



Fortran Library Reference

Forte Developer 7

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-2455-10
May 2002, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright © 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Forte, Java, Solaris, iPlanet, NetBeans, and docs.sun.com are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

Netscape and Netscape Navigator are trademarks or registered trademarks of Netscape Communications Corporation in the United States and other countries.

Sun f90/f95 is derived in part from Cray CF90™, a product of Cray Inc.

libdwarf and liredblack are Copyright 2000 Silicon Graphics Inc. and are available under the GNU Lesser General Public License from <http://www.sgi.com>.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats - Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Forte, Java, Solaris, iPlanet, NetBeans, et docs.sun.com sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Netscape et Netscape Navigator sont des marques de fabrique ou des marques déposées de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays.

Sun f90/f95 est dérivée d'une part de Cray CF90™, un produit de Cray Inc.

libdwarf et liredblack sont Copyright 2000 Silicon Graphics Inc., et sont disponible sur GNU General Public License à <http://www.sgi.com>.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Contents

Before You Begin ix

Typographic Conventions x

Shell Prompts xi

Accessing Forte Developer Development Tools and Man Pages xi

Accessing Forte Developer Documentation xiii

Accessing Related Solaris Documentation xvi

Sending Your Comments xvi

1. Fortran Library Routines 1-1

1.1 Data Type Considerations 1-1

1.2 64-Bit Environments 1-2

1.3 Fortran Math Functions 1-3

1.4 `abort`: Terminate and Write Core File 1-11

1.5 `access`: Check File Permissions or Existence 1-11

1.6 `alarm`: Call Subroutine After a Specified Time 1-12

1.7 `bit`: Bit Functions: `and`, `or`, ..., `bit`, `setbit`, ... 1-13

1.8 `chdir`: Change Default Directory 1-16

1.9 `chmod`: Change the Mode of a File 1-17

1.10 `date`: Get Current Date as a Character String 1-18

- 1.11 `date_and_time`: Get Date and Time 1-19
- 1.12 `dtime, etime`: Elapsed Execution Time 1-20
- 1.13 `exit`: Terminate a Process and Set the Status 1-23
- 1.14 `fdate`: Return Date and Time in an ASCII String 1-24
- 1.15 `flush`: Flush Output to a Logical Unit 1-25
- 1.16 `fork`: Create a Copy of the Current Process 1-26
- 1.17 `fseek, ftell`: Determine Position and Reposition a File 1-27
- 1.18 `fseeko64, ftello64`: Determine Position and Reposition a Large File 1-29
- 1.19 `getarg, iargc`: Get Command-Line Arguments 1-31
- 1.20 `getc, fgetc`: Get Next Character 1-32
- 1.21 `getcwd`: Get Path of Current Working Directory 1-35
- 1.22 `getenv`: Get Value of Environment Variables 1-36
- 1.23 `getfd`: Get File Descriptor for External Unit Number 1-37
- 1.24 `getfilep`: Get File Pointer for External Unit Number 1-38
- 1.25 `getlog`: Get User's Login Name 1-39
- 1.26 `getpid`: Get Process ID 1-40
- 1.27 `getuid, getgid`: Get User or Group ID of Process 1-40
- 1.28 `hostnm`: Get Name of Current Host 1-42
- 1.29 `idate`: Return Current Date 1-42
- 1.30 `ieee_flags, ieee_handler, sigfpe`: IEEE Arithmetic 1-43
- 1.31 `index, rindex, lnblnk`: Index or Length of Substring 1-48
- 1.32 `inmax`: Return Maximum Positive Integer 1-50
- 1.33 `itime`: Current Time 1-51
- 1.34 `kill`: Send a Signal to a Process 1-51
- 1.35 `link, symlink`: Make a Link to an Existing File 1-52
- 1.36 `loc`: Return the Address of an Object 1-54

- 1.37 long, short: Integer Object Conversion 1-54
 - 1.38 longjmp, setjmp: Return to Location Set by setjmp 1-56
 - 1.39 malloc, malloc64, realloc, free: Allocate/Reallocate/Deallocate Memory 1-58
 - 1.40 mvbits: Move a Bit Field 1-62
 - 1.41 perror, gerror, ierrno: Get System Error Messages 1-63
 - 1.42 putc, fputc: Write a Character to a Logical Unit 1-65
 - 1.43 qsort, qsort64: Sort the Elements of a One-Dimensional Array 1-67
 - 1.44 ran: Generate a Random Number Between 0 and 1 1-69
 - 1.45 rand, drand, irand: Return Random Values 1-71
 - 1.46 rename: Rename a File 1-72
 - 1.47 secnds: Get System Time in Seconds, Minus Argument 1-73
 - 1.48 sh: Fast Execution of an sh Command 1-74
 - 1.49 signal: Change the Action for a Signal 1-75
 - 1.50 sleep: Suspend Execution for an Interval 1-76
 - 1.51 stat, lstat, fstat: Get File Status 1-76
 - 1.52 stat64, lstat64, fstat64: Get File Status 1-80
 - 1.53 system: Execute a System Command 1-80
 - 1.54 time, ctime, ltime, gmtime: Get System Time 1-81
 - 1.55 ttyname, isatty: Get Name of a Terminal Port 1-85
 - 1.56 unlink: Remove a File 1-86
 - 1.57 wait: Wait for a Process to Terminate 1-87
- 2. Fortran 95 Intrinsic Functions 2-1**
 - 2.1 Standard Fortran 95 Generic Intrinsic Functions 2-1
 - 2.2 Non-Standard Fortran 95 Intrinsic Functions 2-13
- 3. FORTRAN 77 and VMS Intrinsic Functions 3-1**
 - 3.1 Arithmetic and Mathematical Functions 3-2

3.2	Character Functions	3-11
3.3	Miscellaneous Functions	3-12
3.4	Remarks	3-14
3.5	VMS Intrinsic Functions	3-21
Index		Index-1

