	INTEGER*2	INTEGER*2
	JIBITS			INTEGER*4	INTEGER*4
	KIBITS			INTEGER*8	INTEGER*8
IBSET [†]	IIBSET	Bit set (returns the value of a_1 with bit a_2 of a_1 set to 1)	2	INTEGER*2	INTEGER*2
	JIBSET			INTEGER*4	INTEGER*4
	KIBSET			INTEGER*8	INTEGER*8
BTEST [†]	BITEST	Bit test (returns TRUE if bit a_2 of argument a_1 equals 1)	2	INTEGER*2	LOGICAL*2
	BJTEST			INTEGER*4	LOGICAL*4
	BKTEST			INTEGER*8	LOGICAL*8

Table 24 (continued)
Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
<i>IBCLR</i> [†]	<i>IIBCLR</i> <i>JIBCLR</i> <i>KIBCLR</i>	Bit clear (returns the value of a_1 with bit a_2 of a_1 set to 0)	2	<i>INTEGER</i> *2 <i>INTEGER</i> *4 <i>INTEGER</i> *8	<i>INTEGER</i> *2 <i>INTEGER</i> *4 <i>INTEGER</i> *8
<i>ISHFTC</i> [†]	<i>IISHFTC</i> <i>JISHFTC</i> <i>KISHFTC</i>	Bitwise circular shift (circularly shifts rightmost a_3 bits of argument a_1 by a_2)	3	<i>INTEGER</i> *2 <i>INTEGER</i> *4 <i>INTEGER</i> *8	<i>INTEGER</i> *2 <i>INTEGER</i> *4 <i>INTEGER</i> *8
15	LEN	Length (returns length of the character expression)	1	CHARACTER	<i>INTEGER</i> *4
15	INDEX	Index(c_1, c_2) (returns the position of the substring c_2 in the character expression c_1)	2	CHARACTER	<i>INTEGER</i> *4
15 19	CHAR	Character (returns a character that has the ASCII value specified by the argument)	1	<i>LOGICAL</i> *1 <i>INTEGER</i> *2 <i>INTEGER</i> *4 <i>INTEGER</i> *8	CHARACTER CHARACTER CHARACTER CHARACTER
15	ICHAR	ASCII value (returns the ASCII value of the argument, which must be a character expression of length 1)	1	CHARACTER	<i>INTEGER</i> *4
	LLT LLE LGT LGE	Character relationals (ASCII collating sequence)	2	CHARACTER CHARACTER CHARACTER CHARACTER	<i>LOGICAL</i> *4 <i>LOGICAL</i> *4 <i>LOGICAL</i> *4 <i>LOGICAL</i> *4
<i>DOT_PRODUCT</i> 18		<i>Dot product of two rank-one arrays</i>	2	<i>Varies</i> ^{††}	<i>Varies</i> ^{††}
<i>MATMUL</i> 18		<i>Matrix multiply</i>	2	<i>Varies</i> ^{††}	<i>Varies</i> ^{††}
<i>ALL</i> 18		<i>True if all array elements along optional dimension are true</i>	1,2	<i>Varies</i> ^{††}	<i>Varies</i> ^{††}
<i>ANY</i> 18		<i>True if any array element along optional dimension is true</i>	1,2	<i>Varies</i> ^{††}	<i>Varies</i> ^{††}
<i>COUNT</i> 18		<i>Counts number of true elements along optional dimension in array</i>	1,2	<i>Varies</i> ^{††}	<i>Varies</i> ^{††}

Table 24 (continued)
 Generic and specific intrinsics

Intrinsics		Function	No. args	Argument type	Result type
Generic	Specific				
MAXVAL 18		Maximum value along optional dimension using optional mask	1-3	Varies ^{††}	Varies ^{††}
MINVAL 18		Minimum value along optional dimension using optional mask	1-3	Varies ^{††}	Varies ^{††}
PRODUCT 18		Product of array elements along optional dimension using optional mask	1-3	Varies ^{††}	Varies ^{††}
SUM 18		Sum all the elements of an array along optional dimension, with an optional mask	1-3	Varies ^{††}	Varies ^{††}
MERGE 18		Merge under mask	3	Varies ^{††}	Varies ^{††}
PACK 18		Pack an array into an array of rank one under mask	3	Varies ^{††}	Varies ^{††}
SPREAD 18		Replicates array by adding a dimension	3	Varies ^{††}	Varies ^{††}
UNPACK 18		Unpack an array of rank one into an array under a mask	3	Varies ^{††}	Varies ^{††}
CSHIFT 18		Circular shift array elements	2-3	Varies ^{††}	Varies ^{††}
EOSHIFT 18		End-off shift array elements	2-3	Varies ^{††}	Varies ^{††}
TRANSPOSE 18		Transpose of an array of rank two	1	Varies ^{††}	Varies ^{††}
MAXLOC 18		Location of a maximum value in an array under an optional mask	2,3	Varies ^{††}	Varies ^{††}
MINLOC 18		Location of a minimum value in an array under an optional mask	2,3	Varies ^{††}	Varies ^{††}

[†] Arguments after the first are converted to the type of the first argument.

^{††} The argument list to the corresponding library routines is variable for these Fortran 90 intrinsics. The types of the arguments to these functions are not predetermined.

Notes

1. The SQRT, DSQRT, or QSQRT argument must be greater than or equal to zero. The result is the principal value where the real part is greater than or equal to zero. If the real part is zero, the result is the principal value and the imaginary part greater than or equal to zero.
2. The ALOG, DLOG, QLOG, ALOG10, DLOG10, QLOG10, ATAND, ATAN2D, ASIND, DASIND, QASIND, ACOSD, DACOSD, or QACOSD argument, must be greater than zero. The CLOG or CDLOG argument cannot be (0., 0.).
3. The SIN, DSIN, QSIN, COS, DCOS, QCOS, TAN, DTAN, or QTAN argument must be in radians, and is treated as modulo 2π . The SIND, COSD, or TAND argument must be in degrees; the argument is treated as modulo 360. The value of the sine and cosine functions for very large arguments is not meaningful because of the accuracy of the argument reduction. For single precision, the maximum value is $\pi * 2^{23}$; for double precision, the maximum value is $\pi * 2^{52}$; for quadruple precision, the maximum value is $\pi * 2^{112}$. If the argument exceeds the maximum, it is replaced with 0 and evaluation continues.
4. The absolute value of the ASIN, DASIN, QASIN, ACOS, DACOS, QACOS, ASIND, DASIND, QASIND, ACOSD, DACOSD, or QACOSD argument must be less than or equal to 1.
5. The result of ASIN, DASIN, QASIN, ACOS, DACOS, QACOS, ATAN, DATAN, QATAN, ATAN2, DATAN2, or QATAN2 is in radians, and that of ASIND, DASIND, QASIND, ACOSD, DACOSD, QACOSD, ATAND, DATAND, QATAND, ATAN2D, DATAN2D, or QATAN2D is in degrees.
6. If the value of the first ATAN2, DATAN2, or QATAN2 argument is positive, the result is positive; if it is zero, the result is zero if the second argument is positive, and π if the second argument is negative. A negative value for the first argument determines a negative result. A zero value for the second argument results in the absolute value of $\pi/2$. Both arguments cannot have the value zero. The range of the result for ATAN2, DATAN2, or QATAN2 is $-\pi < \text{result} \leq \pi$.
7. If the value of the first ATAN2D, DATAN2D, or QATAN2D argument is positive, the result is positive. If it is zero, the result is zero if the second argument is positive, and

- 180 degrees if the second argument is negative. A negative value for the first argument means the result is negative. A zero value for the second argument results in the absolute value of 90 degrees. Both arguments cannot be zero. The range of the result for `ATAN2`, `DTAN2D`, or `QATAN2D` is: $-180 \text{ degrees} < \text{result} \leq 180 \text{ degrees}$.
8. The absolute value of a complex number, (X, Y) , is the real value: $(X**2 + Y**2)**1/2$.
 9. Define $[x]$ as the largest integer whose magnitude is not greater than the magnitude of x , and whose sign matches that of x . For example $[5.7]$ equals 5 and $[-5.7]$ equals -5.
 10. Functions used to convert one data type to another have the same effect as the implied conversion in assignment statements. The functions `REAL` with a real argument, `DBLE` with a double-precision argument, and `INT` with an integer argument return the value of the argument without conversion.
 11. If `CMPLX` or `DCMPLX` has only one argument, the argument converts into the real part of a complex value, and zero is assigned to the imaginary part. If there are two arguments (not complex), conversion of the first argument into the real part of the value, and the second argument into the imaginary part, produces a complex value.
 12. This function causes the return of the maximum value from the argument list; there must be at least two arguments.
 13. This function causes the return of the minimum value from the argument list; there must be at least two arguments.
 14. This function converts to the default `INTEGER` type.
 15. The `INT`, `IDINT`, `IQINT`, `NINT`, `IDNINT`, `IQNINT`, `IFIX`, `MAX1`, `MINI`, and `ZEXT` functions return `INTEGER*4` values if the `/I4` (`COVUE` only) or `-i4` flag is in effect, `INTEGER*2` values if the `/NOI4` (`COVUE` only) or `-i2` flag is in effect, or `INTEGER*8` values if the `-i8`, `-p8`, or `-pd8` flag is in effect.

16. These functions shift binary patterns—positive, left (a_1 is >0) and negative, right ($a_2 < 0$). Since ISHFT indicates a logical shift, the bits shifted out of one end are lost and zeros are shifted in at the other end. As ISHFTC specifies a circular shift, the bits shifted out at one end are shifted back in at the other end.
17. In CONVEX FORTRAN, bits within a word are numbered from right to left. For example, in REAL*4 data, the rightmost (least significant) bit is 0; the leftmost (most significant) bit is 31.

Bit extraction proceeds from right to left; thus, a_2 is the rightmost bit extracted and a_2+a_3-1 is the leftmost bit. The extracted field is shifted right before being placed into the receiving field. If the extracted field is shorter than the receiving field, zeros are filled in from the left.

The following example illustrates bit extraction:

```

PROGRAM TEST
  INTEGER*4 I1, J1, K1, L1
  I1 = 'FFFFFFFF'X

  J1 = IBITS (I1,1,4)
  K1 = IBITS (I1,29,4)
  L1 = IBITS (I1,16,6)

  WRITE (6,100) 'I1 = ', I1
  WRITE (6,100) 'J1 = ', J1
  WRITE (6,100) 'K1 = ', K1
  WRITE (6,100) 'L1 = ', L1

100 FORMAT (A4,Z8)
  END

```

When the program is executed, the output displays the original bit pattern and the extracted bit fields, as follows:

```

I1 =FFFFFFFF
J1 =      F
K1 =      7
L1 =     3F

```

18. These are Fortran 90 intrinsics. They have no user-callable specific functions. For a more detailed explanation of their use, refer to Appendix C, "Fortran 90 compatibility."
19. No user-callable specific functions are available for these argument/result pairs.

Commonly used library routines

Commonly used library subroutines and functions are listed below. Functions are denoted as such; all other listed utilities are subroutines. For more information refer to the *CONVEX FORTRAN User's Guide*, Chapter 6, "System utilities."

DATE (*buf*)

Returns current date as *dd-mmm-yy*, where *dd* and *yy* are numeric characters and *mmm* is a three-letter abbreviation for the month. *buf* is of type CHAR*9. This subroutine behaves differently in *-cfc* mode; refer to Appendix D, "Cray FORTRAN compatibility."

IDATE (*iarray*)

Returns current date in *iarray*, a 3-element array of type INTEGER. *iarray*(1) contains the day, *iarray*(2) contains the month as a value between 1 and 12 and *iarray*(3) contains the year as a value ≥ 1969 .

ERRSNS (*fnum*, *rmssts*, *rmsstv*, *iunit*, *condval*)

Returns information about last runtime error in *fnum*; remaining arguments are not used. All arguments are of type INTEGER*4.

EXIT (*status*)

Ends a process and makes the argument *status* available to the parent process. *status* is of type INTEGER*4.

SECNDS (*x*)

(function) Returns system time minus the value of its argument in seconds. Both *x* and *secnds* (*x*) are of type REAL*4.

TIME (*buf*)

Returns current system time in an ASCII string as *hh:mm:ss*. *buf* is of type CHAR*8.

RAN(*i*)

(function) Returns random integer; similar to ConvexOS utility rand, except that ran returns type REAL*4.

MVBITS(*m*, *i*, *len*, *n*, *j*)

Transfers *len* bits from position *i* through *i+len-1* of source location *m* to position *j* through *j+len-1* of destination location *n*; $0 \leq i, j < 32$; $i+len, j+len \leq 32$. All arguments are of type INTEGER*4.

FORTRAN 66 compatibility

B

The CONVEX FORTRAN compiler adheres to the American National Standard FORTRAN 77, X3.9-1978, ISO 1539-1980(E). The default language interpretations are FORTRAN 77. The compiler can, however, compile FORTRAN 66 programs. There are five incompatibilities between American National Standard FORTRAN 77 and FORTRAN 66, X3.9-1966:

- EXTERNAL statement
- DO loop minimum iteration count
- OPEN statement BLANK keyword default
- OPEN statement STATUS keyword default
- X format edit descriptor

The first two incompatibilities are interpreted by the compiler; the rest are interpreted by the runtime system. If your program uses the OPEN statement and you want FORTRAN 66 interpretation rules at runtime, either include a call to `ioinit` in your main program or include the library `I66` in your link by using the option `-LI66` on the `fc` command line. The X format edit descriptor usage must be modified.

Compiling FORTRAN 66 programs

To compile a FORTRAN 66 program, you can modify the program using the following procedure, transforming it into a FORTRAN 77 program, or use the `-F66` option.

1. Use `grep` to identify OPEN statements in which a STATUS keyword is to be added, EXTERNAL statements that must be changed to INTRINSIC statements, and FORMAT statements using the X edit descriptor. These changes can then be made manually in an editor or automatically through use of a shell script.

2. Use the `-F66` option or `OPTIONS` statement to select FORTRAN 66 language interpretations. The `-F66` option allows for the interpretation of `EXTERNAL` statements, `DO` loop minimum iteration counts, and `BLANK` and `STATUS` keyword defaults in `OPEN`. It does not affect the `X` format edit descriptor. If you are running the `C` shell, you can avoid including the `-F66` option in the `fc` command each time you wish to invoke the FORTRAN 66 compiler by using an alias. `alias` is a `C` shell command with the following form:

```
alias fc 'fc -F66'
```

You can include this format in your `.cshrc` file, with the `$` parameter representing specified files. For more information, see the `cs`h(1) man page.

EXTERNAL statement

In FORTRAN 66, the `EXTERNAL` statement specifies that a symbolic name is the name of either a user-defined external procedure or a FORTRAN-supplied function. In FORTRAN 77, two statements accomplish this function:

- The `INTRINSIC` statement specifies that the procedure is a FORTRAN-supplied intrinsic procedure, such as `SQRT`.
- The `EXTERNAL` statement specifies that the procedure is user-supplied.

Because of the exact specification of these two procedures, you cannot modify the `EXTERNAL` statements in your program so that the same source program works with FORTRAN 77 and FORTRAN 66. You must substitute an equivalent statement to include the changes, as shown below.

FORTAN 66	FORTAN 77
<code>EXTERNAL USER</code>	<code>EXTERNAL USER</code> (no change)
<code>EXTERNAL SQRT</code>	<code>INTRINSIC SQRT</code>
<code>EXTERNAL *SQRT</code>	<code>EXTERNAL SQRT</code> (where <code>SQRT</code> is a user function, not the intrinsic for the square root)

DO loop minimum iteration count

In FORTRAN 66, the body of a `DO` loop is always executed; in FORTRAN 77, the body of the `DO` loop is not executed if the end condition of the loop is already satisfied when the `DO` statement is executed. To run a FORTRAN 66 program with the FORTRAN

77 compiler, you can either use the -F66 option, or modify the DO statements in the program to ensure a minimum loop count of 1.

For example, in FORTRAN 77, the loop

```
DO 20 J=INIT, LAST
```

is not executed if INIT is greater than LAST, but is executed once in FORTRAN 66.

If this DO statement occurs in a FORTRAN 66 program, its equivalent FORTRAN 77 statement is as follows:

```
DO 20 J=INIT, MAX (INIT, LAST)
```

OPEN statement keywords

While FORTRAN 66 does not contain an OPEN statement, it does allow for many implementations based on FORTRAN 66 which contain an OPEN statement. Both the BLANK and STATUS keywords in OPEN for FORTRAN 77 differ from the implementations that are used under FORTRAN 66.

BLANK keyword

The BLANK keyword affects the treatment of blanks in numeric input fields read with the D, E, F, G, I, O, and Z field descriptors. In FORTRAN 77, unless the -vfc flag is specified or the COVUE shell is used, the OPEN statement BLANK keyword defaults to BLANK='NULL' (which means that blanks in numeric fields are ignored). The FORTRAN 66 interpretation of blanks in numeric input fields is equivalent to BLANK='ZERO'.

When a logical unit is opened without an explicit OPEN statement and the -F66 option is specified on the compiler command line, CONVEX FORTRAN provides a default that is equivalent to BLANK='ZERO'.

The use of BLANK='NULL' causes embedded and trailing blanks to be ignored and the value converted as if the nonblank characters were right justified in the field. However, the use of BLANK='ZERO' causes embedded and trailing blanks to be treated as zeros.

If your program treats blanks in numeric input fields as zeros, and you do not want to use -I166 or ioinit, include BLANK='ZERO' in the OPEN statement.

STATUS keyword

The OPEN statement STATUS keyword in FORTRAN 77 specifies the initial status of the file (OLD, NEW, SCRATCH, or UNKNOWN); its default value is UNKNOWN. In FORTRAN 66, where STATUS is called TYPE, the default value is NEW.


If your program assumes that the default value for TYPE is NEW and you do not want to use -LI66 or ioinit, put STATUS = 'NEW' in the OPEN statement.

X descriptor

The FORTRAN 66 implementation of the X format edit descriptor writes blanks to the output record and can extend it. The FORTRAN 77 version does not modify character positions that are skipped and does not, as a result, affect the length of the output record.

Format code separators

Formats without format code separators are supported.



The CONVEX FORTRAN compiler adheres to the American National Standard programming language FORTRAN 77, X3.9-1978, ISO 1539-1980(E). The default language interpretations are FORTRAN 77 with default CONVEX extensions; however, when the `-f90` flag is used on the compiler command line, certain features of the International Fortran Standard, ISO/IEC 1539.1991, which is identical to the ANSI Fortran 90 programming language, ANSI X3.198-1992, are enabled. These features include:

- Array sections
- Automatic arrays
- Allocatable arrays
- Certain Fortran 90 array manipulation intrinsics
- Array assignments
- the `WHERE` statement and construct

Each of these features is explained in the appropriate section of this manual. This appendix consolidates detailed descriptions of each feature, including usage examples where appropriate.

Array sections

Use of the `-f90` compiler flag allows you to use array sections as operands in assignment statements and as arguments to many intrinsic functions as specified in Appendix A, "Intrinsics and commonly used library routines." An array section is a piece of an array bounded by specified elements in each dimension. The typical example of a rectangular section of a two-dimensional array is presented further on. An array section is denoted with the following form:

$$\text{arrname} (i:j[:\text{step}] [, \dots])$$

where

arrname

is the array name.

i

is a constant, variable, or expression describing the element at which the section starts for that particular dimension of the array. *i* defaults to the declared lower bound of that dimension of the array.

j

is a constant, variable, or expression describing the element at which the section ends for that particular dimension of the array. *j* defaults to the declared upper bound of that dimension of the array.

step

is a constant, variable, or expression describing the number of elements to step along that particular dimension when selecting elements for the array section. *step* defaults to 1.

Given the defaults for each of these values, array section specifiers such as $X(: :)$ and $X(:)$ are allowed. These particular examples are equivalent to the entire array, which can also be denoted by X .

Note

Array section notation has similar syntax to rank definition, but differs according to usage context.

Following is an example of how a rectangular section of a two-dimensional array is denoted:

$$Y = X(5:10, 1:20:2)$$

This example assigns a section of the array X consisting of rows 5 through 10 and odd-numbered columns 1 through 20 to the array Y (a step of 2 starting at column 1 yields odd-numbered columns). This section has rank 2 and shape (6,10). Y must be declared as a two-dimensional array with exactly the same shape and rank as the section. Any time an array section is assigned to another array section, the array sections must have identical shape. Array sections of shape (1)—for example, $A(3:3, 1:1)$ —are still considered arrays and therefore are not accepted where a scalar variable is required. In other words, $A(3:3, 1:1)$ is not equivalent to $A(3, 1)$; the former describes a single element array section and the latter describes a scalar value.

Automatic arrays

The `-f90` and `-cfc` compiler options enable the use of automatic arrays. Any array declared local to a subroutine with one or more non-constant dimensions is considered an automatic array. Memory for automatic arrays is dynamically allocated on the stack on entry into the subroutine based on a variable dimension value passed into the subroutine (this is explained in detail further on). This stack space is freed on exit from the subroutine; automatic arrays cannot be saved using the `SAVE` statement.

Automatic arrays are declared with the following form:

```
type arrname (n [, . . . ])
```

or

```
DIMENSION arrname (n [, . . . ])
```

where

type

is any valid CONVEX FORTRAN type statement, such as `INTEGER`, `REAL`, `COMPLEX`, `CHARACTER*n`, and so on.

arrname

is the name of an array that is local to a subroutine.

n

is a constant, variable, or expression consisting of or derived from an integer argument passed to the subroutine in which the automatic array resides. *n* represents the desired size of a certain dimension of the automatic array. *n* can be of the form *lb:ub*, where *lb* and *ub* specify dimension bounds as defined for conventional arrays. *lb* and *ub* can be constants, variables or expressions. *ub* must be greater than or equal to *lb*. At least one of the specified dimensions must be non-constant for *arrname* to be an automatic array.

Different dimensions of a multidimensional automatic array can be declared to be different sizes based on one or more dimension arguments passed to the subroutine.

The following example illustrates two ways of declaring multidimensional automatic arrays.

```

SUBROUTINE SUB(I,J,K)
.
.
.
DIMENSION IARR(I,J,2*I) !IMPLICITLY TYPED INTEGER
REAL*8 X(J+I,K)
.
.
.

```

Here the arrays IARR and X are local to the subroutine SUB. I, J and K are integer arguments passed to the subroutine that are used in allocating the arrays. The arrays are automatically allocated at runtime on entry into the subroutine. IARR has one dimension of length I, one of length J, and one of length 2*I. X has one dimension of length J+I and one of length K. Both are automatically deallocated on exit from the subroutine.

Automatic arrays can only be used in subroutines and functions; they are not allowed in MAIN programs or BLOCK DATA subprograms or in COMMON blocks, nor can they be used as dummy arguments.

Allocatable arrays

Use of the `-f90` or `-cfc` flag enables the use of Fortran 90 allocatable arrays. Storage for allocatable arrays is dynamically allocated on the heap at runtime when an `ALLOCATE` statement is executed. The heap space must be deallocated through use of the `DEALLOCATE` statement when the allocatable array is no longer needed. Allocatable arrays must be declared as such with the other specification statements in the program before they can be allocated.

The `ALLOCATABLE` statement

Allocatable arrays are declared with the `ALLOCATABLE` statement. The array's rank must be supplied either in the `ALLOCATABLE` statement, in a `DIMENSION` statement, or in the array's type declaration, which, if used, precedes the `ALLOCATABLE` statement. Rank definitions have the following form:

```
arrname ( : [ , : . . . ] )
```

where *arrname* is the name of the array, which is followed by parentheses containing one colon for each dimension of the array. Multiple colons are separated by commas. Code examples

of this are given further on in this section. Note that the above example is not a Fortran 90 statement, but rather a form for use within certain Fortran 90 statements where a rank definition is required. Array rank and shape are discussed in detail in Chapter 3, "Arrays and substrings."

Allocatable array declarations have the following form:

```
[type arrexp]  
ALLOCATABLE (arrexp [, ...])
```

where

type

is an optional type definition for the array. The form for this is identical to the form for a standard array type definition, except instead of specifying constant dimension parameters, you must either supply a rank definition or provide no parameters.

arrexp

is the array name or rank definition if it was not given in a preceding type statement. The array name may occur in both the type and ALLOCATABLE statements; the rank definition must occur in exactly one, or the compiler flags an error.

Example:

```
INTEGER A(:), B  
ALLOCATABLE (A, B(:, :), X(:, :, :), I(:))
```

In this example, arrays A and B are declared type INTEGER and A is given rank one in the type declaration. In the ALLOCATABLE statement, then, A cannot be given a rank. B is given rank two, and it will take the type INTEGER from the type statement above. X is given rank three and implicitly typed REAL. I is implicitly typed INTEGER, and given rank one.

The ALLOCATE and DEALLOCATE statements

After declaring an allocatable array, you must allocate storage for it before you can use it. This is done with the ALLOCATE statement, which has the following form:

```
ALLOCATE (arrdef [, ...])
```

where *arrdef* is an array definition with dimension values or ranges for all dimensions. Dimension ranges are described under the "Array declaration" section of Chapter 3.

Example:

```
ALLOCATE (A(6),B(-9:0,0:9),X(5,10,-50:-40),I(10))
```

Recall that the array types have either been declared in preceding type statements or are implicitly typed. A is allocated space for a 6-element one-dimensional array; B is allocated space for a 10-by-10 two-dimensional array, with row subscripts numbered from -9 to 0 and column subscripts numbered from 0 to 9; X is allocated space for a 5-by-10-by-11 three-dimensional array, with subscripts of the last dimension numbered from -50 to -40; and I is allocated space for a ten-element one-dimensional array.

When you are done using an allocatable array, deallocate the array's storage space using the DEALLOCATE command. You can deallocate multiple arrays with one statement, or, if you finish with them at different points in the program, individually. The DEALLOCATE statement has the following form:

```
DEALLOCATE (arrname [, ...])
```

where *arrname* is the name of the array to be deallocated. *arrname* must be free of subscripts.

Example:

```
DEALLOCATE(A, B, X, I)
```

This example deallocates space for the arrays A, B, X, and I. Because information contained in a deallocated array is lost, an error occurs if you try to access a deallocated array.

You must manually deallocate arrays that are allocated local to a subroutine before exiting the subroutine.

To change the size or subscript range of a presently allocated allocatable array, you must first deallocate the array, then allocate it with the new information.

Fortran 90 array manipulation intrinsics

CONVEX FORTRAN includes a subset of Fortran 90 array manipulation intrinsics which can be accessed through use of the `-f90` command line option. In CONVEX FORTRAN, Fortran 90 intrinsics cannot be nested. These intrinsics are explained in the following subsections.

Many Fortran 90 intrinsics provide for optional arguments. The ANSI Fortran 90 standard allows the *keyword = argument* syntax in the argument list so that these optional arguments can be skipped. CONVEX FORTRAN does not support this syntax. In CONVEX FORTRAN, if you want to omit an optional argument that is not the final argument in the argument list, you must supply a 0 in its place in the argument list. Optional arguments that fall at the end of the argument list do not require a place holder; they can be left out of the list.

Vector and matrix multiply functions

These routines are used to multiply vectors by vectors, matrices by matrices and vectors by matrices. A *vector* is defined as an array of rank one. A *matrix* is defined as an array of rank greater than or equal to one.

DOT_PRODUCT

DOT_PRODUCT performs dot-product multiplication of two numeric or logical vectors. It has the following form:

DOT_PRODUCT (*vecta*, *vectb*)

where

vecta

must be an array valued vector of numeric type (INTEGER, REAL or COMPLEX) or of type LOGICAL.

vectb

must be an array valued vector of numeric type if *vecta* is of numeric type or of type LOGICAL if *vecta* is of type LOGICAL. *vectb* must be the same size as *vecta*.

DOT_PRODUCT returns a scalar value. For numeric arguments, the type of this value is the same as the type of (*vecta* × *vectb*); for logical arguments, the result is the same as the type of (*vecta* .AND. *vectb*).

MATMUL

MATMUL performs matrix multiplication of two numeric or logical matrices. It has the following form:

MATMUL (*mata*, *matb*)

where

mata

must be an array valued rank one or rank two matrix of numeric type or type LOGICAL.

matb

must be an array valued rank one or rank two matrix of numeric type if *mata* is of numeric type or of type LOGICAL if *matb* is of type LOGICAL. If *mata* has rank one, *matb* must have rank two; if *matb* has rank one, *mata* must have rank two. The size of the first (or only) dimension of *matb* must equal the size of the last (or only) dimension of *mata*.

MATMUL returns an array with the same type as that of (*mata* × *matb*); for logical arguments, the result is the same as the type of (*mata* .AND. *matb*). The shape of the array is determined by the arguments:

- If *mata* has shape (*n,m*) and *matb* has shape (*m,k*), the result has shape (*n,k*).
- If *mata* has shape (*m*) and *matb* has shape (*m,k*), the result has shape (*k*).
- If *mata* has shape (*n,m*) and *matb* has shape (*m*), the result has shape (*n*).

Reduction functions

These functions perform various reduction operations on arrays. All reduction functions have at least one argument (an array) and can have one or two additional optional arguments.

In CONVEX FORTRAN, when a function with two optional arguments is called and the middle argument is to be omitted, the middle argument must be given a value of 0 as a place holder.

ALL

ALL determines whether all values of an array along an optional dimension are true. ALL has the following form:

$$\text{ALL}(\text{mask}[, \text{dim}])$$

where

mask

an array of type LOGICAL.

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq \text{dim} \leq n$, where n is the rank of *mask*. If *dim* is omitted, ALL is applied to the entire array and yields a scalar result.

ALL always yields a result of type LOGICAL. It is scalar if *dim* is omitted or if *mask* has rank one. Otherwise, the result is an array of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{\text{dim}-1}, d_{\text{dim}+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *mask*. In other words, the shape is identical to the shape of *mask* except it is missing the dimension specified in *dim*.

ANY

ANY determines whether any value in an array along an optional dimension is true. ANY has the following form:

$$\text{ANY}(\text{mask}[, \text{dim}])$$

where

mask

an array of type LOGICAL.

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq \text{dim} \leq n$, where n is the rank of *mask*. If *dim* is omitted, ANY is applied to the entire array and yields a scalar result.

ANY always yields a result of type LOGICAL. It is scalar if *dim* is omitted or if *mask* has rank one. Otherwise, the result is an array of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{\text{dim}-1}, d_{\text{dim}+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *mask*.

COUNT

COUNT counts the true elements in an array along an optional dimension. COUNT has the following form:

COUNT (*mask* [, *dim*])

where

mask

an array of type LOGICAL.

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq dim \leq n$, where n is the rank of *mask*. If *dim* is omitted, ANY is applied to the entire array and yields a scalar result.

COUNT always yields a result of type INTEGER. It is scalar if *dim* is omitted or if *mask* has rank one. Otherwise, the result is an array of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{dim-1}, d_{dim+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *mask*.

MAXVAL

MAXVAL returns the maximum value of the elements of an array along an optional dimension corresponding to the true elements of an optional mask. MAXVAL has the following form:

MAXVAL (*array* [, *dim* [, *mask*]])

where

array

an array of type INTEGER or REAL.

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq dim \leq n$, where n is the rank of *mask*.

mask (optional)

must be an array of type LOGICAL that is conformable with *array*.

While both *dim* and *mask* are optional arguments, in CONVEX FORTRAN, if you want to omit *dim* and supply *mask*, you must supply a 0 in *dim*'s position.