Tables

TABLE 1-1	Library Routines for 64-bit Environments	1-3
TABLE 1-2	Single-Precision Math Functions	1-4
TABLE 1-3	Double Precision Math Functions	1-7
TABLE 1-4	Quadruple-Precision libm Functions	1-10
TABLE 1-5	IEEE Arithmetic Support Routines	1-43
TABLE 1-6	<code>ieee_flags(action,mode,in,out)</code> Parameters and Actions	1-44
TABLE 1-7	<code>ieee_handler(action,in,out)</code> Parameters	1-45
TABLE 2-1	Specific and Generic Names for Fortran 95 Intrinsic Functions	2-10
TABLE 2-2	BLAS Intrinsics	2-13
TABLE 2-3	Intrinsic functions from Cray CF90 and other compilers	2-15
TABLE 3-1	Fortran 77 Arithmetic Functions	3-3
TABLE 3-2	Fortran 77 Type Conversion Functions	3-5
TABLE 3-3	Fortran 77 Trigonometric Functions	3-7
TABLE 3-4	Other Fortran 77 Mathematical Functions	3-9
TABLE 3-5	Fortran 77 Character Functions	3-11
TABLE 3-6	Fortran 77 Bitwise Functions	3-12
TABLE 3-7	Fortran 77 Environmental Inquiry Functions	3-13
TABLE 3-8	Fortran 77 Memory Functions	3-14
TABLE 3-9	VMS Double-Precision Complex Functions	3-21
TABLE 3-10	VMS Degree-Based Trigonometric Functions	3-22

TABLE 3-11 VMS Bit-Manipulation Functions 3-23

TABLE 3-12 VMS Integer Functions 3-24

Before You Begin

The *Fortran Library Reference* describes the intrinsic functions and routines in the Forte™ Developer Fortran libraries. This reference manual is intended for programmers with a working knowledge of the Fortran language and the Solaris™ operating environment.

This guide is intended for scientists, engineers, and programmers who have a working knowledge of the Fortran language and wish to learn how to use the Sun Fortran compilers effectively. Familiarity with the Solaris operating environment or UNIX® in general is also assumed.

Discussion of Fortran programming issues on Solaris operating environments, including input/output, application development, library creating and use, program analysis, porting, optimization, and parallelization can be found in the companion Forte Developer *Fortran Programming Guide*.

Typographic Conventions

The following are the typographic conventions used in the text and code examples in this manual:

TABLE P-1 Typeface Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
<i>AaBbCc123</i>	Command-line placeholder text; replace with a real name or value	To delete a file, type rm <i>filename</i> .

- The symbol Δ stands for a blank space where a blank is significant:

ΔΔ36.001

- The FORTRAN 77 standard used an older convention, spelling the name “FORTRAN” capitalized. The Forte Developer Fortran compiler documentation uses both “FORTRAN” and “Fortran”. The current convention is to use lower case: “Fortran 95”
- References to online man pages appear with the topic name and section number. For example, a reference to the library routine GETENV will appear as `getenv(3F)`, implying that the man command to access this man page would be:
`man -s 3F getenv.`

TABLE P-2 Code and Command-Line Conventions

Code Symbol	Meaning	Notation	Code Example
[]	Brackets contain arguments that are optional.	<code>0[n]</code>	<code>04, 0</code>
{ }	Braces contain a set of choices for required option.	<code>d{y n}</code>	<code>dy</code>
	The “pipe” or “bar” symbol separates arguments, only one of which may be chosen.	<code>B{dynamic static}</code>	<code>Bstatic</code>
:	The colon, like the comma, is sometimes used to separate arguments.	<code>Rdir[:dir]</code>	<code>R/local/libs:/U/a</code>
...	The ellipsis indicates omission in a series.	<code>xinline=f1[,...fn]</code>	<code>xinline=alpha,dos</code>

Shell Prompts

Shell	Prompt
C shell	<code>%</code>
Bourne shell and Korn shell	<code>\$</code>
C shell, Bourne shell, and Korn shell superuser	<code>#</code>

Accessing Forte Developer Development Tools and Man Pages

The Forte Developer product components and man pages are not installed into the standard `/usr/bin/` and `/usr/share/man` directories. To access the Forte Developer compilers and tools, you must have the Forte Developer component

directory in your `PATH` environment variable. To access the Forte Developer man pages, you must have the Forte Developer man page directory in your `MANPATH` environment variable.

For more information about the `PATH` variable, see the `csh(1)`, `sh(1)`, and `ksh(1)` man pages. For more information about the `MANPATH` variable, see the `man(1)` man page. For more information about setting your `PATH` and `MANPATH` variables to access this Forte Developer release, see the installation guide or your system administrator.

Note – The information in this section assumes that your Forte Developer products are installed in the `/opt` directory. If your product software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

Accessing Forte Developer Compilers and Tools

Use the steps below to determine whether you need to change your `PATH` variable to access the Forte Developer compilers and tools.

▼ To Determine Whether You Need to Set Your `PATH` Environment Variable

1. Display the current value of the `PATH` variable by typing the following at a command prompt:

```
% echo $PATH
```

2. Review the output for a string of paths that contain `/opt/SUNWspro/bin/`.

If you find the path, your `PATH` variable is already set to access Forte Developer development tools. If you do not find the path, set your `PATH` environment variable by following the instructions in the next section.

▼ To Set Your `PATH` Environment Variable to Enable Access to Forte Developer Compilers and Tools

1. If you are using the C shell, edit your home `.cshrc` file. If you are using the Bourne shell or Korn shell, edit your home `.profile` file.
2. Add the following to your `PATH` environment variable.

```
/opt/SUNWspro/bin
```

Accessing Forte Developer Man Pages

Use the following steps to determine whether you need to change your MANPATH variable to access the Forte Developer man pages.

▼ To Determine Whether You Need to Set Your MANPATH Environment Variable

1. Request the dbx man page by typing the following at a command prompt:

```
% man dbx
```

2. Review the output, if any.

If the dbx(1) man page cannot be found or if the man page displayed is not for the current version of the software installed, follow the instructions in the next section for setting your MANPATH environment variable.

▼ To Set Your MANPATH Environment Variable to Enable Access to Forte Developer Man Pages

1. If you are using the C shell, edit your home .cshrc file. If you are using the Bourne shell or Korn shell, edit your home .profile file.
2. Add the following to your MANPATH environment variable.
/opt/SUNWspro/man

Accessing Forte Developer Documentation

You can access Forte Developer product documentation at the following locations:

- The product documentation is available from the documentation index installed with the product on your local system or network at
/opt/SUNWspro/docs/index.html.

If your product software is not installed in the /opt directory, ask your system administrator for the equivalent path on your system.

- Most manuals are available from the docs.sun.comsm web site. The following titles are available through your installed product only:

- *Standard C++ Library Class Reference*
- *Standard C++ Library User's Guide*
- *Tools.h++ Class Library Reference*
- *Tools.h++ User's Guide*

The docs.sun.com web site (<http://docs.sun.com>) enables you to read, print, and buy Sun Microsystems manuals through the Internet. If you cannot find a manual, see the documentation index installed with the product on your local system or network.