MAXVAL yields a result of the same type as *array*. It is scalar if *dim* is omitted or if *array* has rank one. Otherwise, the result is an array of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{dim-1}, d_{dim+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *array*.

MINVAL

MINVAL returns the minimum value of the elements of an array along an optional dimension corresponding to the true elements of an optional mask. MINVAL has the following form:

MINVAL (*array* [, *dim* [, *mask*]])

where

array

an array of type INTEGER or REAL.

dim (optional)

must be a scalar of type INTEGER with a value in the range $(1 \leq dim \leq n)$, where n is the rank of *mask*.

mask (optional)

must be an array of type LOGICAL that is conformable with *array*.

While both *dim* and *mask* are optional arguments, in CONVEX FORTRAN, if you want to omit *dim* and supply *mask*, you must supply a 0 in *dim*'s position.

MINVAL yields a result of the same type as *array*. It is scalar if *dim* is omitted or if *array* has rank one. Otherwise, the result is an array of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{dim-1}, d_{dim+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *array*.

PRODUCT

PRODUCT returns the product of the elements of an array along an optional dimension corresponding to the true elements of an optional mask. PRODUCT has the following form:

PRODUCT (*array* [, *dim* [, *mask*]])

where

array

an array of type INTEGER, REAL or COMPLEX.

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq dim \leq n$, where n is the rank of *mask*.

mask (optional)

must be an array of type LOGICAL that is conformable with *array*.

While both *dim* and *mask* are optional arguments, in CONVEX FORTRAN, if you want to omit *dim* and supply *mask*, you must supply a 0 in *dim*'s position.

PRODUCT yields a result of the same type as *array*. It is scalar if *dim* is omitted or if *array* has rank one. Otherwise, the result is an array of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{dim-1}, d_{dim+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *array*.

SUM

SUM returns the sum of the elements of an array along an optional dimension corresponding to the true elements of an optional mask. SUM has the following form:

```
SUM(array[ , dim[ , mask] ])
```

where

array

is an array of type INTEGER or REAL.

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq dim \leq n$, where n is the rank of *mask*.

mask (optional)

must be an array of type LOGICAL that is conformable with *array*.

While both *dim* and *mask* are optional arguments, in CONVEX FORTRAN, if you want to omit *dim* and specify a *mask*, you must supply a 0 in *dim*'s position.

SUM yields a result of the same type as *array*. It is scalar if *dim* is omitted or if *array* has rank one. Otherwise, the result is an array of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{dim-1}, d_{dim+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *array*.

Construction functions

The functions listed in this section are used to construct arrays of all types.

MERGE

MERGE constructs an array by merging two source arrays under a mask. MERGE has the following form:

MERGE(*tsource*, *fsource*, *mask*)

where

tsource

is an array or scalar of any type. The elements in this array are masked by the `.TRUE.` elements of *mask*. If *tsource* is a scalar and *fsource* is an array, the value of *tsource* is broadcast to cover the shape of *fsource*.

fsource

is an array or scalar of the same type as *tsource*. In CONVEX FORTRAN, *fsource* must be conformable with *tsource*. The elements in this array are masked by the `.FALSE.` elements of *mask*. If *fsource* is a scalar and *tsource* is an array, the value of *fsource* is broadcast to cover the shape of *tsource*.

mask

is an array of type LOGICAL that represents a mask under which *tsource* and *fsource* are merged. *mask* must be conformable with *tsource*.

MERGE returns an array of the same type as *tsource*. The resulting array consists of elements of *tsource* that correspond to the `.TRUE.` elements of *mask* and elements of *fsource* that correspond to the other elements of *mask*. If *tsource* is an array, the result is of the same shape; if *tsource* is a scalar but *mask* is an array, the result is of the same shape as *mask*. If all three arguments are scalars, the result is *tsource* if *mask* is `.TRUE.`, *fsource* otherwise.

PACK

Packs the elements of an array into an array of rank one under control of a mask. **PACK** has the following form:

PACK (*array*, *mask*, *vector*)

where

array

is an array of any type.

mask

must be an array or scalar of type LOGICAL and must be conformable with *array*. The resulting array will contain only those elements of *array* that correspond to .TRUE. elements of *mask*.

vector

is a rank one array of the same type as *array*, containing at least as many elements as there are .TRUE. elements in *mask*. If *mask* is a scalar, *vector* must contain at least as many elements as *array*.

PACK returns a rank one array with the same type as *array*. The resulting array is the same size as *vector*. If *vector* contains more elements than those masked out of *array*, the masked *array* elements fill *vector* consecutively from its beginning and any leftover *vector* elements are unchanged. Note that while *vector* is an optional argument according to the ANSI Fortran 90 standard, it is required in the CONVEX FORTRAN implementation of **PACK**.

SPREAD

SPREAD replicates an array into a specified new dimension a specified number of times. **SPREAD** has the following form:

SPREAD (*source*, *dim*, *ncopies*)

where

source

is an array or scalar of any type. It must have a rank of less than seven.

dim

must be a scalar of type INTEGER with a value in the range $1 \leq dim \leq n+1$, where n is the rank of *source*. *dim* specifies the dimension to add.

ncopies

must be a scalar of type INTEGER. Specifies the number of copies of *source* to be made.

SPREAD returns an array of the same type as *source* with rank $n+1$, where n is the rank of *source*. If *source* is scalar, the shape of the result is $(\text{MAX}(0, \text{ncopies})$; if *source* is an array with shape (d_1, d_2, \dots, d_n) , the shape of the result is $(d_1, d_2, \dots, d_{dim-1}, \text{MAX}(0, \text{ncopies}), d_{dim}, \dots, d_n)$

UNPACK

Unpacks a rank one array into an array under control of a mask. UNPACK has the following form:

UNPACK (*vector*, *mask*, *field*)

where

vector

is a rank one array of any type. It must have at least as many elements as there are .TRUE. elements in *mask*. *vector* contains the array to be unpacked.

mask

must be an array of type LOGICAL. *mask* provides the shape that *vector* will be unpacked into.

field

must be of the same type as *vector* and must be conformable with *mask*; scalar values are acceptable.

UNPACK returns an array of the same type as *vector* having the shape of *mask*. The positions in the result that correspond to the positions of the .TRUE. elements of *mask* are filled, in array element order, with the values from *vector*. All other elements of the result contain the value of *field* if *field* is a scalar, or corresponding elements of *field* if *field* is an array.

Manipulation functions

These functions are used to manipulate arrays.

CSHIFT

CSHIFT performs a circular shift on the positions of elements parallel to a specified dimension of the array; elements shifted off one end reappear at the other end.

CSHIFT has the following form:

CSHIFT (*array*, *shift*, *dim*)

where

array

is an array of any type.

shift

must be of type integer and must be scalar if *array* has rank one; otherwise, *shift* must be scalar or of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{dim-1}, d_{dim+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *array*.

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq dim \leq n$, where n is the rank of *array*. If *dim* is omitted, it is as if it were present with the value 1.

CSHIFT returns an array of the same type and shape as *array*.

EOSHIFT

EOSHIFT performs an end-off shift on an array expression of rank one or on each complete rank-one section along a given dimension of an array expression of rank two or greater. Elements are shifted off at one end and copies of a boundary value are shifted into vacant elements at the other end. Different sections may have different boundary values and may be shifted by different amounts and in different directions.

EOSHIFT has the following form:

EOSHIFT (*array*, *shift*, *boundary*, *dim*)

where

array

is an array of any type.

shift

must be of type integer and must be scalar if *array* has rank one; otherwise, *shift* must be scalar or of rank $n-1$ and of shape $(d_1, d_2, \dots, d_{dim-1}, d_{dim+1}, \dots, d_n)$, where (d_1, d_2, \dots, d_n) is the shape of *array*.

boundary (optional)

must be of the same type and type parameters as *array* and must be scalar if *array* has rank one; otherwise, it must be either scalar or of rank $n - 1$ and of shape (d_1, d_2, \dots, d_n) . *boundary* may be omitted for the data types shown below and, in this case, it is as if it were present with the scalar value shown.

Type of <i>array</i>	Value of <i>boundary</i>
INTEGER	0
REAL	0.0
COMPLEX	(0.0, 0.0)
LOGICAL	.FALSE.
CHARACTER* (<i>len</i>)	<i>len</i> blanks

dim (optional)

must be a scalar of type INTEGER with a value in the range $1 \leq dim \leq n$, where n is the rank of *array*. If *dim* is omitted, it is as if it were present with the value 1.

EOSHIFT returns an array of the same type and shape as *array*.

TRANPOSE

TRANPOSE returns the transpose of an array of rank two. TRANPOSE has the following form:

TRANPOSE (*matrix*)

where

matrix

is a matrix of any type. Must be of rank two.

TRANPOSE returns a rank two array of the same type as *matrix* and with shape (m, n) where (n, m) is the shape of *matrix*. In other words, the rows and columns of *matrix* are swapped.

Location functions

These functions are used to locate the maximum and minimum values of the array.

MAXLOC

MAXLOC locates the maximum value in an array (along an optional mask if present). MAXLOC has the following form:

MAXLOC (*array* [, *mask*])

where

array

must be an array of type INTEGER or REAL. *array* must not be a scalar.

mask (optional)

must be a type LOGICAL array that is conformable with *array*.

MAXLOC returns a rank one array of type INTEGER and of size equal to the rank of *array*. The elements of the result are the subscripts of *array* which locate the maximum value of *array*. If this value occurs in more than one element of *array*, the location of the first occurrence is returned. If *mask* is present, the element returned is the maximum of those elements of *array* that correspond to true elements of *mask*.

MINLOC

MINLOC locates the minimum value in an array (along an optional mask if present). MINLOC has the following form:

MINLOC (*array* [, *mask*])

where

array

must be an array of type INTEGER or REAL. *array* must not be a scalar.

mask (optional)

must be a type LOGICAL array that is conformable with *array*.

MINLOC returns a rank one array of type INTEGER and of size equal to the rank of *array*. The elements of the result are the subscripts of *array* which locate the minimum value of *array*. If this value occurs in more than one element of *array*, the location of the first occurrence is returned. If *mask* is present, the element returned is the minimum of those elements of *array* that correspond to true elements of *mask*.

Array assignments

CONVEX FORTRAN supports both the use of array-valued expressions in assignment statements and masked array assignments (WHERE statements and constructs) under the `-f90` compiler option. Both these methods of manipulating array data are discussed in this section.

Array-valued expressions

The ability to use array-valued expressions in assignment statements is supported by CONVEX FORTRAN under the `-f90` flag. This feature allows you to assign a value or expression to an entire array or array section with one assignment, using the following form:

$$arr = ex$$

where *arr* is the name of an array or an array section description and *ex* is an expression.

As with assignment to a scalar variable, mismatched expression types are converted to the type of the argument on the left before assignment.

Example:

```
INTEGER IX(10)
REAL X
      .
      .
      .
IX = IX * X
```

Here, each element of *IX* is multiplied by the *REAL* variable *X* and truncated to an *INTEGER*. The result then replaces the original element in the array.

Note

Fortran 90 array assignments are translated into loops by the compiler. Any optimization options specified at compile time are then applied to the generated loop; for instance, the loop would be vectorized at optimization options `-o2` and higher.

Array sections can be similarly assigned. Refer to the section "Array sections" in this appendix for more information.

Masked array assignment

When the `-f90` compiler option is specified, CONVEX FORTRAN allows assignment of values to an array under a mask specified via the `WHERE` statement or `WHERE` construct. The `WHERE` statement evaluates a logical expression to determine to which elements the assignment is being applied. The `WHERE` construct works similarly, but is terminated with the `ENDWHERE` statement. It can contain several assignments and an `ELSEWHERE` statement, which allows alternate assignments to be applied to the complement of the mask-expression.

The `WHERE` statement has the following form:

```
WHERE (mask-expr) assignment-stmt
```

The `WHERE` construct has the following form:

```
WHERE (mask-expr)  
  [assignment-stmt]  
  [...]  
  [ELSEWHERE  
    [assignment-stmt]  
    [...]]  
ENDWHERE
```

where

mask-expr

is a logical expression of the same shape as the array(s) being manipulated in the *assignment-stmt*(s).

assignment-stmt

is an array assignment. The array must be the same shape as the array in *mask-expr*. If a function is used here, any array arguments it takes must also be of the same shape.

On execution, the *mask-expr* is evaluated. Any following assignments are executed only on array elements corresponding to those elements for which *mask-expr* evaluated to `.TRUE.` If an `ELSEWHERE` statement is present, its assignments are applied to array elements corresponding to those elements for which *mask-expr* evaluated to `.FALSE.`

In the WHERE construct, *mask-expr* is evaluated once at the beginning of the construct and the result stored for use in every *assignment-stmt*. Execution of the construct then proceeds as if each *assignment-stmt* was part of a WHERE statement using the original evaluation of *mask-expr*, as demonstrated below.

```
WHERE (mask-expr)  
    assignment-stmt1  
    assignment-stmt2  
ELSEWHERE  
    assignment-stmt3  
ENDWHERE
```

is equivalent to:

```
WHERE (mask-expr) assignment-stmt1  
WHERE (mask-expr) assignment-stmt2  
WHERE (.NOT. (mask-expr)) assignment-stmt3
```

Each WHERE statement is then further decomposed into a DO-loop by the compiler to facilitate the actual assignment of values into the array elements. It is important to note that:

- *mask-expr* is evaluated only once before entering the body of the WHERE and does not change over the course of the construct, even if the array on which it is based is changed in the WHERE body.
- Each *assignment-stmt* is executed in full independently of the other *assignment-stmts*. This means that if *assignment-stmt* contains an array assignment, the compiler may generate a DO-loop specifically for that assignment. Do not assume that multiple *assignment-stmts* that involve the same array section will execute sequentially as if part of a single DO-loop.
- The compiler-generated DO-loops are subject to the same optimizations as hand-written DO-loops.

Examples of the WHERE statement and construct follow.

Example 1:

```
REAL DATA(1000), OFFSET(10000), ADJUSTED(1000), LIMIT
LOGICAL NORMAL(1000)
LIMIT = 130.0
.
.
.
WHERE (DATA .LE. LIMIT) NORMAL = .TRUE.
```

In this example, all elements of the logical array `NORMAL` that have the same index as elements in the real array `DATA` that are less than or equal to `LIMIT` are set to `.TRUE.`

The following example assumes the same variable declarations as Example 1.

Example 2:

```
WHERE (DATA .GT. LIMIT)
  ADJUSTED = FIX (DATA)
  NORMAL = .FALSE.
ELSEWHERE
  ADJUSTED = 0.0
  NORMAL = .TRUE.
ENDWHERE
```

In this example, elements of `ADJUSTED` corresponding to elements of `DATA` greater than 130.0 are set to `FIX (DATA)`, and all other elements of `ADJUSTED` are set to 0.0. Similarly, all elements of `NORMAL` corresponding to `DATA` elements greater than 130.0 are set to `.FALSE.` and all other elements of `NORMAL` are set to `.TRUE.`

Array sections and subscript expressions can be used with the `WHERE` construct to change the correspondence of elements between arrays. The following example assumes the same variable declarations as Example 1.

Example 3:

```
WHERE (DATA .LE. LIMIT)
  OFFSET ( ISET : ISET+999 ) = DATA
ELSEWHERE
  OFFSET ( ISET : ISET+999 ) = 0.0
ENDWHERE
```

In this example, elements of OFFSET starting at the value ISET and going through ISET+999 are set to correspond to elements of DATA where the DATA elements are less than or equal to LIMIT, and are set to 0.0 where the DATA elements are greater than LIMIT.

Note

It is possible to assign values to the same array element in both the `.TRUE.` and `.FALSE.` branches of the `WHERE` construct when using array sections or subscript expressions that differ in each branch of the construct. This action can inhibit vectorization and parallelization of the `WHERE` construct, and should therefore be avoided.

Remember, the `WHERE` construct differs from the block-`IF` in that both clauses can and often do execute. It is therefore possible for assignments that execute in the `ELSEWHERE` clause of the `WHERE` construct to change values assigned in the `WHERE` clause when array section notation is used. The following example, which illustrates this, assumes the same variable declarations as Example 1.

Example 4:

```
WHERE (DATA .LE. LIMIT)
  OFFSET (ISET:ISET+999:2) = DATA(1:1000:2)
ELSEWHERE
  OFFSET (ISET+1:ISET+1000:2) =
^  OFFSET (ISET:ISET+999:2)
ENDWHERE
```

In this example, every other element from `OFFSET (ISET)` through `OFFSET (ISET+999)` is set in the `WHERE` clause. Then in the `ELSEWHERE` clause, because of the sectioning notation used, each value assigned in the `WHERE` clause is copied into the element immediately following it in `OFFSET`. The end result is that `OFFSET (ISET) = OFFSET (ISET+1) = DATA (1)`, `OFFSET (ISET+2) = OFFSET (ISET+3) = DATA (3)` and so on for every other element of `OFFSET`.

Cray FORTRAN compatibility

D

To facilitate porting code written for the Cray FORTRAN compiler, the `-cfc` option can be specified on the `fc` command line. This appendix explains the effects this option has on the compiler.

Compiler defaults

The `-cfc` option changes compiler defaults in the following ways:

- Default data types are as follows:

Type	Default length
Integer	INTEGER*8
Real	REAL*8
Complex	COMPLEX*16
Logical	LOGICAL*8
Double precision	REAL*16

- Constants are stored in the default INTEGER or REAL type.
- Constants written in single-precision exponential form (such as 1.23E4) are stored in REAL*8 format.
- Ininsics that work with default integer and single-precision types work with Cray default types (8-byte quantities).
- LOGICAL*2 and *4, INTEGER*4, and REAL*4 are treated as LOGICAL*8, INTEGER*8, and REAL*8.
- Double-precision (REAL*16) data types are supported. Constants written with a D in the exponent or objects declared double-precision are treated as 16-byte objects.
- Logical constants .T. and .F. are supported.

Note

The CONVEX/VAX representation of REAL*16 data gives more precision (113 bits) than the Cray, IBM, and IEEE specifications. This greater accuracy makes the CONVEX REAL*16 software much slower (possibly 40 times slower) than Cray double-precision software because of the greater number of intermediate results that must be generated and saved internally.

Unsupported Cray features

Only those Cray-specific intrinsics enumerated in the "Supported Cray intrinsics" section of this appendix are supported. Cray-specific directives are not supported.

In addition, when a program uses the LOC function, Cray FORTRAN returns a word address; CONVEX FORTRAN returns a byte address.

Cray POINTER support

The Cray POINTER statement is supported. However, by default pointer arithmetic computes byte addresses rather than word addresses. Due to this difference you must either multiply your pointer offsets by eight or use the `-cfcwpa` compiler option. For pointers that are specifically declared using the POINTER statement, addition or subtraction of pointers will compute word addresses under the `-cfcwpa` option. The compiler will issue a warning when it encounters other arithmetic operations on pointers. If you do not specify the `-cfcwpa` option and the compiler encounters an explicitly declared pointer in an arithmetic context, a warning message is issued. The `-cfcwpa` option must be used with the `-cfc` option.

Note

The `-cfcwpa` option does not recognize pointers that have not been specifically declared using the POINTER statement. Byte arithmetic will be performed on pointers that are not explicitly declared.

For more information on the `-cfcwpa` option refer to Chapter 1, "Compiling programs," of the *CONVEX FORTRAN User's Guide*.

The syntax for the POINTER statement is

```
POINTER (p, s)
```

where *p* is a pointer, and *s* is the name of a local variable or array. In Cray terminology, *s* is the pointee. *s* cannot be associated with any other known piece of named and referenced storage except through assignments to *p* or by associating two or more pointees with one pointer.

No two pointers can point to the same storage, but this optimization requirement is not enforced by the compiler.

Pointees are assumed *not* to overlap. This means that a value stored indirectly through one pointee must not be accessed via another pointee even when both pointees are associated with the same pointer.

Debugging code containing Cray pointers

`csd`, the CONVEX Symbolic Debugger, is an optional product available as part of the CONVEX Consultant package. It cannot access Cray pointers. However, you can dump the contents of the pointer in `csd`; this gives you the raw memory address of the pointee. You can then view the contents of this address in `csd`.

Example:

```
PROGRAM POINT
REAL A(10)
POINTER (IPA, A)
CALL HPALLOC (IPA, 10, IA, IR)
DO I = 1, 10
  A(I) = I
END DO
PRINT *, A
END
```

To examine the contents of `A` from within `csd`, first get its address (the value of `IPA`):

```
(csd) p ipa
2148233220
```

Now you can examine the memory locations beginning at 2148233220 to see the contents of the array `A`:

```
(csd) 2148233220,10?F
800b9004:  1.000000  2.000000
800b9014:  3.000000  4.000000
800b9024:  5.000000  6.000000
800b9034:  7.000000  8.000000
800b9044:  9.000000 10.000000
```

Refer to the *CONVEX Consultant User's Guide* for more information on the use of `csd`.

Cray automatic arrays

Cray automatic arrays are supported, under both the `-cfc` and `-f90` compiler options. Any array declared local to a subroutine with one or more non-constant dimensions is considered an automatic array. Memory for automatic arrays is dynamically allocated on the stack on entry into the subroutine based on a variable dimension value passed into the subroutine. This stack space is freed on exit from the subroutine; automatic arrays cannot be saved using the `SAVE` statement.

Automatic arrays are described in detail in Chapter 3, "Arrays and substrings" and Appendix C, "FORTRAN 90 compatibility."

Cray BUFFERIN, BUFFEROUT support

CONVEX provides a look-alike implementation of the Cray `BUFFERIN`, `BUFFEROUT` feature and its accompanying routines. This feature is provided mainly to aid in porting existing Cray code; it is not likely to provide noticeably superior performance compared to conventional CONVEX I/O methods. `BUFFEROUT` will always be slightly slower than unformatted fixed record length I/O.

Related statements and routines

The following statements are available for use with `BUFFERIN`, `BUFFEROUT`:

`BUFFERIN (unit, mode) (beginLoc, endLoc)`

`BUFFEROUT (unit, mode) (beginLoc, endLoc)`

These library routines are also available: `UNIT (unit)`, `LENGTH (unit)`, `GETPOS (unit)` and `SETPOS (unit, pos)`. These routines reside in `libcfc.a`, which is implicitly loaded when linking with the `-cfc` option; if you do not compile using `-cfc` these routines will not be available. Old object files will link and run with the new libraries.

Note

If present, data format conversions specified in the `OPEN` statement do not affect data read or written with `BUFFERIN` or `BUFFEROUT` statements.

Restrictions

Note the following restrictions:

- The SETPOS function cannot be used with magnetic tape; doing so produces an error message.
- Using BUFFERIN and BUFFEROUT with data types less than eight bytes produces a fatal compiler error.
- Other FORTRAN I/O statements (READ, WRITE, PRINT, ACCEPT, TYPE) cannot be used on the same unit as BUFFERIN, BUFFEROUT.
- BACKSPACE does not work with BUFFERIN, BUFFEROUT files.

Cray unformatted file support

CONVEX FORTRAN is capable of reading and writing unformatted Cray data files through use of the `FORMAT = UNFORMATTED/CRAY` and `FORMAT = UNFORMATTED/CRAYUB` keyword definitions in the OPEN statement. Formatted files are supported regardless of file structure or access mode; unformatted files are partially supported depending on file structure and access mode. Table 25 shows the degree of support provided given various combinations of file structures and access modes.

Table 25
Unformatted Cray files
readable by CONVEX
FORTRAN

Cray file structure	Access mode	
	Sequential	Direct
Unblocked (pure)	Does not comply with ANSI FORTRAN 77 standard; no record boundaries inherent in file; BACKSPACE not permitted; use <code>FORMAT=UNFORMATTED/CRAYUB</code> in OPEN statement.	CONVEX FORTRAN reads directly.*
Blocked	Must run <code>fcUnblock</code> [†] utility, supplied with CONVEX FORTRAN V7.0, on these files to convert them into a form readable by CONVEX FORTRAN. Use <code>FORMAT=UNFORMATTED/CRAY</code> .*	Cray does not allow these.

*These are the default file structures written in Cray FORTRAN.

[†]For more information on the `fcUnblock` utility, see the `fcUnblock (1F)` man page.

Supported Cray library routines

The following library routines are available using the `-cfc` option:

cft\$bool	cft\$dprod	clock	date	dump
etime	findch	gather	hpalloc	hpcheck
hpc1move	hpdeallc	hpdump	hpnewlen	hpshrink
iceil	igtbyt	ihplen	ihpstat	ilsum
int24	jdate	komstr	lint	mvc
pack	putbyt	rbn	rnb	scopy
sdot	second	ssum	strmov	timef
tr	unpack			

Additionally, most `libU77.a` routines can be accessed in Cray mode. The exceptions are: `dbesj0`, `dbesj1`, `dbesjn`, `dbesy0`, `dbesy1`, `dbesyn`, `derf`, and `derfc`. For more information about `libU77.a` routines, refer to the `intro.3f` man page.

Note

A Cray mode call to `date` accesses the Cray specific `date` routine, which returns the date in a format different from that returned from the `libU77.a` `date` routine. Similarly, a Cray-mode call to `rand` accesses the Cray routine and returns a different sequence of numbers than the corresponding call to the `libU77.a` `rand` routine.

Supported Cray intrinsics

The following intrinsics are available using the `-cfc` option:

and	compl	cot	csmg	cvmgm
cvmgm	cvmgp	cvmgt	cvmgz	dcot
eqv	leadz	loc	mask	movbit
neqv	numarg	or	popcnt	poparr
ranf	ranget	ranset	shift	shiftl
shiftr	xor			

Additionally, all Military Standard (MIL-STD-1753) routines are supported from Cray mode.

Cray TASK COMMON support

CONVEX FORTRAN supports Cray TASK COMMON blocks under the `-cfc` option. These blocks are created using the TASK COMMON statement, which has the following form:

```
TASK COMMON /cbn/nlist [ , /cbn/nlist ] . . .
```

where

cbn

is a symbolic name for a task common block. Unnamed TASK COMMON blocks are not allowed.

nlist

is a list of variable names, array names, and array declarators. These variables cannot appear in a DATA statement, but otherwise can be used like any variables in COMMON storage.

All occurrences of the TASK COMMON block must be declared TASK COMMON; a common block cannot be declared both COMMON and TASK COMMON. TASK COMMON blocks can only be declared in functions, subprograms and BLOCK DATA subprograms. Variables in TASK COMMON blocks are provided thread-local storage. A program should already be running multiple threads before calling a subroutine that contains a TASK COMMON block.

Cray Boolean octal constant support

CONVEX FORTRAN supports Cray Boolean octal constants under the `-cfc` option. Cray Boolean octal constants consist of one or more octal digits followed by the letter B. A Boolean octal constant has the following form:

```
cc...cB
```

where *c* represents an octal digit. There can be up to 22 octal digits in an octal constant. An octal digit can range from 0 to 7.

Like all Cray constants, octal constants occupy 8 bytes (equivalent to one 64-bit Cray word) in memory. It follows that octal constants can be 22 octal digits long, but, because 22 octal digits represent 66 bits in memory, the leftmost octal digit must be either 0 or 1. Octal constants are right justified and zero-filled to the left. An octal constant that occupies all 22 digits cannot have a value greater than 1 as its leftmost octal digit. Blanks within an octal constant are ignored.

Examples:

Valid	Invalid	Reason
765B	835B	8 not in range 0 to 7
1126752354176524376524B	3126752354176524376524B	Leftmost digit > 1

Cray Hollerith constants

CONVEX FORTRAN supports left-justified Cray Hollerith constants of the following form:

$nLcc...c$

or

$'cc...c'L$

where

n

specifies the number of the characters in the constant (including spaces and tabs). The value of n must be an unsigned positive integer greater than zero.

c

is a printable ASCII character.

Cray Hollerith constants occupy 8 bytes (equivalent to one 64-bit Cray word) in memory. Each byte contains one ASCII character code, and Hollerith constants using this form cannot exceed eight characters in length. This form left-justifies the constant in memory and zero-fills its 8-byte storage space to the right. You must supply the the `-cfc` option when using this form.

Example:

Valid	Invalid	Reason
4LHelp	4L'Help'	No quote in this form.
'foo'L	3'foo'L	Number of characters specifier not allowed in this form.

This appendix describes compatibility between VAX FORTRAN and CONVEX FORTRAN. To facilitate porting code written for the VAX FORTRAN compiler, certain VAX features are supported as described below.

Supported features

CONVEX FORTRAN supports the following VAX FORTRAN features only when the `-vfc` option is specified on the `fc` command line:

- VAX INCLUDE statement
- REAL*16 data type
- The alternate form (without parentheses) of the PARAMETER statement with only one constant specified
- The 'r form of the record specifier
- Octal constants in the form `*m`, where `m` is a string of octal digits
- Default file names in the form `FOR0m.DAT`, where the number `m` corresponds to unit number `m`
- VAX FORTRAN records
- Hollerith constants where a CHARACTER value is expected
- The RECL= specifier used in the OPEN and INQUIRE statements. This specifier returns the number of VAX words rather than bytes for unformatted files. (This is not true for programs compiled and loaded separately unless `-vfc` is specified for the load phase.)
- VAX FORTRAN RECORDS, including the RECORD, STRUCTURE, UNION and MAP statements.

The organization and structure of VAX FORTRAN records is discussed briefly at the end of this appendix.

Unsupported features

CONVEX FORTRAN does not support the following VAX FORTRAN features:

- %DESCR
- RMS calls
- The VOLATILE statement
- The zccc...c form of hexadecimal constants
- Interactive display of NAMELIST group and values or end-of-line comments (!) in the NAMELIST input data
- Extra parentheses in WRITE statements (allowed in VAX FORTRAN)
- VMS file names
- Byte ordering with respect to passing characters and parameters
- Logical values. On CONVEX computers, a logical value is true if all the bits are 1. On VAX, it is true in a test if the low-order bit is 1, for example, IF (A) is equivalent to IF (A .AND. 1).
- Numerical differences. The accuracy of CONVEX floating-point representation and the rounding method used cause these differences. Refer to the *CONVEX Architecture Reference* for further information.
- Automatic conversion of REAL*8 (in caller) to REAL*4 (in called subroutine) across subroutine calls
- Calling a function as a subroutine
- Radix 50 constants
- The variable on the left side of a character assignment statement appearing on the right side (VAX FORTRAN extension)
- MOD function defined for a zero denominator
- Modifying an argument within a subroutine if the subroutine was called with a constant in the argument list (the CONVEX FORTRAN compiler enforces this rule, which is an ANSI standard)
- DO statements of the form:

```
DO 714 J=1, 100 WHILE (Q.NE.Z)
```

- VMS path names—The FILE name you specify when using the INQUIRE statement under the COVUE shell must be the absolute UNIX path name.
- The FILE keyword in the OPEN statement. The keyword must specify the name of the file to open as a character expression; numeric variables, arrays or array elements containing the file name cannot be substituted.

CONVEX FORTRAN does not support these VAX FORTRAN I/O extensions:

- REWRITE, DELETE, and UNLOCK statements
- Indexed I/O (key-indexed files)
- File sharing
- DEFINEFILE statement
- OPEN keywords (PRINT and SUBMIT values for DISPOSE; USEROPEN; INITIALSIZE; EXTENDSIZE; BUFFERCOUNT; SEGMENTED for RECORDTYPE; and ORGANIZATION)
- CLOSE keywords (PRINT and SUBMIT values for STATUS)
- ASCII null as a carriage-control character

The internal format of variable-length type records of VAX FORTRAN and CONVEX FORTRAN differ when RECORDTYPE=VARIABLE.