Note – Sun is not responsible for the availability of third-party web sites mentioned in this document and does not endorse and is not responsible or liable for any content, advertising, products, or other materials on or available from such sites or resources. Sun will not be responsible or liable for any damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services available on or through any such sites or resources.

Product Documentation in Accessible Formats

Forte Developer 7 product documentation is provided in accessible formats that are readable by assistive technologies for users with disabilities. You can find accessible versions of documentation as described in the following table. If your product software is not installed in the /opt directory, ask your system administrator for the equivalent path on your system.

Type of Documentation	Format and Location of Accessible Version
Manuals (except third-party manuals)	HTML at http://docs.sun.com
Third-party manuals: • <i>Standard C++ Library Class Reference</i> • <i>Standard C++ Library User's Guide</i> • <i>Tools.h++ Class Library Reference</i> • <i>Tools.h++ User's Guide</i>	HTML in the installed product through the documentation index at <code>file:/opt/SUNWspro/docs/index.html</code>
Readmes and man pages	HTML in the installed product through the documentation index at <code>file:/opt/SUNWspro/docs/index.html</code>
Release notes	Text file on the product CD at <code>/cdrom/devpro_v10n1_sparc/release_notes.txt</code>

Related Forte Developer Documentation

The following table describes related documentation that is available at `file:/opt/SUNWspro/docs/index.html`. If your product software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

Document Title	Description
<i>Fortran Programming Guide</i>	Describes how to write effective Fortran code on Solaris environments; input/output, libraries, performance, debugging, and parallel processing.
<i>Fortran User's Guide</i>	Describes the compile-time environment and command-line options for the f95 compiler. Also includes guidelines for migrating legacy f77 programs to f95.
<i>OpenMP API User's Guide</i>	Summary of the OpenMP multiprocessing API, with details on the Forte Developer implementation.
<i>Numerical Computation Guide</i>	Describes issues regarding the numerical accuracy of floating-point computations.

Accessing Related Solaris Documentation

The following table describes related documentation that is available through the docs.sun.com web site.

Document Collection	Document Title	Description
Solaris Reference Manual Collection	See the titles of man page sections.	Provides information about the Solaris operating environment.
Solaris Software Developer Collection	<i>Linker and Libraries Guide</i>	Describes the operations of the Solaris link-editor and runtime linker.
Solaris Software Developer Collection	<i>Multithreaded Programming Guide</i>	Covers the POSIX and Solaris threads APIs, programming with synchronization objects, compiling multithreaded programs, and finding tools for multithreaded programs.

Sending Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. Email your comments to Sun at this address:

docfeedback@sun.com

Fortran Library Routines

This chapter describes the Fortran library routines.

All the routines described in this chapter have corresponding man pages in section 3F of the man library. For example, **man -s 3F access** will display the man page entry for the library routine `access`.

See also the *Numerical Computation Guide* for additional math routines that are callable from Fortran and C. These include the standard math library routines in `libm` and `libsunmath` (see `Intro(3M)`), optimized versions of these libraries, the SPARC vector math library, `libmvec`, and others.

See Chapter 3 for details on Fortran 77 and VMS intrinsic functions implemented by the `f95` compiler.

1.1 Data Type Considerations

Unless otherwise indicated, the function routines listed here are not intrinsics. That means that the type of data a function returns may conflict with the implicit typing of the function name, and require explicit type declaration by the user. For example, `getpid()` returns `INTEGER*4` and would require an `INTEGER*4 getpid` declaration to ensure proper handling of the result. (Without explicit typing, a `REAL` result would be assumed by default because the function name starts with `g`.) As a reminder, explicit type statements appear in the function summaries for these routines.

Be aware that `IMPLICIT` statements and the `-dbl` and `-xtypemap` compiler options also alter the data typing of arguments and the treatment of return values. A mismatch between the expected and actual data types in calls to these library routines could cause unexpected behavior. Options `-xtypemap` and `-dbl` promote the data type of `INTEGER` functions to `INTEGER*8`, `REAL` functions to `REAL*8`, and

DOUBLE functions to REAL*16. To protect against these problems, function names and variables appearing in library calls should be explicitly typed with their expected sizes, as in:

```
integer*4 seed, getuid
real*4 ran
...
seed = 70198
val = getuid() + ran(seed)
...
```

Explicit typing in the example protects the library calls from any data type promotion when the `-xtypemap` and `-dbl` compiler options are used. Without explicit typing, these options could produce unexpected results. See the *Fortran User's Guide* and the `f95(1)` man page for details on these options.

The Fortran 95 compiler, `f95`, provides an include file, `system.inc`, that defines the interfaces for most non-intrinsic library routines. Include this file to insure that functions you call and their arguments are properly typed, especially when default data types are changed with `-xtypemap`.

```
include 'system.inc'
integer(4) mypid
mypid = getpid()
print *, mypid
```

You can catch many issues related to type mismatches over library calls by using the Fortran compilers' global program checking option, `-Xlist`. Global program checking by the `f95` compiler is described in the *Fortran User's Guide*, the *Fortran Programming Guide*, and the `f95(1)` man page.

1.2 64-Bit Environments

Compiling a program to run in a 64-bit operating environment (that is, compiling with `-xarch=v9`, `v9a`, or `v9b` and running the executable on a SPARC platform running the 64-bit enabled Solaris operating environment) changes the return values of certain functions. These are usually functions that interface standard system-level routines, such as `malloc(3F)` (see “`malloc`, `malloc64`, `realloc`, `free`: Allocate/Reallocate/Deallocate Memory” on page 58), and may take or return 32-bit or 64-bit values depending on the environment. To provide portability of code

between 32-bit and 64-bit environments, 64-bit versions of these routines have been provided that always take and/or return 64-bit values. The following table identifies library routine provided for use in 64-bit environments:

TABLE 1-1 Library Routines for 64-bit Environments

Function	Description
<code>malloc64</code>	Allocate memory and return a pointer
<code>fseeko64</code>	Reposition a large file
<code>ftello64</code>	Determine position of a large file
<code>stat64</code> , <code>fstat64</code> , <code>lstat64</code>	Determine status of a file
<code>time64</code> , <code>ctime64</code> , <code>gmtime64</code> , <code>ltime64</code>	Dissect system time, convert to character
<code>qsort64</code>	Sort the elements of an array

1.3 Fortran Math Functions

The following functions and subroutines are part of the Fortran math libraries. They are available to all programs compiled with `f95`. These routines are non-intrinsics that take a specific data type as an argument and return the same. Non-intrinsics do have to be declared in the routine referencing them.