VAX and CONVEX versions of the STOP message differ as follows:

Statement	CONVEX message	VAX message
STOP	STOP:	FORTRAN STOP
STOP 4	STOP: 4	4
STOP 'here'	STOP: here	here

Integer overflow traps are turned off by default in `fc`, whereas VAX FORTRAN enables them by default. The main reasons for this are:

- Overflow is turned off in C, and many users mix C and FORTRAN code.
- It is difficult to optimize integer expressions if integer overflow is turned on because most addresses are negative integers near overflow on a CONVEX machine.

When you use the H descriptor on input, the first character transferred appears immediately after the letter H. The characters that are in the field descriptor before input are replaced (overwritten) by the new input characters.

Miscellaneous differences

The following miscellaneous differences exist between VAX FORTRAN and CONVEX FORTRAN:

- Unit numbers in VAX FORTRAN range from 0 to 99 and in CONVEX FORTRAN from 0 to 255.
- VAX FORTRAN supports the use of Hollerith and apostrophe edit descriptors during formatted input. CONVEX FORTRAN allows Hollerith descriptors but does not allow apostrophe edit descriptors.
- In VAX FORTRAN, 'cc...c' O octal constants are of type INTEGER; in CONVEX FORTRAN, the 'cc...c' form of octal constants is typeless.
- In the ASSIGN statement, ASSIGN *s* TO *i*, where *s* is the label of an executable statement in the current program unit and *i* is an integer variable, VAX FORTRAN sets the integer variable to 1 at initial execution of the ASSIGN statement; CONVEX FORTRAN sets it to 0.
- If invalid data is encountered on a VAX FORTRAN READ statement, all variables on the *iolist* are assigned except those corresponding to the bad data. If the same error occurs in CONVEX FORTRAN, the READ statement ends at once and the remaining variables in the *iolist* are unchanged.
- VAX FORTRAN correctly handles REAL*16 constants that do not contain Q in the exponent. CONVEX FORTRAN does not; if Q is not specified, the constant is considered to be REAL*4.
- CONVEX FORTRAN follows the ANSI standard in that any function referenced directly or indirectly in an I/O statement cannot contain another I/O statement. VAX FORTRAN may not always follow this standard.
- Before VAX binary files can be read by CONVEX FORTRAN, the *cvbin* utility must be used to convert them. Refer to the *cvbin(1)* man page or to the *CONVEX COVUEbinary User's Guide* for more information.

VAX FORTRAN records

A VAX FORTRAN record is a derived data type containing one or more fields. Fields within a record are defined by a structure declaration that defines field names, types of data within fields, and order and alignment of fields within a record.

Structure declaration

A structure declaration is bounded by `STRUCTURE` and `END STRUCTURE` statements and has one or more field declarations. The order in which the field declarations occur determines the order of fields within the structure. At least one field declaration must be specified or an error condition occurs.

A structure declaration has the following format:

```
STRUCTURE /structure-name/
    field declaration
    .
    .
    .
END STRUCTURE
```

A structure declaration does not create a variable. A variable is created by a `RECORD` statement containing the name of a previously declared structure. The `RECORD` statement has the following form:

```
RECORD /structure-name/record-namelist
```

where *structure-name* is the name of a previously declared structure and *record-namelist* is a list of variable names, array names, or array declarations, separated by commas.

Records must be read and written using unformatted I/O. For example:

```
C DEFINITION OF THE NAME STRUCTURE
STRUCTURE / NAME /
    CHARACTER*5 LAST
    CHARACTER*5 FIRST
    CHARACTER*1 INITIAL
END STRUCTURE
```

```

C DEFINITION OF THE DATE STRUCTURE
  STRUCTURE /DATE/
    CHARACTER*2 MONTH
    CHARACTER*2 DATE
    CHARACTER*2 YEAR
  END STRUCTURE

C DEFINITION OF THE PERSON STRUCTURE
  STRUCTURE / PERSON /
    RECORD / NAME / NAME
    LOGICAL*1 SEX
    RECORD / DATE / BIRTH
  END STRUCTURE

C SET UP NUMBER OF EMPLOYEES TO BE HANDLED
  RECORD / PERSON / EMPLOYEES(3)

C READ EMPLOYEE DATA
  DO I = 1,3
    READ(2) EMPLOYEES(1)
  ENDDO

```

Field declaration

A field declaration can be any combination of the following:

- A typed data declaration
- A substructure declaration
- A union declaration.

A typed data declaration is the same as a normal FORTRAN type statement. As with FORTRAN typed data statements, field declarations can contain initializers. The name %FILL can be used in place of a field name to create space in the structure for padding. This space cannot be initialized.

A substructure must be declared by a *RECORD* statement that creates an instance of a previously declared structure, not by a nested *STRUCTURE* statement. The preceding example shows the proper declaration of a substructure within the *PERSON* structure.

A union declaration is bounded by *UNION* and *END UNION* statements and defines a data area that can be shared during program execution. A union declaration must contain at least two map declarations (as indicated below) or an error condition occurs. A union declaration has the following form:

```

UNION
    map declaration
    map declaration
    .
    .
    .
END UNION

```

A *map-declaration* defines a unique group of fields and is bounded by MAP and END MAP statements. A map declaration must contain at least one field declaration or an error condition occurs. A map declaration has the following form:

```

MAP
    field declaration
    .
    .
    .
END MAP

```

VAX floating point data

CONVEX FORTRAN can read VAX floating point data in any format using the `vax_d` (`vax-d`) or `vax_g` (`vax-g`) data format names. The correct data format name is specified in the program's OPEN statement or in a shell variable. Chapter 9, "Input/output statements," explains in detail how data format names can be specified.

Table 26 lists various VAX floating point formats, their respective CONVEX equivalents, and the correct CONVEX data format name to use to read and write the format.

Table 26
VAX floating point data
format names

VAX floating point format	CONVEX equivalent	CONVEX data format name
F_floating	REAL*4	vax_g, vax-g, vax_d or vax-d
D_floating	REAL*8 [†]	vax_d or vax-d
G_floating	REAL*8 [†]	vax_g or vax-g
H_floating	REAL*16 ^{††}	vax_g or vax-g
F_floating	COMPLEX*8	vax_g, vax-g, vax_d or vax-d
D_floating	COMPLEX*16	vax_d or vax-d
G_floating	COMPLEX*16	vax_g or vax-g

[†] All REAL*8 data in a particular datafile must be read/written in either vax_d or vax_g format. The formats can not be mixed.

^{††} The REAL*16 data type is supported only in native floating point mode.

Supported VAX intrinsic

The following VAX intrinsics are available using the `-vfc` option:

ACOSD	AIMAX0	AIMIN0	AJMAX0	AJMIN0
AKMAX0	AKMIN0	ASIND	ATAN2D	ATAND
BITEST	BJTEST	BKTEST	CDABS	CDCOS
CDEXP	CDLOG	CDSIN	CDSQRT	COSD
DACOSD	DASIND	DATAN2D	DATAND	DBLEQ
DCMLPX	DCONJG	DCOSD	DFLOAT	DFLOTI
DFLOTJ	DFLOTK	DIMAG	DIMAX0	DIMIN0
DJMAX0	DJMIN0	DKMAX0	DKMIN0	DMAX0
DMIN0	DREAL	DSIND	DTAND	FLOATI
FLOATJ	FLOATK	IDMAX1	IDMIN1	IIABS
IIAND	IIBCLR	IIBITS	IIBSET	IIDIM
IIDINT	IIDMAX1	IIDMIN1	IIDNNT	IIEOR
IIFIX	IINT	IIOR	IIQINT	IIQNNT
IISHFT	IISHFTC	IISIGN	IMAX0	IMAX1
IMIN0	IMIN1	IMOD	ININT	INOT
INT1	INT2	INT4	INT8	IQINT
IQNINT	IZEXT	JIABS	JIAND	JIBCLR
JIBITS	JIBSET	JIDIM	JIDINT	JIDMAX1
JIDMIN1	JIDNNT	JIEOR	JIFIX	JINT
JIOR	JIQINT	JIQNNT	JISHFT	JISHFTC
JISIGN	JMAX0	JMAX1	JMIN0	JMIN1
JMOD	JNINT	JNOT	JZEXT	KIABS
KIAND	KIBCLR	KIBITS	KIBSET	KIDIM
KIDINT	KIDMAX1	KIDMIN1	KIDNNT	KIEOR
KIFIX	KINT	KIOR	KIQINT	KIQNNT
KISHFT	KISHFTC	KISIGN	KMAX0	KMAX1
KMIN0	KMIN1	KMOD	KNINT	KNOT
KZEXT	QABS	QACOS	QACOSD	QASIN
QASIND	QATAN	QATAN2	QATAN2D	QATAND
QCOS	QCOSD	QCOSH	QDIM	QEXP
QEXT	QEXTD	QFLOAT	QFLOTI	QFLOTJ
QFLOTK	QINT	QLOG	QLOG10	QMAX1
QMIN1	QMOD	QNINT	QSIGN	QSIN
QSIND	QSINH	QSQRT	QTAN	QTAND
QTANH	SIND	SINGLQ	TAND	ZEXT

Sun FORTRAN compatibility

F

To facilitate porting code written for the Sun FORTRAN compiler, the `-sfc` option can be specified on the `fc` command line. When used, this option changes certain aspects of the CONVEX FORTRAN compiler as follows:

- As in the C language, escape sequences using a backslash (\) are supported in character strings to define nonprintable characters. Table 27 lists the supported sequences.

Table 27
Supported Sun FORTRAN
escape sequences

Character	Sequence
newline	\n
tab	\t
form feed	\f
NUL	\0
single quote	\'
double quote	\"
backslash	\\

- The ampersand (&) character in the first nonblank column of a source line indicates that the line is a continuation, regardless of what is in column 6.
- Recursive subroutines and functions are allowed.
- The declarations `AUTOMATIC` and `STATIC` are supported.

The preceding features, which are implemented by CONVEX FORTRAN, represent a subset of Sun FORTRAN features. The Sun definition of real constants as 64-bit numbers is not supported.

Index

Symbols

- " (double quote) character
 - and SUN FORTRAN 271
 - as delimiter 19
- \$ edit descriptor 173
- \$ record delimiter 181
- & (ampersand) character
 - in Sun FORTRAN 271
- ' (single quote) character
 - as delimiter 19
 - in character assignments 66
 - SUN FORTRAN 271
- * character
 - and assumed-size arrays 190 to 191
 - and ENTRY statement 203
 - as descriptor 155
 - in dummy argument list 194
 - runtime format descriptor 177
- * edit descriptor 155, 177
- : edit descriptor 174
- \ character
 - in SUN FORTRAN 271

A

- A edit descriptor 153
 - default value 175
- ABS intrinsic 210
- absolute value 210
- ACCEPT statement 94, 109
- ACCESS keyword 97, 122
- accessing files 97
- ACOS intrinsic 209
 - argument 219
 - result 219
- ACOSD intrinsic 209
 - argument 219
 - result 219
- actual arguments 187 to 189
 - & 202
 - * 202
 - common block elements as 190
 - Hollerith constants 192
 - passing by reference 196
 - passing by value 196
 - to subroutines 201
- address of variable
 - finding with LOC 46, 197
- adjustable arrays 190
- ALL intrinsic 217, 237
- allocatable array declarations 232

- allocatable arrays 50, 232
 - declaration 25, 232
 - Fortran 90 26
 - rank 232
- ALLOCATABLE statement 50, 232
 - form 25, 233
- ALLOCATE statement 26, 233
 - form 26, 233
- allocating memory
 - dynamically via pointers 47
- allocation
 - dynamic memory 47
- ALOG intrinsic
 - argument 219
- ALOG10 intrinsic
 - argument 219
- alternate PARAMETER statement 42
- alternate return arguments 194, 200
- AMAX0 intrinsic 214
- AMIN0 intrinsic 215
- AND
 - bitwise intrinsic 216
 - .AND. operator 34
 - ANINT intrinsic 211
 - ANSI FORTRAN 66 standard 225
 - ANSI FORTRAN 77 standard xvii
 - ANSI Fortran 90 standard 235
 - ANSI standard
 - data type correspondence to CONVEX types 10
 - data types 9
 - formatting 6
 - Fortran 90 229
 - VAX adherence 264
 - ANSI X3.9-1978 xvii
 - ANY intrinsic 217, 237
- apostrophe edit descriptor 155
- arguments 199 to 200
 - actual 187 to 188, 201
 - alternate return 194, 200
 - changing passing convention 196
 - character 192
 - dummy 187 to 188, 199 to 200, 202
 - Hollerith 192
 - passing address 196
 - passing by reference 196
 - passing by value 196
 - procedures as 194
 - variables as dummy 189
- arithmetic expressions 31
 - data types 31, 33
 - involving constants 33
- arithmetic IF statement 78
- arithmetic operators 31

- array assignment
 - Fortran 90 247
 - masked 67, 248
- array declarations 22
 - examples 23
 - form 22
- array elements
 - referencing 27
- array intrinsics
 - Fortran 90 235
- array sections
 - defaults 230
 - example 230
 - form 27, 229
 - rank 230
 - shape 230
- array storage 28
- array subscripts 27
 - defined 27
 - lower bound 23, 49
 - upper bound 23, 49
- array-valued expressions
 - in assignments 67, 247
- arrays 21
 - adjustable 190
 - allocatable 25 to 26, 232
 - and NAMELIST statement 182
 - as dummy arguments 190, 192
 - assigning Fortran 90 66, 247
 - assigning values to 182
 - assumed size 190
 - assumed-length character 193
 - assumed-size 191
 - automatic 22 to 23, 25, 49, 231, 256
 - circular element shift 243
 - conformability 21 to 22
 - construction functions 241
 - dimension limits 49
 - dimensioning 22, 49
 - dynamic 47
 - dynamic via pointers 48
 - end-off element shift 244
 - equivalencing 51
 - finding any true value 237
 - finding dot product 235
 - finding true values 237
 - Fortran 90 assignments 247
 - Fortran 90 sections 229
 - initializing in type-declaration statements 44
 - locating maximum element 246
 - locating minimum element 246
 - location functions 245
 - manipulation functions 241, 243
 - matrix multiplication 236
 - matrix multiply functions 235
 - maximum value 238
 - merging 241
 - minimum value 239
 - multidimensional dynamic 48
 - packing 242
 - product of elements 239
 - rank 21
 - reduction functions 236
 - referencing elements 27
 - replicating 242
 - sections 27, 229 to 230
 - shape 21 to 22
 - specifying bounds 49
 - storage 28
 - subscript bounds 23
 - summing 240
 - transposing 245
 - typing 22, 49
 - unpacking 243
 - vector multiply functions 235
- ASCII characters
 - 8-bit 62
- ASCII values
 - finding 217
 - finding corresponding characters 217
- ASIN intrinsic 208
 - argument 219
 - result 219
- ASIND intrinsic 208
 - argument 219
 - result 219
- ASSIGN statement 73
 - VAX interpretation 264
- assigned GOTO statement 77
- assignment statements 65
 - character 66
 - Fortran 90 array 66
 - truncation in 70
 - using arrays in 67, 247
- assignments
 - Fortran 90 array 247
- associated documents xx
 - American National Standard programming language FORTRAN, ANSI X3.9-1978 xxi
 - CONVEX adb Debugger User's Guide xx
 - CONVEX Application Compiler User's Guide xx
 - CONVEX Compiler Utilities User's Guide xx
 - CONVEX Consultant User's Guide xx
 - CONVEX COVUEshell Reference Manual xxi
 - CONVEX CXdb Concepts xxi
 - CONVEX FORTRAN Optimization Guide xx
 - CONVEX FORTRAN User's Guide xx
 - CONVEX Performance Analyzer (CXpa) User's Guide xxi
 - ConvexOS Man Pages for Programmers xx
 - ConvexOS Man Pages for Users xx
 - ConvexOS Primer xx
 - CXmetrics User's Guide xxi
 - Fortran 90 standard xxi
 - ISO/IEC 1539:1991 xxi
- ASSOCIATEVARIABLE keyword 122
- assumed-length arrays
 - character 193
- assumed-length character argument 45
- assumed-length function name 200

assumed-size arrays 190 to 191
 using * to dimension 190 to 191
asterisk character (*)
 and assumed-size arrays 190 to 191
 and ENTRY statement 203
 as descriptor 155
 in dummy argument list 194
 runtime format descriptor 177
ATAN intrinsic 209
 result 219
ATAN2 intrinsic 209
 result 219
ATAN2D intrinsic 209
 argument 219
 result 219
ATAND intrinsic 209
 argument 219
 result 219
automatic arrays 22, 25, 231
 Cray 256
 dimensioning 23, 49
 form 231
 multidimensional 23, 49
AUTOMATIC statement
 Sun FORTRAN 271
auxiliary input/output statements 119

B

B edit descriptor 166
backslash character (\)
 SUN FORTRAN 271
BACKSPACE statement 136
binary data file format conversions 137
 and Cray data types 140
 and Cray floating point data 141
 and Cray unformatted files 141
 and error handling 141
 and Hollerith data 140
 and optimization 141
 and REAL*16 data 140
 and SIGFPE signal 141
 and VAX data types 140
 restrictions 140
 sample user-defined routine 143
 user-defined 142
 user-supplied routine names 144
 using a shell variable 146
 VAX data 139
 via the OPEN statement 139
 when to use 138
binary data files
 VAX 264
bits
 circular shift 217
 extraction 221
 extraction intrinsic 216
 numbering 221
 setting 216

 testing 216
bitwise AND 216
bitwise circular shift 217
bitwise complement 216
bitwise OR 216
bitwise shift 216
bitwise XOR 216
blank characters
 in numeric formatted input fields 123
blank common storage 39
BLANK keyword 123
 FORTRAN 66 227
 FORTRAN 66 interpretation 227
BLOCK DATA statement 187
 form 187
 in function subprograms 200
block data subprogram 187
 and COMMON blocks 188
block IF statement 80
 nested 83
blocks
 COMMON 38, 187
blocksize
 specifying 123
BLOCKSIZE keyword 123
BN edit descriptor 166
BTEST intrinsic 216
BUFFERIN statement 256
BUFFEROUT statement 256
built-in functions 196
 %LOC 197
 %REF 196
 %VAL 196
bypassing input records 175
BYTE data type 9 to 10
BZ edit descriptor 166

C

C functions
 calling from FORTRAN 48
 f_init 111
CALL statement 89, 201
 form 201
calling FORTRAN from C 48
 initializing FORTRAN I/O 111
carriage control
 characters 186
 in formatted sequential WRITES 173
carriage returns
 suppressing in formatted sequential WRITES 173
CARRIAGECONTROL keyword 124
CDLOG intrinsic
 argument 219
-cfc option 22, 155, 177, 231, 253
 and allocatable arrays 26, 50, 232
 and compiler defaults 253
 and constants 253
 and default data types 253

- and intrinsics 253
- Boolean octal constants 259
- Hollerith constants 260
- intrinsics 258
- library routines 258
- libU77.a 258
- TASK COMMON 259
- cfcwpa option 254
- CHAR intrinsic 217
- character arguments 192
 - assumed-length 192
 - lengths 192
- character arrays
 - as dummy arguments 193
- character assignments 66
- character concatenation 35
- character constants 19
 - formatting 155
- character conversions 66
- character data
 - formatting 153
- CHARACTER data type 9 to 10
 - as arguments to procedures 192
 - as dummy arguments 193
 - type-declaration statements 44
- character descriptor A 153
- character equivalence 53
- character expressions 35
 - finding length of 217
- CHARACTER FUNCTION statement 200
- character set 1 to 2
 - extended 2
 - standard 1
- character substrings 28
 - finding position in character expression 217
 - referencing 28
- CHARACTER variables
 - declaring 44
- character-per-column formatting 4
- characters
 - carriage-control 186
 - finding ASCII values of 217
 - finding from ASCII value 217
- CLOG intrinsic
 - argument 219
- CLOSE statement 131
 - form 131
- closing files 131
- CMPLX intrinsic 213
 - use with one argument 220
- colon edit descriptor (:) 174
- comma field separator 176
- comment indicators 2
- comment line 2
- comments 2
- COMMON
 - Cray TASK 39
- common block storage 39
- COMMON blocks 38
 - Cray TASK COMMON 259
- equivalencing 53
 - in block data subprogram 188
- initializing arrays 39
- initializing variables 39
- initializing via block data subprograms 187
- initializing via DATA statements 63
- COMMON statement 19, 38
- commonly used library routines 222
- compatibility
 - Fortran 90 229
- compatibility modes
 - Cray FORTRAN 253, 258 to 260
 - FORTRAN 66 225
 - Sun FORTRAN 271
 - VAX FORTRAN 261, 263 to 267, 269
- compiler directives 6
- compiler options
 - 72 6, 18
 - cfc 22, 26, 155, 177, 231 to 232, 253, 258
 - dfc 139
 - F66 85, 225 to 226
 - f90 22, 25 to 27, 67, 229, 231 to 232, 247 to 248
 - li66 225
 - nosc 84
 - re 205
 - sfc 271
 - vfc 261
- compiling FORTRAN 66 programs 225
- complement
 - bitwise 216
- complex constants 14
- complex data
 - formatting 162, 164
 - scaling with FORMAT 168
- COMPLEX data type 10
- complex descriptor 153
- complex number
 - absolute value of 220
- COMPLEX*16 constants 14
- COMPLEX*16 data type 9, 32
- COMPLEX*8 constants 14
- COMPLEX*8 data type 9 to 10, 32 to 33
- COMPLEX-to-COMPLEX conversions 12
- COMPLEX-to-noncomplex conversions 12
- computed GOTO statement 76
- concatenation 35
- conditional evaluation
 - short circuiting 83
- conformability
 - defined 21
 - scalar value 22
- CONJG intrinsic 214
- constant values 41
- constants 9, 13
 - and -cfc 253
 - character 19
 - complex 14
 - Cray Hollerith 17, 260
 - Cray octal 15, 259
 - double-precision 14

- hexadecimal 16
- Hollerith 17
- integer 13
- logical 18
- octal 15
- real 13
- VAX octal 15
- contact utility xdi
- continuation indicator 5
- continuation line 5
 - in Sun FORTRAN 271
- CONTINUE statement 89
- control statement 75
- conversion
 - automatic 12
 - binary data file format 137
 - character 66
 - data type 65, 70
 - restrictions 140
 - user-defined 142
 - using a shell variable 146
- conversion feature
 - when to use 138
- conversion of data types 12
- conversion using OPEN statement 139
- CONVEX extensions typeface xix
- COS intrinsic 208
 - argument 219
- COSD intrinsic 208
 - argument 219
- COSH intrinsic 209
- COUNT intrinsic 217, 238
- COVUshell 8
 - OPEN statement operation 120
- Cray compatibility 253
 - and * descriptor 155, 177
 - and -cfcwpa 254
 - automatic arrays 256
 - floating point conversions 141
 - Hollerith constants 17
 - intrinsic 258
 - library routines 258
 - libU77.a 258
 - octal constants 15, 259
 - pointer arithmetic 254
 - POINTER statement 254
 - TASK COMMON 259
 - unformatted file support 257
 - unsupported features 254
- Cray functions
 - GETPOS 256
 - LENGTH 256
 - SETPOS 256
 - UNIT 256
- CSHIFT intrinsic 218, 243
- cvbin utility 138, 264

D

- D edit descriptor 157, 162
 - and comma field separator 176
 - default value 175
 - with B descriptor 166
 - with S descriptor 169
- D indicator 5
- DACOS intrinsic
 - argument 219
 - result 219
- DACOSD intrinsic
 - argument 219
 - result 219
- DASIN intrinsic
 - argument 219
 - result 219
- DASIND intrinsic
 - argument 219
 - result 219
- data format 96 to 97, 147
- DATA statement 59
 - and COMMON variables 63
 - data-type conversion 62
 - extensions 62
 - form 59
 - in block data subprogram 188
- data type conversion 220
 - array assignments 67, 247
- data type priority 32
- data types 9
 - arithmetic expression 31
 - BYTE 9
 - CHARACTER 9 to 10
 - COMPLEX*16 9
 - COMPLEX*8 9
 - controlling storage requirements 10
 - conversion 12, 70
 - converting 65
 - correspondence to ANSI standard 10
 - defaults under -cfc 253
 - DOUBLE COMPLEX 9
 - DOUBLE PRECISION 9
 - equivalenced 51
 - for hexadecimal constants 16
 - for octal constants 16
 - INTEGER*1 9
 - INTEGER*2 9
 - INTEGER*4 9
 - INTEGER*8 9
 - LOCIGAL*1 9
 - LOGICAL*2 9
 - LOGICAL*4 9
 - LOGICAL*8 9
 - of arithmetic expressions 33
 - of Hollerith constants 18
 - REAL*16 9
 - REAL*4 9
 - REAL*8 9
 - RECORD 9

standard 9
 data-type length specifiers 43
 dataformat 127
 DATAN intrinsic
 result 219
 DATAN2 intrinsic
 result 219
 DATAN2D intrinsic
 result 219
 DATAND intrinsic
 result 219
 date subroutine 222
 DBLE intrinsic 212
 argument 220
 -dc option 5
 DCMPLX intrinsic 213
 use with one argument 220
 DCOS intrinsic
 argument 219
 DEALLOCATE statement 26, 233
 form 26, 234
 debug statements 5
 debugging Cray pointers 255
 declaration statements
 type 43
 declaring allocatable arrays 25, 232
 declaring arrays 22, 49
 DECODE statement 117
 form 117
 see also sequential-access internal READ statement
 default edit descriptor values 175
 default files
 specifying 124
 DEFAULTFILE keyword 124
 delimiters
 string 19
 descriptors
 see edit descriptors
 -dfc option 139
 DFLOAT intrinsic 213
 DIM intrinsic 215
 dimension declarator 22
 DIMENSION statement 22, 49
 form 22
 dimensioning arrays 22, 49
 direct access 98
 direct-access files
 maximum number of records 127
 positioning 118
 specifying 122
 direct-access internal READ statement 108
 direct-access internal WRITE statement 114
 directives
 compiler 6
 disconnecting files 131
 DISPOSE keyword 125
 DLOG intrinsic
 argument 219
 DLOG10 intrinsic
 argument 219
 DO list
 implied 99
 implied 61
 DO loop 84
 control variable 84
 extended range 86
 FORTRAN 66 behavior 85
 FORTRAN 66 minimum iteration count 226
 increment 85
 iteration count 85
 nested 86
 shared terminal statements 86
 terminating 88
 DO statement 84
 extended range 87
 DO WHILE loop
 terminating 88
 DO WHILE statement 87
 DO-variable 84
 dollar sign (\$)
 edit descriptor 173
 record delimiter 181
 DOT_PRODUCT intrinsic 217, 235
 DOUBLE COMPLEX data type 32
 DOUBLE PRECISION data type 9 to 10, 32
 double-precision constants 13
 double-precision data
 formatting 162, 164
 DSIN intrinsic
 argument 219
 DTAN intrinsic
 argument 219
 dummy arguments 187 to 188, 202
 adjustable arrays 190
 and ENTRY statement 203
 arrays as 190, 192
 CHARACTER 192
 in statement functions 199
 NAMELIST statement 55
 procedures as 194
 to functions 199 to 200
 variables as 189
 dummy procedures 194
 dynamic arrays 47, 50
 dynamic memory allocation 47

E
 E edit descriptor 162
 and comma field separator 176
 default value 175
 with B descriptor 166
 with S descriptor 169
 edit descriptors 153
 \$(dollar sign) 173
 *(asterisk) 155, 177
 /(slash) 175
 :(colon) 174
 A 153

apostrophe 155
 B 166
 Cray compatibility 155
 D 157, 162, 166, 169
 default values 175
 E 162, 166, 169
 F 160, 166, 169
 for character constant 155
 for character data 153
 for complex data 164
 for double-precision data 162, 164
 for Hollerith data 153
 for integer data 157
 for literal text 155 to 156
 for logical data 156
 for numeric data 166
 for octal data 158
 for positioning 171
 for quad-precision data 164
 for real data 160, 162
 for single-precision data 164
 for unsigned hexadecimal data 159
 forcing sign 169
 FORTRAN 66 225, 228
 G 164, 166
 H 156
 I 166, 169 to 170
 L 156
 nonrepeatable 151
 O 158, 166
 overriding specified field width 176
 P 168
 Q 174
 R 170
 repeatable 151 to 152
 S 169
 scale factor 168
 suppressing newline 173
 T 171
 X 171
 Z 159, 166
 ELSE IF statement 80
 ELSE statement 80
 ELSEWHERE statement 68, 248
 ENCODE statement 115
 form 115
 see also sequential-access internal WRITE statement
 END DO statement 88
 END IF statement 80
 END MAP statement 267
 END specifier 103
 END statement 91
 in function subprograms 200
 with SUBROUTINE statement 202
 END STRUCTURE statement 265
 END UNION statement 266
 end-of-file specifier 103
 ENDFILE record 95
 ENDFILE statement 95, 137
 entry points
 alternative 202
 ENTRY statement 202
 and * 203
 and dummy arguments 203
 form 203
 in function subprogram 201
 EOSHIFT intrinsic 218, 244
 EQUIVALENCE statement 19, 51, 53
 and data file conversions 140
 equivalencing
 arrays 51
 common blocks 53
 substrings 53
 .EQV. operator 34
 ERR keyword 126
 ERR specifiers 102
 error specifier 102
 errsns subroutine 222
 escape character
 in Sun FORTRAN 271
 evaluation
 short circuit 83
 exclusive OR 216
 executable program 1
 executable statements 3
 externally-defined procedures
 identifying 55
 exit subroutine 222
 EXP intrinsic 208
 expressions 31
 arithmetic 31
 character 35
 logical 34
 relational 33
 extended-range DO loop 86
 external files
 accessing 96
 connecting to units 96
 manipulating 135
 opening 96
 external READ statement 105
 formatted 105
 list-directed sequential-access 106
 namelist-directed 106
 namelist-directed sequential-access 112
 unformatted 105
 EXTERNAL statement 55, 202
 FORTRAN 66 225 to 226
 external WRITE statement
 list-directed sequential-access 112

F

F edit descriptor 160
 and comma field separator 176
 default value 175
 with B descriptor 166
 with S descriptor 169
 -F66 option 225 to 226