Many of these routines are "wrappers", Fortran interfaces to routines in the C language library, and as such are non-standard Fortran. They include IEEE recommended support functions, and specialized random number generators. See the *Numerical Computation Guide* and the man pages `libm_single(3F)`, `libm_double(3F)`, `libm_quadruple(3F)`, for more information about these libraries.

1.3.1 Single-Precision Functions

These subprograms are single-precision math functions and subroutines.

In general, the functions below provide access to single-precision math functions that do *not* correspond to standard Fortran generic intrinsic functions—data types are determined by the usual data typing rules.

These functions need not be explicitly typed with a `REAL` statement as long as default typing holds. (Names beginning with “r” are `REAL`, with “i” are `INTEGER`.)

For details on these routines, see the C math library man pages (3M). For example, for `r_acos(x)` see the `acos(3M)` man page.

TABLE 1-2 Single-Precision Math Functions

Function Name	Return Type	Description
<code>r_acos(x)</code>	REAL	arc cosine
<code>r_acosd(x)</code>	REAL	--
<code>r_acosh(x)</code>	REAL	arc cosh
<code>r_acosp(x)</code>	REAL	--
<code>r_acospi(x)</code>	REAL	--
<code>r_atan(x)</code>	REAL	arc tangent
<code>r_atand(x)</code>	REAL	--
<code>r_atanh(x)</code>	REAL	arc tanh
<code>r_atanp(x)</code>	REAL	--
<code>r_atanpi(x)</code>	REAL	--
<code>r_asin(x)</code>	REAL	arc sine
<code>r_asind(x)</code>	REAL	--
<code>r_asinh(x)</code>	REAL	arc sinh
<code>r_asinp(x)</code>	REAL	--
<code>r_asinpi(x)</code>	REAL	--
<code>r_atan2((y, x)</code>	REAL	arc tangent
<code>r_atan2d(y, x)</code>	REAL	--
<code>r_atan2pi(y, x)</code>	REAL	--
<code>r_cbrt(x)</code>	REAL	cube root
<code>r_ceil(x)</code>	REAL	ceiling
<code>r_copysign(x, y)</code>	REAL	--
<code>r_cos(x)</code>	REAL	cosine
<code>r_cosd(x)</code>	REAL	--
<code>r_cosh(x)</code>	REAL	hyperb cos
<code>r_cosp(x)</code>	REAL	--
<code>r_cospi(x)</code>	REAL	--
<code>r_erf(x)</code>	REAL	err function
<code>r_erfc(x)</code>	REAL	--

TABLE 1-2 Single-Precision Math Functions (*Continued*)

Function Name	Return Type	Description
r_expm1(x)	REAL	(e**x)-1
r_floor(x)	REAL	floor
r_hypot(x, y)	REAL	hypotenuse
r_infinity()	REAL	--
r_j0(x)	REAL	Bessel
r_j1(x)	REAL	--
r_jn(x)	REAL	--
ir_finite(x)	INTEGER	--
ir_fp_class(x)	INTEGER	--
ir_ilogb(x)	INTEGER	--
ir_rint(x)	INTEGER	--
ir_isinf(x)	INTEGER	--
ir_isnan(x)	INTEGER	--
ir_isnormal(x)	INTEGER	--
ir_issubnormal(x)	INTEGER	--
ir_iszero(x)	INTEGER	--
ir_signbit(x)	INTEGER	--
r_addran()	REAL	random number generators
r_addrans(x, p, l, u)	subroutine	
r_lcran()	REAL	
r_lcrans(x, p, l, u)	subroutine	
r_shufrans(x, p, l, u)	subroutine	
r_lgamma(x)	REAL	log gamma
r_logb(x)	REAL	--
r_log1p(x)	REAL	--
r_log2(x)	REAL	--
r_max_normal()	REAL	
r_max_subnormal()	REAL	
r_min_normal()	REAL	
r_min_subnormal()	REAL	
r_nextafter(x, y)	REAL	
r_quiet_nan(n)	REAL	
r_remainder(x, y)	REAL	
r_rint(x)	REAL	
r_scalb(x, y)	REAL	
r_scalbn(x, n)	REAL	
r_signaling_nan(n)	REAL	
r_significand(x)	REAL	

TABLE 1-2 Single-Precision Math Functions (*Continued*)

Function Name	Return Type	Description
r_sin(x)	REAL	sine
r_sind(x)	REAL	--
r_sinh(x)	REAL	hyperb sin
r_sinp(x)	REAL	--
r_sinpi(x)	REAL	--
r_sincos(x, s, c)	subroutine	sine & cosine
r_sincosd(x, s, c)	subroutine	--
r_sincosp(x, s, c)	subroutine	--
r_sincospi(x, s, c)	subroutine	--
r_tan(x)	REAL	tangent
r_tand(x)	REAL	--
r_tanh(x)	REAL	hyperb tan
r_tanp(x)	REAL	--
r_tanpi(x)	REAL	--
r_y0(x)	REAL	bessel
r_y1(x)	REAL	--
r_yn(n, x)	REAL	--

- Variables c, l, p, s, u, x, and y are of type REAL.
- Type these functions as explicitly REAL if an IMPLICIT statement is in effect that types names starting with “r” to some other data type.
- sind(x), asind(x), ... take *degrees* rather than *radians*.

See also: intro(3M) and the *Numerical Computation Guide*.

1.3.2 Double-Precision Functions

The following subprograms are double-precision math functions and subroutines.

In general, these functions do *not* correspond to standard Fortran generic intrinsic functions—data types are determined by the usual data typing rules.

These DOUBLE PRECISION functions need to appear in a DOUBLE PRECISION statement.

Refer to the C library man pages for details: the man page for `d_acos(x)` is `acos(3M)`

TABLE 1-3 Double Precision Math Functions

Function Name	Return Type	Description
<code>d_acos(x)</code>	DOUBLE PRECISION	arc cosine
<code>d_acosd(x)</code>	DOUBLE PRECISION	--
<code>d_acosh(x)</code>	DOUBLE PRECISION	arc cosh
<code>d_acosp(x)</code>	DOUBLE PRECISION	--
<code>d_acospi(x)</code>	DOUBLE PRECISION	--
<code>d_atan(x)</code>	DOUBLE PRECISION	arc tangent
<code>d_atand(x)</code>	DOUBLE PRECISION	--
<code>d_atanh(x)</code>	DOUBLE PRECISION	arc tanh
<code>d_atanp(x)</code>	DOUBLE PRECISION	--
<code>d_atanpi(x)</code>	DOUBLE PRECISION	--
<code>d_asin(x)</code>	DOUBLE PRECISION	arc sine
<code>d_asind(x)</code>	DOUBLE PRECISION	--
<code>d_asinh(x)</code>	DOUBLE PRECISION	arc sinh
<code>d_asinp(x)</code>	DOUBLE PRECISION	--
<code>d_asinpi(x)</code>	DOUBLE PRECISION	--
<code>d_atan2((y, x)</code>	DOUBLE PRECISION	arc tangent
<code>d_atan2d(y, x)</code>	DOUBLE PRECISION	--
<code>d_atan2pi(y, x)</code>	DOUBLE PRECISION	--
<code>d_cbrt(x)</code>	DOUBLE PRECISION	cube root
<code>d_ceil(x)</code>	DOUBLE PRECISION	ceiling
<code>d_copysign(x, x)</code>	DOUBLE PRECISION	--
<code>d_cos(x)</code>	DOUBLE PRECISION	cosine
<code>d_cosd(x)</code>	DOUBLE PRECISION	--
<code>d_cosh(x)</code>	DOUBLE PRECISION	hyperb cos
<code>d_cosp(x)</code>	DOUBLE PRECISION	--
<code>d_cospi(x)</code>	DOUBLE PRECISION	--
<code>d_erf(x)</code>	DOUBLE PRECISION	error func
<code>d_erfc(x)</code>	DOUBLE PRECISION	--
<code>d_expm1(x)</code>	DOUBLE PRECISION	$(e^x)-1$
<code>d_floor(x)</code>	DOUBLE PRECISION	floor
<code>d_hypot(x, y)</code>	DOUBLE PRECISION	hypotenuse
<code>d_infinity()</code>	DOUBLE PRECISION	--

TABLE 1-3 Double Precision Math Functions (*Continued*)

Function Name	Return Type	Description
d_j0(x)	DOUBLE PRECISION	Bessel
d_j1(x)	DOUBLE PRECISION	--
d_jn(x)	DOUBLE PRECISION	--
id_finite(x)	INTEGER	
id_fp_class(x)	INTEGER	
id_ilogb(x)	INTEGER	
id_rint(x)	INTEGER	
id_isinf(x)	INTEGER	
id_isnan(x)	INTEGER	
id_isnormal(x)	INTEGER	
id_issubnormal(x)	INTEGER	
id_iszero(x)	INTEGER	
id_signbit(x)	INTEGER	
d_addran()	DOUBLE PRECISION	random number generators
d_addrans(x, p, l, u)	subroutine	
d_lcran()	DOUBLE PRECISION	
d_lcrans(x, p, l, u)	subroutine	
d_shufrans(x, p, l,u)	subroutine	
d_lgamma(x)	DOUBLE PRECISION	log gamma
d_logb(x)	DOUBLE PRECISION	--
d_log1p(x)	DOUBLE PRECISION	--
d_log2(x)	DOUBLE PRECISION	--
d_max_normal()	DOUBLE PRECISION	
d_max_subnormal()	DOUBLE PRECISION	
d_min_normal()	DOUBLE PRECISION	
d_min_subnormal()	DOUBLE PRECISION	
d_nextafter(x, y)	DOUBLE PRECISION	
d_quiet_nan(n)	DOUBLE PRECISION	
d_remainder(x, y)	DOUBLE PRECISION	
d_rint(x)	DOUBLE PRECISION	
d_scalb(x, y)	DOUBLE PRECISION	
d_scalbn(x, n)	DOUBLE PRECISION	
d_signaling_nan(n)	DOUBLE PRECISION	
d_significand(x)	DOUBLE PRECISION	

TABLE 1-3 Double Precision Math Functions (*Continued*)

Function Name	Return Type	Description
d_sin(x)	DOUBLE PRECISION	sine
d_sind(x)	DOUBLE PRECISION	--
d_sinh(x)	DOUBLE PRECISION	hyper sine
d_sinp(x)	DOUBLE PRECISION	--
d_sinpi(x)	DOUBLE PRECISION	--
d_sincos(x, s, c)	subroutine	sine & cosine
d_sincosd(x, s, c)	subroutine	--
d_sincosp(x, s, c)	subroutine	--
d_sincospi(x, s, c)	subroutine	--
d_tan(x)	DOUBLE PRECISION	tangent
d_tand(x)	DOUBLE PRECISION	--
d_tanh(x)	DOUBLE PRECISION	hyperb tan
d_tanp(x)	DOUBLE PRECISION	--
d_tanpi(x)	DOUBLE PRECISION	--
d_y0(x)	DOUBLE PRECISION	bessel
d_y1(x)	DOUBLE PRECISION	--
d_yn(n, x)	DOUBLE PRECISION	--

- Variables c, l, p, s, u, x, and y are of type DOUBLE PRECISION.
- Explicitly type these functions on a DOUBLE PRECISION statement or with an appropriate IMPLICIT statement).
- sind(x), asind(x), ... take *degrees* rather than *radians*.

See also: intro(3M) and the *Numerical Computation Guide*.

1.3.3 Quad-Precision Functions

These subprograms are quadruple-precision (REAL*16) math functions and subroutines.

In general, these do *not* correspond to standard generic intrinsic functions; data types are determined by the usual data typing rules.

The quadruple precision functions must appear in a `REAL*16` statement

TABLE 1-4 Quadruple-Precision `libm` Functions

Function Name	Return Type
<code>q_copysign(x, y)</code>	<code>REAL*16</code>
<code>q_fabs(x)</code>	<code>REAL*16</code>
<code>q_fmod(x)</code>	<code>REAL*16</code>
<code>q_infinity()</code>	<code>REAL*16</code>
<code>iq_finite(x)</code>	<code>INTEGER</code>
<code>iq_fp_class(x)</code>	<code>INTEGER</code>
<code>iq_ilogb(x)</code>	<code>INTEGER</code>
<code>iq_isinf(x)</code>	<code>INTEGER</code>
<code>iq_isnan(x)</code>	<code>INTEGER</code>
<code>iq_isnormal(x)</code>	<code>INTEGER</code>
<code>iq_issubnormal(x)</code>	<code>INTEGER</code>
<code>iq_iszero(x)</code>	<code>INTEGER</code>
<code>iq_signbit(x)</code>	<code>INTEGER</code>
<code>q_max_normal()</code>	<code>REAL*16</code>
<code>q_max_subnormal()</code>	<code>REAL*16</code>
<code>q_min_normal()</code>	<code>REAL*16</code>
<code>q_min_subnormal()</code>	<code>REAL*16</code>
<code>q_nextafter(x, y)</code>	<code>REAL*16</code>
<code>q_quiet_nan(n)</code>	<code>REAL*16</code>
<code>q_remainder(x, y)</code>	<code>REAL*16</code>
<code>q_scalbn(x, n)</code>	<code>REAL*16</code>
<code>q_signaling_nan(n)</code>	<code>REAL*16</code>

- The variables `c`, `l`, `p`, `s`, `u`, `x`, and `y` are of type quadruple precision.
- Explicitly type these functions with a `REAL*16` statement or with an appropriate `IMPLICIT` statement.
- `sind(x)`, `asind(x)`, ... take *degrees* rather than *radians*.