- and DO loops 85
- f90 option 22, 27, 67, 229, 231
 - and allocatable arrays 25 to 26, 50, 232
 - and array-valued expressions 247
 - and masked array assignments 248
- f_init function 111
- field declaration 266
- field descriptors 153
- field separators
 - external 176
- fields
 - in VAX records 266
 - statement 3
- file access
 - direct 98
 - sequential 97
- file access keyword 122
- file format
 - converting from VAX 138
- FILE keyword 126
- file positioning 152
- file-positioning statements 135
 - form 135
- files 95
 - accessing 97
 - binary conversions 137
 - closing 131
 - Cray unformatted 257
 - determining access mode 132
 - determining attributes 132
 - determining block size 132
 - disconnecting from units 131
 - formatted 101
 - internal 95
 - large 95
 - manipulating external 135
 - opening 96
 - opening for I/O 119
 - positioning at initial point 136
 - positioning to end 137
 - positioning to preceding record 136
 - sequential access 97
 - shared 129
 - specifying name for unit 126
 - specifying read-only 128
 - specifying status 129
- %FILL
 - in VAX records 266
- FIND statement 118
 - form 118
- fixed-length I/O records
 - specifying 129
- FLOAT intrinsic 213
- floating point data
 - VAX 267
- FMT specifier 101
- form feed character
 - SUN FORTRAN 271
- FORM keyword 126, 139
- format code separators 228
- FORMAT control 151
 - ending with : descriptor 174
- format conversions
 - binary data file 137
- format specifications 149, 177
- format specifier 101
- FORMAT statement 101, 149
 - \$ (dollar sign) descriptor 173
 - * (asterisk) descriptor 155, 177
 - / (slash) descriptor 174
 - : (colon) descriptor 174
 - A descriptor 153
 - and signed data 169
 - apostrophe descriptor 155
 - B descriptor 166
 - changing radix for integer I/O 170
 - comma field separator 176
 - Cray compatibility 155
 - D descriptor 162
 - default descriptor values 175
 - E descriptor 162
 - edit descriptors 153
 - F descriptor 160
 - form 149
 - G descriptor 164
 - H descriptor 156
 - I descriptor 157, 170
 - L descriptor 156
 - labeled 149
 - O descriptor 158
 - overriding input field width 176
 - P descriptor 168
 - positioning descriptor 171
 - Q descriptor 174
 - R descriptor 170
 - repeatable edit descriptors 152
 - S descriptor 169
 - setting record position 171
 - suppressing newline 173
 - T descriptor 171
 - variable 177
 - X descriptor 171
 - Z descriptor 159
- formats
 - variable 177
- formatted files
 - specifying 101
- formatted I/O
 - direct access WRITE 112
 - external direct access READ 107
 - internal sequential-access READ 108
 - sequential READ statement 105
 - sequential-access WRITE statement 111
 - specifying 126
- formatted I/O records 94
- formatted I/O statements 149
- formatted input 94
- formatted sequential-access WRITE statement 111
 - suppressing newlines 173
- formatting

ANSI standard 6
 carriage-control characters 186
 character-per-column 4
see also I/O formatting
 list-directed 179, 184
 NAMELIST statement 181
 namelist-directed 181, 184
 tab-key 6
 vertical 186

FORTRAN 66
 BLANK keyword 225, 227
 compatibility 225
 compiling programs 225
 differences 225
 DO loop 225 to 226
 DO loop behavior 85
 EXTERNAL statement 225 to 226
 STATUS keyword 225
 X edit descriptor 228
 X format edit descriptor 225

FORTRAN 77
 ANSI standard xvii
 FORTRAN 77 formatting 6

Fortran 90
 ANSI standard 235
 allocatable arrays 26, 50, 232
 array sections 27, 229
 automatic arrays 231

Fortran 90 array assignments 66, 247
 form 247
 masked 248

Fortran 90 arrays
 masked assignment 67, 248

Fortran 90 compatibility 229
 allocatable arrays 232
 array assignments 247
 array sections 229
 automatic arrays 231
 intrinsics 235
 masked array assignments 248

Fortran 90 intrinsics 222, 235
 ALL 217, 237
 ANY 217, 237
 array construction 241
 array location functions 245
 array manipulation 241
 array manipulation functions 243
 COUNT 217, 238
 CSHIFT 218, 243
 DOT_PRODUCT 217, 235
 EOSHIFT 218, 244
 MATMUL 217, 236
 MAXLOC 218, 246
 MAXVAL 218, 238
 MERGE 218, 241
 MINLOC 218, 246
 MINVAL 218, 239
 optional arguments 235
 PACK 218, 242
 PRODUCT 218, 239
 reduction functions 236
 SPREAD 218, 242
 SUM 218, 240
 TRANSPOSE 218, 245
 UNPACK 218, 243
 vector and matrix multiply 235

Fortran 90 standard 229
FORTTRAN preprocessor 8
 fpp (FORTRAN preprocessor) 8
 fpr utility 186

FUNCTION statement 199
 CHARACTER form 200
 data type 200
 dummy arguments 199 to 200
 form 199
 function name 200

function subprogram 199
 assumed length name 200
 data type 200
 dummy arguments 200
 names 200

functions 195
 built-in 196
 intrinsic 195
 %LOC 197
 %REF 196
 statement 197
 %VAL 196

G

G edit descriptor 164
 and comma field separator 176
 default value 175
 with B descriptor 166

generic intrinsics 195, 207

GETPOS function 256

GOTO
 assigned 77
 computed 76
 unconditional 75

GOTO statement 75

H

H edit descriptor 156
 VAX interpretation 264

hexadecimal constants 16
 data types 16

hexadecimal data
 formatting 159

Hollerith constants 17
 as actual arguments 192
 as arguments 192
 continuing across lines 18
 Cray 17, 260

data types 18
VAX 264
hyperbolic cosine 209
hyperbolic sine 209
hyperbolic tangent 209

I

I edit descriptor 157, 170
and comma field separator 176
default value 175
with B descriptor 166
with S descriptor 169

I/O

large files 95
unformatted Cray 257
I/O formatting 149, 151
\$ (dollar sign) delimiter 181
\$ (dollar sign) descriptor 173
* (asterisk) descriptor 177
* (asterisk) format indicator 179
/ (slash) descriptor 175
: (colon) descriptor 174
bypassing input records 175
carriage-control characters 186
changing radix 170
character constants 155
character data 153
comma field separator 176
complex data 162, 164
creating empty records 175
data-type based 179
default descriptor values 175
descriptors 153
double-precision data 162, 164
forcing sign 169
hexadecimal data 159
in character variable 149, 177
integer data 157
list-directed 149, 179
logical data 156
NAMELIST statement 181
namelist-directed 181, 184
numeric data 166
octal data 158
Q descriptor 174
quad-precision data 164
real data 160, 162
repeat count 152
runtime 177
scale factor 168
setting record position 171
single-precision data 164
suppressing newlines in 173
variable 177
vertical 186
I/O list specifiers 100
I/O records 94
\$ (dollar sign) delimiter 181

bypassing on input 175
creating empty 175
determining number of unread characters 174
ENDFILE 95
formatted 94
in internal files 96
list-directed 179
setting position in 171
specifying fixed-length 129
specifying maximum size 128
specifying size 128
specifying variable-length 129
unformatted 95
variable format 177
I/O statements 93
ACCEPT 109
auxiliary 119
DECODE 117
ENCODE 115
FIND 118
format 98
formatting 149, 177
OPEN 119
PRINT 114
READ 104
special 115
suppressing newlines in 173
TYPE 114
variable format 177
WRITE 110
-i2 option
and intrinsic functions 220
-i4 option
and intrinsic functions 220
IABS intrinsic 210
IAND intrinsic 216
IBITS intrinsic 216
IBSET intrinsic 216
ICHAR intrinsic 217
idate subroutine 222
IDIM intrinsic 215
IDINT intrinsic
result 220
IDNINT intrinsic 211
result 220
IEEE mode 137
IEOR intrinsic 216
IF statement 78
arithmetic 78
block 80
nested block 83
short circuiting 83
IF THEN statement 80
IFIX intrinsic 213
result 220
IMPLICIT NONE statement 40
IMPLICIT statement 39
and intrinsic functions 196
implicit typing 12
defined 40

- disabling 40
- overriding 39, 43
- implied-DO list 61, 99
- INCLUDE statement 8
- inclusive OR 216
- INDEX intrinsic 217
- Inf operand 162, 164, 166
- initial line 5
- initializing
 - arrays in type-declarations 44
 - character variables in type-declarations 44
 - variables in type-declarations 43
- input
 - formatted 94
 - input fields
 - \$ (dollar sign) delimiter 181
 - delimiters 180
 - input files
 - field delimiters 180
 - input formatting
 - comma field separator 176
 - list-directed 179
 - NAMELIST statement 181
 - namelist-directed 181
 - input records
 - comma field separator 176
 - null fields 177
 - input/output
 - and program mode 137
 - list-directed formatting 179
 - input/output lists 99
 - input/output statements 93
 - INQUIRE statement 132
 - and large files 133
 - form 132
 - INT intrinsic 210
 - argument 220
 - result 220
 - integer
 - finding nearest 211
 - integer constants 13
 - integer data
 - formatting 157
 - integer descriptor 153
 - INTEGER variables
 - declaring 43
 - INTEGER*1 data type 9, 32
 - INTEGER*2 data type 9, 32
 - INTEGER*4 data type 9, 32
 - and -cfc 253
 - INTEGER*8 data type 9, 32
 - INTEGER-to-REAL conversions 12
 - internal files 95
 - I/O records in 96
 - internal READ statement 108
 - direct-access 108
 - sequential-access 108
 - internal WRITE statement
 - direct-access 114
 - sequential-access 113
 - internal WRITE statements 113
 - intrinsic functions 195, 207
 - and IMPLICIT statement 196
 - Cray 258
 - generic names 195
 - specific names 195
 - using as actual arguments 56
 - INTRINSIC statement 56, 226
 - intrinsics
 - ABS 210
 - ACOS 209
 - ACOSD 209
 - ALL 217
 - AMAX0 214
 - AMIN0 215
 - ANINT 211
 - ANY 217
 - ASIN 208
 - ASIND 208
 - ATAN 209
 - ATAN2 209
 - ATAN2D 209
 - ATAND 209
 - bit extraction 221
 - BTEST 216
 - CHAR 217
 - CMPLX 213
 - CONJG 214
 - COS 208
 - COSD 208
 - COSH 209
 - COUNT 217
 - CSHIFT 218
 - DBLE 212
 - DCMPLX 213
 - DFLOAT 213
 - DIM 215
 - DOT_PRODUCT 217
 - EOSHIFT 218
 - EXP 208
 - FLOAT 213
 - Fortran 90 222, 235
 - generic 207
 - IABS 210
 - IAND 216
 - IBITS 216
 - IBSET 216
 - ICHAR 217
 - IDIM 215
 - IDNINT 211
 - IEOR 216
 - IFIX 213
 - INDEX 217
 - INT 210
 - IOR 216
 - IQINT 211
 - IQNINT 211
 - ISHFT 216
 - ISHFTC 217
 - ISIGN 216

LEN 217
LOG 207
LOG10 207
MATMUL 217
matrix multiply 235
MAX 214
MAX0 214
MAX1 214
MAXLOC 218
MAXVAL 218
MERGE 218
MIN 215
MIN0 215
MIN1 215
MINLOC 218
MINVAL 218
MOD 215
NINT 211
NOT 216
PACK 218
PRODUCT 218
QEXT 213
QFLOAT 213
REAL 212
shift 221
SIGN 216
SIN 208
SIND 208
SINH 209
specific 207
SPREAD 218
SQRT 207, 219
SUM 218
TAN 208
TAND 208
TANH 209
TRANSPOSE 218
truncation 220
under -cfc 253
UNPACK 218
VAX 269
vector multiply 235
ZEXT 212
IOR intrinsic 216
IOSTAT keyword 127
IOSTAT specifier 102
IQINT intrinsic 211
result 220
IQNINT intrinsic 211
ISHFT intrinsic 216
ISHFTC intrinsic 217
ISIGN intrinsic 216
iteration count in DO loop 85

K

keywords

ACCESS 97, 122
ASSOCIATEVARIABLE 122

BLANK 123, 225
BLOCKSIZE 123
CARRIAGECONTROL 124
DEFAULTFILE 124
DISPOSE 125
ERR 126
FILE 126
FORM 126, 139
FORTRAN 66 OPEN statement 227
FORTRAN 66 STATUS 228
IOSTAT 127
MAXREC 127
NAME 126
NOSPANBLOCKS 128
OPEN statement 120
READONLY 128
RECL 128
RECORDSIZE 128
RECORDTYPE 129
SHARED 129
STATUS 129, 225
TYPE 129
UNIT 130

L

L edit descriptor 156
and comma field separator 176
default value 175
labels
statement 4
large files 95
and INQUIRE 133
and REC specifier 102
LEN intrinsic 217
LENGTH function 256
length of character expression 217
length specifiers
see data-type length specifiers
-li66 option 225
libcfc.a 256
library routines 207, 222
Cray 258
date 222
errsns 222
exit 222
idate 222
mvbits 223
ran 223
secsds 222
time 222
libU77.a
Cray support 258
limits
file size 95
list-directed formatting 149, 179
character input 180
complex input 180
input 179

- null value 180
 - output 184
 - slashes 180
- list-directed input 179
- list-directed output 184
- list-directed records 179
 - separators 179
- list-directed sequential-access READ statement 106
- list-directed sequential-access WRITE statement 112
- lists
 - implied-DO 61, 99
 - input/output 99
- LOC function 46, 197
 - and -cfc 254
 - example 46
- location functions 245
- location of storage elements
 - finding 197
- LOG intrinsic 207
- LOG10 intrinsic 207
- logical constants 18
- logical data
 - formatting 156
- logical descriptor 153
- logical elements 34
- logical entities
 - in arithmetic context 33
- logical expressions 34
- logical IF statement 79
- logical operator .XOR. 35
- logical operators 34
 - .AND. 34
 - .EQV. 34
 - .NEQV. 34
 - .NOT. 34
 - .OR. 34
 - .XOR. 34
- LOGICAL variables
 - declaring 43
- LOGICAL*1 data type 9, 32
- LOGICAL*2 data type 9, 32
 - and -cfc 253
- LOGICAL*4 data type 9, 32
 - and -cfc 253
- LOGICAL*8 data type 9, 32
- loop
 - DO 84
 - DO WHILE 87
 - nested DO 86

M

- main program 1, 38
- malloc
 - calling from FORTRAN 47
- manipulation functions 243
- manual organization xvii
- map declarations
 - VAX 266

- MAP statement 267
 - form 267
- masked array assignment 67, 248
 - form 248
- MATMUL intrinsic 217, 236
- matrix
 - defined 235
- matrix multiplication 236
- matrix multiply functions 235
- MAX intrinsic 214
- MAX0 intrinsic 214
- MAX1 intrinsic 214
 - result 220
- maximum
 - finding 214
 - locating array element 246
- MAXLOC intrinsic 218, 246
- MAXREC keyword 127
- MAXVAL intrinsic 218, 238
- MERGE intrinsic 218, 241
- MIN intrinsic 215
- MIN0 intrinsic 215
- MIN1 intrinsic 215
- MINI intrinsic
 - result 220
- minimum
 - finding 215
 - locating array element 246
- MINLOC intrinsic 218, 246
- MINVAL intrinsic 218, 239
- MOD intrinsic 215
- mode
 - floating point 137
 - IEEE 137
 - NATIVE 137
 - program 137
- multidimensional dynamic arrays 48
- multiple statements 5
- mvbits subroutine 223

N

- NAME keyword 126
- namelist specifier 103
- NAMELIST statement 54, 181
 - and arrays 182
 - namelist-directed formatting 181, 184
 - and arrays 182
 - form 181
 - NAMELIST statement 181
- namelist-directed output 184
- namelist-directed READ statement 106
- namelist-directed sequential-access READ statement 112
- NaN operand 162, 164, 166
- NATIVE mode 137
- nearest integer intrinsic (NINT) 211
- .NEQV. operator 34
- nested block IF statement 83
- nested DO loops 86

- newline character
 - Sun FORTRAN 271
- newline suppression in formatted WRITE 173
- NEXTREC specifier
 - and large files 133
- NINT intrinsic 211
 - result 220
- NML specifier 103
- noncomplex-to-COMPLEX conversions 12
- nonexecutable statements 3
- nonrepeatable edit descriptors 151
- nosc option 84
- NOSPANBLOCKS keyword 128
- NOT intrinsic 216
- .NOT. operator 34
- notational conventions xviii
- Notes
 - explanation xix
- NUL character
 - SUN FORTRAN 271
- null values 180
- numbering of bits 221
- numeric data
 - formatting 166
- numeric type-declaration statements 43

O

- O edit descriptor 158
 - and comma field separator 176
 - default value 175
 - with B descriptor 166
- octal constants 15
 - Cray 259
 - data types 16
 - VAX 264
- octal data
 - formatting 158
- OPEN statement 96, 119
 - data conversions 139
 - errors 126
 - form 119
 - FORTRAN 66 keywords 227
 - keywords 120
 - user_defined data format 142
- opening files for I/O 119
- operands
 - Inf 162, 164, 166
 - NaN 162, 164, 166
 - reserved 162, 164, 166
 - Rop 162, 164, 166
- operator precedence 32
- operators
 - arithmetic 31
 - logical 34
 - relational 33
- optional arguments
 - to Fortran 90 intrinsics 235
- OPTIONS statement 7