If you need to use any other quadruple-precision `libm` function, you can call it using `$PRAGMA C(fcn)` before the call. For details, see the chapter on the C–Fortran interface in the *Fortran Programming Guide*.

1.4 abort: Terminate and Write Core File

The subroutine is called by:

call abort

`abort` flushes the I/O buffers and then aborts the process, possibly producing a core file memory dump in the current directory. See `limit(1)` about limiting or suppressing core dumps.

1.5 access: Check File Permissions or Existence

The function is called by:

INTEGER*4 access status = access (name, mode)			
name	character	Input	File name
mode	character	Input	Permissions
Return value	INTEGER*4	Output	status=0: OK status>0: Error code

`access` determines if you can access the file *name* with the permissions specified by *mode*. `access` returns zero if the access specified by *mode* would be successful. See also `gerror(3F)` to interpret error codes.

Set *mode* to one or more of *r*, *w*, *x*, in any order or combination, or blank, where *r*, *w*, *x* have the following meanings:

'r'	Test for read permission
'w'	Test for write permission
'x'	Test for execute permission
' '	Test for existence of the file

Example 1: Test for read/write permission:

```
INTEGER*4 access, status
status = access ( 'taccess.data', 'rw' )
if ( status .eq. 0 ) write(*,*) "ok"
if ( status .ne. 0 ) write(*,*) 'cannot read/write', status
```

Example 2: Test for existence:

```
INTEGER*4 access, status
status = access ( 'taccess.data', ' ' ) ! blank mode
if ( status .eq. 0 ) write(*,*) "file exists"
if ( status .ne. 0 ) write(*,*) 'no such file', status
```

1.6 alarm: Call Subroutine After a Specified Time

The function is called by:

INTEGER*4 alarm n = alarm (<i>time</i> , <i>sbrtn</i>)			
<i>time</i>	INTEGER*4	Input	Number of seconds to wait (0=do not call)
<i>sbrtn</i>	Routine name	Input	Subprogram to execute must be listed in an external statement.
Return value	INTEGER*4	Output	Time remaining on the last alarm

Example: `alarm`—wait 9 seconds then call `sbrtn`:

```
integer*4 alarm, time / 1 /
common / alarmcom / i
external sbrtn
i = 9
write(*,*) i
nseconds = alarm ( time, sbrtn )
do n = 1,100000          ! Wait until alarm activates sbrtn.
  r = n                  ! (any calculations that take enough time)
  x=sqrt(r)
end do
write(*,*) i
end
subroutine sbrtn
common / alarmcom / i
i = 3                    ! Do no I/O in this routine.
return
end
```

See also: `alarm(3C)`, `sleep(3F)`, and `signal(3F)`. Note the following restrictions:

- A subroutine cannot pass its own name to `alarm`.
- The `alarm` routine generates signals that could interfere with any I/O. The called subroutine, *sbrtn*, must not do any I/O itself.
- Calling `alarm()` from a parallelized or multi-threaded Fortran program may have unpredictable results.

1.7 bit: Bit Functions: `and`, `or`, ..., `bit`, `setbit`, ...

The definitions are:

<code>and(word1, word2)</code>	Computes bitwise <i>and</i> of its arguments.
<code>or(word1, word2)</code>	Computes bitwise <i>inclusive or</i> of its arguments.
<code>xor(word1, word2)</code>	Computes bitwise <i>exclusive or</i> of its arguments.
<code>not(word)</code>	Returns bitwise <i>complement</i> of its argument.
<code>lshift(word, nbits)</code>	Logical left shift with no end around carry.

<code>rshift(word, nbits)</code>	Arithmetic right shift with sign extension.
<code>call bis(bitnum, word)</code>	Sets bit <i>bitnum</i> in <i>word</i> to 1.
<code>call bic(bitnum, word)</code>	Clears bit <i>bitnum</i> in <i>word</i> to 0.
<code>bit(bitnum, word)</code>	Tests bit <i>bitnum</i> in <i>word</i> and returns <code>.true.</code> if the bit is 1, <code>.false.</code> if it is 0.
<code>call setbit(bitnum,word,state)</code>	Sets bit <i>bitnum</i> in <i>word</i> to 1 if <i>state</i> is nonzero, and clears it otherwise.

The alternate external versions for MIL-STD-1753 are:

<code>iband(m, n)</code>	Computes the bitwise <i>and</i> of its arguments.
<code>ior(m, n)</code>	Computes the bitwise <i>inclusive or</i> of its arguments.
<code>ieor(m, n)</code>	Computes the bitwise <i>exclusive or</i> of its arguments.
<code>ishft(m, k)</code>	Is a logical shift with no end around carry (left if $k > 0$, right if $k < 0$).
<code>ishftc(m, k, ic)</code>	Circular shift: right-most <i>ic</i> bits of <i>m</i> are left-shifted circularly <i>k</i> places.
<code>ibits(m, i, len)</code>	Extracts bits: from <i>m</i> , starting at bit <i>i</i> , extracts <i>len</i> bits.
<code>ibset(m, i)</code>	Sets bit: return value is equal to word <i>m</i> with bit number <i>i</i> set to 1.
<code>ibclr(m, i)</code>	Clears bit: return value is equal to word <i>m</i> with bit number <i>i</i> set to 0.
<code>btest(m, i)</code>	Tests bit <i>i</i> in <i>m</i> ; returns <code>.true.</code> if the bit is 1, and <code>.false.</code> if it is 0.

See also “mvbits: Move a Bit Field” on page 62, and Chapters 2 and 3 for other functions that manipulate bit fields.

1.7.1 Usage: and, or, xor, not, rshift, lshift

For the intrinsic functions:

```

x = and( word1, word2 )
x = or( word1, word2 )
x = xor( word1, word2 )
x = not( word )
x = rshift( word, nbits )
x = lshift( word, nbits )

```

word, *word1*, *word2*, *nbits* are integer input arguments. These are intrinsic functions expanded inline by the compiler. The data type returned is that of the first argument.