- OR
 - bitwise intrinsic 216
 - inclusive intrinsic 216
- .OR. operator 34
- order of statments and lines 6
- organization of manual xvii
- output formatting
 - list-directed 184
 - namelist-directed 184

P

- P edit descriptor 168
- p8 option
 - and intrinsic functions 220
- PACK intrinsic 218, 242
- PARAMETER statement 41
 - alternate 42
 - standard 41
- parameters 41
- PAUSE statement 90
- pd8 option
 - and intrinsic functions 220
- pointees 45
- pointer arithmetic
 - and -cfc 254
- POINTER statement 45
 - and -cfc 254
 - and -cfcwpa 254
- pointers 45
 - and dynamic arrays 47
 - and dynamic memory allocation 47
 - and pointees 45
 - debugging 255
 - declaring 45
 - dimensioning dynamic arrays 48
 - example 46 to 47
- positioning direct-access files 118
- positive difference intrinsic 215
- precedence
 - operator 32
- precision
 - of operands in expressions 32
- preprocessor 8
- PRINT statement 94, 114
- printing
 - and CARRIAGECONTROL keyword 124
- priority
 - data type 32
- procedures 188
 - actual arguments 188
 - adjustable arrays as dummy arguments 190
 - arrays as dummy arguments to 190, 192
 - as dummy arguments 194
 - assumed-size arrays as dummy argument 190
 - character arguments 192
 - dummy 194
 - dummy arguments 188
 - function subprograms 199

- functions 195
- intrinsic functions 195
- single-statement 197
- statement functions 197
- subroutine 201
- variables as dummy arguments 189
- PRODUCT intrinsic 218, 239
- program
 - executable 1
 - main 1, 38
 - types of 1
- program control
 - returning from function 91
 - returning from subroutine 89, 91
 - transferring via GOTO statement 75
 - transferring via IF statement 78
- program execution
 - suspending 90
 - terminating 90
- program mode 137
- PROGRAM statement 1, 38
 - in function subprograms 200
- program units 1

- argument 219
- quad-precision data
 - formatting 164

R

- R edit descriptor 170
- 'r' specifier 101
- radix
 - changing for integer I/O 170
- ran function 223
- random number generator 223
- rank 21
 - allocatable arrays 232
 - data type 32
- rank definitions
 - form 25, 232
- re option 205
- READ statement 94, 104
 - and large files 102
 - bypassing records 175
 - direct-access external 107
 - direct-access internal 108
 - external 105
 - external direct-access 107
 - form 104
 - formatted external direct access 107
 - formatted sequential 105
 - internal 108
 - internal direct access 108
 - internal sequential access 108
 - list-directed formatting 179
 - list-directed sequential-access 106
 - namelist-directed 106
 - namelist-directed sequential-access 112
 - sequential-access external 105
 - sequential-access internal 108
 - unformatted external direct access 107
 - unformatted sequential-access 105
 - VAX interpretation 264
- READ statement formatting
 - see I/O formatting
- read-only files
 - specifying 128
- READONLY keyword 128
- real constants 13
 - Sun support 271
- real data
 - formatting 160, 162
 - scaling with FORMAT 168
- REAL data type 10
- real descriptor 153
- REAL intrinsic 212
 - argument 220
- REAL variables
 - declaring 43
- REAL*16 constants
 - VAX interpretation 264
- REAL*16 data type 9, 13, 32

Q

- Q edit descriptor 174
- QACOS intrinsic
 - argument 219
 - result 219
- QACOSD intrinsic
 - argument 219
 - result 219
- QASIN intrinsic
 - argument 219
 - result 219
- QASIND intrinsic
 - argument 219
 - result 219
- QATAN intrinsic
 - result 219
- QATAN2 intrinsic
 - result 219
- QATAN2D intrinsic
 - result 219
- QATAND intrinsic
 - result 219
- QCOS intrinsic
 - argument 219
- QEXT intrinsic 213
- QFLOAT intrinsic 213
- QLOG intrinsic
 - argument 219
- QLOG10 intrinsic
 - argument 219
- QSIN intrinsic
 - argument 219
- QTAN intrinsic

- accuracy 254
 - and -cfc 253 to 254
 - and IEEE mode 137
- REAL*4 data type 9 to 10, 13, 32, 70
 - and -cfc 253
- REAL*8 data type 9 to 10, 13, 32 to 33, 70
- REAL*8-to-REAL*4 conversions 12
- REAL-to-INTEGER conversions 12
- REAL-to-REAL conversions 12
- REC specifier 101
- RECL keyword 128
- RECL specifier
 - VAX form 261
- record data structure
 - VAX 265
- RECORD data type 9
 - declaring 45
- record specifier 101
 - VAX form 102
- RECORD statement
 - field declaration 266
 - form 265
 - in substructure 266
 - VAX 265
- RECORD type-declaration statements 45
- records
 - and large files 102
 - ENDFILE 95
 - formatted I/O 94
 - I/O 94
 - in internal files 96
 - unformatted I/O 95
- RECORDSIZE keyword 128
- RECORDTYPE keyword 129
- recursion
 - in Sun FORTRAN 271
- %REF function 196
- reference
 - passing arguments by 196
- referencing array elements 27
- relational expressions 33
- relational operators 33
- remainder intrinsic 215
- repeatable edit descriptors 151 to 152
- reserved operand 162, 164, 166
- restrictions on conversions 140
- RETURN
 - alternate 201
- RETURN statement 89, 91, 201 to 202, 204
 - form 204
 - in function subprograms 200
- REWIND statement 136
- Rop operand 162, 164, 166
- runtime formats 177
 - and * descriptor 177

S

- S edit descriptor 169
- SAVE statement 56
 - and automatic arrays 231
- saving subprogram variables after return 56
- scalar value conformability 22
- scale factor
 - complex data 168
 - real data 168
- secnds function 222
- semicolon separator 5
- SEPOS function 256
- sequential-access external READ statement 105
- sequential-access files 97
 - specifying 122
- sequential-access internal READ statement 108
 - see also* DECODE statement
- sequential-access internal WRITE statement 113
 - see also* ENCODE statement
- sequential-access WRITE statements 111
- 72 option 6, 18
- sfc option 271
 - and recursion 271
 - AUTOMATIC statement 271
 - escape sequences 271
 - STATIC statement 271
- shape 21
- shared files 129
- SHARED keyword 129
- shell variable conversions 146
- shell variables 96
- shift intrinsics 221
- shifting array elements 243 to 244
- short-circuit evaluation 83
 - disabling 84
 - inhibitors 84
- SIGFPE signal
 - and binary data file format conversions 141
- SIGN intrinsic 216
- signed data
 - formatting output 169
- SIN intrinsic 208
 - argument 219
- SIND intrinsic 208
 - argument 219
- single-precision data
 - formatting 164
- SINH intrinsic 209
- slash (/) edit descriptor 175
 - and direct-access files 175
 - and sequential-access files 175
- SP edit descriptor 169
- special I/O statements 115
- specific intrinsics 195, 207
- specification statements 37
- specifiers
 - control information list 100
 - END 103
 - end-of-file 103

ERR 102
 error 102
 FMT 101
 format 101
 IOSTAT 102
 namelist 103
 NML 103
 'r (record specifier) 101
 REC 101
 record 101
 status 102
 unit 100
 SPREAD intrinsic 218, 242
 SQRT intrinsic 207
 argument 219
 result 219
 SS edit descriptor 169
 standard PARAMETER statement 41
 statement field 5
 statement function 197
 form 198
 reference 199
 statement label 4
 assigning 73
 field 4
 statements
 ACCEPT 109
 ALLOCATABLE 25, 50, 233
 ALLOCATE 26, 233
 arithmetic IF 78
 assigned GOTO 77
 assignment 65
 auxiliary I/O 119
 BACKSPACE 136
 BLOCK DATA 187
 block IF 80
 BUFFERIN 256
 BUFFEROUT 256
 CALL 89, 201
 CLOSE 131
 COMMON 19, 38
 computed GOTO 76
 continuation of 5
 CONTINUE 89
 control 75
 DATA 59
 DEALLOCATE 26, 233
 debugging 5
 DECODE 117
 DIMENSION 22, 49
 DO 84
 DO WHILE 87
 ELSE 80
 ELSE IF 80
 ELSEWHERE 68, 248
 ENCODE 115
 END 91, 202
 END DO 88
 END IF 80
 ENDMAP 267
 END STRUCTURE 265
 END UNION 266
 ENDFILE 95, 137
 ENTRY 202
 EQUIVALENCE 19, 51, 53
 executable 2 to 3
 EXTERNAL 55, 202
 fields 3
 file-positioning 135
 FIND 118
 FORMAT 101, 149
 formatting 4
 FUNCTION 199
 GOTO 75
 I/O 98
 IF 78
 IF THEN 80
 IMPLICIT 39
 IMPLICIT NONE 40
 INCLUDE 8
 input/output 93
 INQUIRE 132
 INTRINSIC 56, 226
 MAP 267
 multiple 5
 NAMELIST 54, 181
 nested block IF 83
 nonexecutable 2 to 3
 OPEN 96, 119
 OPTIONS 7
 order 6
 PARAMETER 41
 PAUSE 90
 PRINT 114
 PROGRAM 38
 READ 104
 RECORD 265
 RETURN 89, 200 to 202, 204
 REWIND 136
 SAVE 56
 special I/O 115
 specification 37
 STOP 90
 STRUCTURE 265
 SUBROUTINE 202
 TASK COMMON 259
 TYPE 114
 type declaration 43
 unconditional GOTO 75
 UNION 266
 WHERE 67, 248
 WRITE 110
 STATIC statement
 Sun FORTRAN 271
 STATUS keyword 129
 FORTRAN 66 interpretation 228
 status specifier 102
 status variable

- for open operation 127
- STOP statement 90
- VAX interpretation 263
- storage
 - array 28
- storage element
 - finding address 197
- storage requirements
 - controlling 10
- string delimiters 19
- STRUCTURE keyword 265
- STRUCTURE statement 265
- structures
 - declaration 265
 - in RECORD declarations 45
- SU descriptor 170
- subprograms 1, 187
 - actual arguments 187
 - block data 187
 - dummy arguments 187
 - entry points 202
 - functions 195, 199
 - intrinsic functions 195
 - procedures 188
 - subroutine 201
- subroutine
 - returning control to main program 89
- SUBROUTINE statement 202
 - in function subprograms 200
 - subroutine name 202
- subroutine subprograms 201
- subscript
 - defined 27
- substrings
 - character 28
 - equivalencing 53
- SUM intrinsic 218, 240
- Sun FORTRAN compatibility 271
 - \ (backslash) character 271
 - and recursion 271
 - AUTOMATIC statement 271
 - escape sequences 271
 - STATIC statement 271
- suspending program execution 90
- symbolic names 7

T

- T edit descriptor 171
 - and record length 173
- tab character 6
 - SUN FORTRAN 271
- tab edit descriptor 171
- tab-key formatting 6
- TAN intrinsic 208
 - argument 219
- TAND intrinsic 208
 - argument 219
- TANH intrinsic 209

- TASK COMMON statement 259
 - form 259
- technical assistance xxi
 - phone number xxi
- terminating program execution 90
- time subroutine 222
- TL edit descriptor 171
- TR edit descriptor 171
- transfer of sign intrinsic 216
- TRANPOSE intrinsic 218, 245
- truncation 210
 - in assignments 70
- truncation intrinsics 220
- TYPE keyword 129
- TYPE statement 94, 114
- type-declaration statements 43
 - CHARACTER 44
 - initializing arrays in 44
 - initializing character variables with 44
 - initializing variables with 43
 - length specifiers 43
 - numeric 43
- RECORD 45

U

- unconditional GOTO statement 75
- unformatted I/O
 - direct access WRITE 113
 - external direct access READ 107
 - sequential-access READ statement 105
 - sequential-access WRITE statement 111
 - specifying 126
- unformatted I/O records 95
 - header 95
 - trailer 95
- unformatted sequential-access WRITE statement 111
- UNION statement 266
 - form 266
- UNIT function 256
- UNIT keyword 130
- unit numbers
 - VAX 264
- unit specifier 100
- units 96
 - connecting to external files 96
 - program 1
 - specifying 130
- UNPACK intrinsic 218, 243
- user-defined conversion routine sample 143
- user-defined conversions 142
- user-supplied conversion routine names 144
- utilities
 - contact xxi
 - cvbin 138
 - fpr 186

V

- %VAL function 196
- value
 - passing arguments by 196
- variable addresses
 - finding with LOC 46, 197
- variable formats 177
- variable-length I/O records
 - specifying 129
- variables 9, 19
 - as dummy arguments 189
 - see also* data types
 - finding address 197
- VAX data
 - decomposing 139
- VAX files
 - converting 138
- VAX FORTRAN compatibility 261
 - and BLANK keyword 167
 - ANSI compliance 264
 - ASSIGN statement 264
 - binary data files 264
 - converting VAX data 139
 - floating point data 267
 - H descriptor 264
 - intrinsic 269
 - miscellaneous differences 264
 - NOSPANBLOCKS keyword 128
 - octal constants 264
 - READ statement 264
 - REAL*16 constants 264
 - RECL specifier 261
 - record data structure 265
 - record specifier 102
 - STOP message 263
 - structures 265
 - UNION statement 266
 - unit numbers 264
 - unsupported features 262
 - variable-length records 263
- VAX octal constants 15
- VAX records 265
 - %FILL 266
 - form 265
- VAX structures
 - declaration 265
- vector
 - defined 235
- vector multiply functions 235
- vertical format controls 186
- vfc option 261
 - ANSI compliance 264
 - ASSIGN statement 264
 - binary data files 264
 - floating point data 267
 - H descriptor 264
 - intrinsic 269
 - miscellaneous differences 264
 - octal constants 264
 - READ statement 264
 - REAL*16 constants 264
 - RECL specifier 261
 - record data structure 265
 - STOP message 263
 - structures 265
 - UNION statement 266
 - unit numbers 264
 - unsupported features 262
 - variable-length records 263
- vfc option
 - and BLANK keyword 167
- VMS FORTRAN compatibility 261
 - ANSI compliance 264
 - ASSIGN statement 264
 - binary data files 264
 - floating point data 267
 - H descriptor 264
 - intrinsic 269
 - miscellaneous differences 264
 - octal constants 264
 - READ statement 264
 - REAL*16 constants 264
 - records 265
 - STOP message 263
 - structures 265
 - UNION statement 266
 - unit numbers 264
 - unsupported features 262
 - variable-length records 263

W

- WHERE construct 67, 248
 - form 248
- WHERE statement 67, 248
 - form 248
- WRITE statement 94, 110, 184
 - list-directed output formats 184
 - and empty records 175
 - and large files 102
 - direct-access external 112
 - direct-access internal 114
 - external direct-access 112
 - formatted direct access 112
 - formatted sequential-access 111
 - internal 113
 - internal direct access 114
 - internal sequential access 113
 - list-directed formatting 184
 - list-directed sequential-access 112
 - sequential access 111
 - sequential-access internal 113
 - suppressing newlines 173
 - unformatted direct access 113
 - unformatted sequential-access 111
- WRITE statement formatting
 - see* I/O formatting


X

- X edit descriptor 171
 - FORTRAN 66 interpretation 228
- X3.198-199x ANSI standard 235
- XOR
 - bitwise intrinsic 216
- .XOR operator 34

Z

- Z edit descriptor 159
 - and comma field separator 176
 - default value 175
 - with B descriptor 166
- zero-extend intrinsic 212
- ZEXT intrinsic 212
 - result 220










Order Number
DSW-037



Document Number
720-002230-007