No test is made for a reasonable value of *nbits*.

Example: and, or, xor, not:

```
demo% cat tandornot.f
      print 1, and(7,4), or(7,4), xor(7,4), not(4)
1      format(4x 'and(7,4)', 5x 'or(7,4)', 4x 'xor(7,4)',
      1      6x 'not(4)'/4o12.11)
      end
demo% f95 tandornot.f
demo% a.out
      and(7,4)      or(7,4)      xor(7,4)      not(4)
000000000004 000000000007 000000000003 377777777773
demo%
```

Example: lshift, rshift:

```
      integer*4 lshift, rshift
      print 1, lshift(7,1), rshift(4,1)
1      format(1x 'lshift(7,1)', 1x 'rshift(4,1)'/2o12.11)
      end
demo% f95 tlrshift.f
demo% a.out
      lshift(7,1) rshift(4,1)
000000000016 000000000002
demo%
```

1.7.2 Usage: bic, bis, bit, setbit

For the subroutines and functions

```
call bic( bitnum, word )
call bis( bitnum, word )
call setbit( bitnum, word, state )
LOGICAL bit
x = bit( bitnum, word )
```

bitnum, *state*, and *word* are INTEGER*4 input arguments. Function *bit()* returns a logical value.

Bits are numbered so that bit 0 is the least significant bit, and bit 31 is the most significant.

bic, *bis*, and *setbit* are external subroutines. *bit* is an external function.

Example 3: *bic*, *bis*, *setbit*, *bit*:

```
integer*4 bitnum/2/, state/0/, word/7/
logical bit
print 1, word
1  format(13x 'word', o12.11)
   call bic( bitnum, word )
   print 2, word
2  format('after bic(2,word)', o12.11)
   call bis( bitnum, word )
   print 3, word
3  format('after bis(2,word)', o12.11)
   call setbit( bitnum, word, state )
   print 4, word
4  format('after setbit(2,word,0)', o12.11)
   print 5, bit(bitnum, word)
5  format('bit(2,word)', L )
   end
<output>
           word 00000000007
after bic(2,word) 00000000003
after bis(2,word) 00000000007
after setbit(2,word,0) 00000000003
bit(2,word) F
```

1.8

chdir: Change Default Directory

The function is called by:

INTEGER*4 chdir			
<i>n</i> = chdir(<i>dirname</i>)			
<i>dirname</i>	character	Input	Directory name
Return value	INTEGER*4	Output	<i>n</i> =0: OK, <i>n</i> >0: Error code

Example: `chdir`—change `cwd` to `MyDir`:

```
INTEGER*4  chdir, n
n =  chdir ( 'MyDir' )
if ( n .ne. 0 ) stop 'chdir: error'
end
```

See also: `chdir(2)`, `cd(1)`, and `gerror(3F)` to interpret error codes.

Path names can be no longer than `MAXPATHLEN` as defined in `<sys/param.h>`. They can be relative or absolute paths.

Use of this function can cause `inquire` by unit to fail.

Certain Fortran file operations reopen files by name. Using `chdir` while doing I/O can cause the runtime system to lose track of files created with relative path names, including the files that are created by open statements without file names.

1.9 `chmod`: Change the Mode of a File

The function is called by:

INTEGER*4 <code>chmod</code> <code>n = chmod(name, mode)</code>			
<i>name</i>	character	Input	Path name
<i>mode</i>	character	Input	Anything recognized by <code>chmod(1)</code> , such as <code>o-w</code> , <code>444</code> , etc.
Return value	INTEGER*4	Output	<i>n</i> = 0: OK; <i>n</i> >0: System error number

Example: `chmod`—add write permissions to `MyFile`:

```
character*18 name, mode
INTEGER*4  chmod, n
name = 'MyFile'
mode = '+w'
n =  chmod( name, mode )
if ( n .ne. 0 ) stop 'chmod: error'
end
```

See also: `chmod(1)`, and `gerror(3F)` to interpret error codes.

Path names cannot be longer than `MAXPATHLEN` as defined in `<sys/param.h>`. They can be relative or absolute paths.

1.10 date: Get Current Date as a Character String

Note – This routine is not “Year 2000 Safe” because it returns only a two-digit value for the year. Programs that compute differences between dates using the output of this routine may not work properly after 31 December, 1999. Programs using this `date()` routine will see a runtime warning message the first time the routine is called to alert the user. See `date_and_time()` as a possible alternate routine.

The subroutine is called by:

call date(c)			
c	CHARACTER*9	Output	Variable, array, array element, or character substring

The form of the returned string *c* is *dd-mmm-yy*, where *dd* is the day of the month as a 2-digit number, *mmm* is the month as a 3-letter abbreviation, and *yy* is the year as a 2-digit number (and is not year 2000 safe!).

Example: date:

```
demo% cat dat1.f
* dat1.f -- Get the date as a character string.
      character c*9
      call date ( c )
      write(*, "(' The date today is: ', A9 )" ) c
      end
demo% f95 dat1.f
demo% a.out
Computing time differences using the 2 digit year from subroutine
      date is not safe after year 2000.
      The date today is: 9-Jan-02
demo%
```

See also `idate()` and `date_and_time()`.

1.11 date_and_time: Get Date and Time

This is a Fortran 95 intrinsic routine, and is Year 2000 safe.

The `date_and_time` subroutine returns data from the real-time clock and the date. Local time is returned, as well as the difference between local time and Universal Coordinated Time (UTC) (also known as Greenwich Mean Time, GMT).

The `date_and_time()` subroutine is called by:

call date_and_time(date, time, zone, values)			
<i>date</i>	CHARACTER*8	Output	Date, in form CCYYMMDD, where CCYY is the four-digit year, MM the two-digit month, and DD the two-digit day of the month. For example: 19980709
<i>time</i>	CHARACTER*10	Output	The current time, in the form hhmmss.sss, where hh is the hour, mm minutes, and ss.sss seconds and milliseconds.
<i>zone</i>	CHARACTER*5	Output	The time difference with respect to UTC, expressed in hours and minutes, in the form hhmm
<i>values</i>	INTEGER*4 VALUES(8)	Output	An integer array of 8 elements described below.

The eight values returned in the `INTEGER*4 values` array are

VALUES(1)	The year, as a 4-digit integer. For example, 1998.
VALUES(2)	The month, as an integer from 1 to 12.
VALUES(3)	The day of the month, as an integer from 1 to 31.
VALUES(4)	The time difference, in minutes, with respect to UTC.
VALUES(5)	The hour of the day, as an integer from 1 to 23.
VALUES(6)	The minutes of the hour, as an integer from 1 to 59.
VALUES(7)	The seconds of the minute, as an integer from 0 to 60.
VALUES(8)	The milliseconds of the second, in range 0 to 999.

An example using `date_and_time`:

```
demo% cat dtm.f
      integer date_time(8)
      character*10 b(3)
      call date_and_time(b(1), b(2), b(3), date_time)
      print *, 'date_time    array values:'
      print *, 'year=', date_time(1)
      print *, 'month_of_year=', date_time(2)
      print *, 'day_of_month=', date_time(3)
      print *, 'time difference in minutes=', date_time(4)
      print *, 'hour of day=', date_time(5)
      print *, 'minutes of hour=', date_time(6)
      print *, 'seconds of minute=', date_time(7)
      print *, 'milliseconds of second=', date_time(8)
      print *, 'DATE=', b(1)
      print *, 'TIME=', b(2)
      print *, 'ZONE=', b(3)
      end
```

When run on a computer in California, USA on February 16, 2000, it generated the following output:

```
date_time    array values:
year= 2000
month_of_year= 2
day_of_month= 16
time difference in minutes= -420
hour of day= 11
minutes of hour= 49
seconds of minute= 29
milliseconds of second= 236
DATE=20000216
TIME=114929.236
ZONE=-0700
```

1.12 `dtime`, `etime`: Elapsed Execution Time

Both functions have return values of elapsed time (or -1.0 as error indicator). The time returned is in seconds.

Versions of `mtime` and `etime` used by Fortran 95 use the system's low resolution clock by default. The resolution is one hundredth of a second. However, if the program is run under the Sun OS™ operating system utility `ptime(1)`, (`/usr/proc/bin/ptime`), the high resolution clock is used.

1.12.1 `mtime`: Elapsed Time Since the Last `mtime` Call

For `mtime`, the elapsed time is:

- First call: elapsed time since start of execution
- Subsequent calls: elapsed time since the last call to `mtime`
- Single processor: time used by the CPU
- Multiple Processor: the sum of times for all the CPUs, which is not useful data; use `etime` instead.

Note – Calling `mtime` from within a parallelized loop gives non-deterministic results, since the elapsed time counter is global to all threads participating in the loop

The function is called by:

<i>e</i> = <code>mtime(tarray)</code>			
<i>tarray</i>	<code>real(2)</code>	Output	<i>e</i> = -1.0: Error: <i>tarray</i> values are undefined <i>e</i> ≠ -1.0: User time in <i>tarray</i> (1) if no error. System time in <i>tarray</i> (2) if no error
Return value	<code>real</code>	Output	<i>e</i> = -1.0: Error <i>e</i> ≠ -1.0: The sum of <i>tarray</i> (1) and <i>tarray</i> (2)

Example: `mtime()`, single processor:

```
real e, mtime, t(2)
print *, 'elapsed:', e, ', user:', t(1), ', sys:', t(2)
do i = 1, 10000
    k=k+1
end do
e = mtime( t )
print *, 'elapsed:', e, ', user:', t(1), ', sys:', t(2)
end
demo% f95 tdmtime.f
demo% a.out
elapsed: 0.0E+0 , user: 0.0E+0 , sys: 0.0E+0
elapsed: 0.03 , user: 0.01 , sys: 0.02
demo%
```

1.12.2 `etime`: Elapsed Time Since Start of Execution

For `etime`, the elapsed time is:

- Single Processor Execution—CPU time for the calling process
- Multiple Processor Execution—wallclock time while processing your program

The runtime library determines that a program is executing in a multiprocessor mode if either the `PARALLEL` or `OMP_NUM_THREADS` environment variables are defined to some integer value greater than 1.

The function is called by:

<code>e = etime(tarray)</code>			
<i>tarray</i>	<code>real(2)</code>	Output	<code>e = -1.0</code> : Error: <i>tarray</i> values are undefined. <code>e ≠ -1.0</code> : Single Processor: User time in <i>tarray</i> (1). System time in <i>tarray</i> (2) Multiple Processor: Wall clock time in <i>tarray</i> (1), 0.0 in <i>tarray</i> (2)
Return value	<code>real</code>	Output	<code>e = -1.0</code> : Error <code>e ≠ -1.0</code> : The sum of <i>tarray</i> (1) and <i>tarray</i> (2)

Take note that the initial call to `etime` will be inaccurate. It merely enables the system clock. Do not use the value returned by the initial call to `etime`.

Example: `etime()`, single processor:

```
real e, etime, t(2)
e = etime(t)          ! Startup etime - do not use result
do i = 1, 10000
  k=k+1
end do
e = etime( t )
print *, 'elapsed:', e, ', user:', t(1), ', sys:', t(2)
end
demo% f95 tetime.f
demo% a.out
elapsed: 0.02 , user: 0.01 , sys: 0.01
demo%
```

See also `times(2)`, and the *Fortran Programming Guide*.

1.13 `exit`: Terminate a Process and Set the Status

The subroutine is called by:

```
call exit( status )
```

<i>status</i>	INTEGER*4	Input
---------------	-----------	-------

Example: `exit()`:

```
...
if(dx .lt. 0.) call exit( 0 )
...
end
```

`exit` flushes and closes all the files in the process, and notifies the parent process if it is executing a `wait`.

The low-order 8 bits of *status* are available to the parent process. These 8 bits are shifted left 8 bits, and all other bits are zero. (Therefore, *status* should be in the range of 256 - 65280). This call will never return.

The C function `exit` can cause cleanup actions before the final system 'exit'.

Calling `exit` without an argument causes a compile-time warning message, and a zero will be automatically provided as an argument. See also: `exit(2)`, `fork(2)`, `fork(3F)`, `wait(2)`, `wait(3F)`.

1.14 fdate: Return Date and Time in an ASCII String

The subroutine or function is called by:

call fdate(<i>string</i>)		
<i>string</i>	character*24	Output

or:

CHARACTER fdate*24 <i>string</i> = fdate()		If used as a function, the calling routine must define the type and size of fdate.
Return value	character*24	
	Output	

Example 1: fdate as a subroutine:

```
character*24 string
call fdate( string )
write(*,*) string
end
```

Output:

```
Wed Aug 3 15:30:23 1994
```

Example 2: `fdate` as a function, same output:

```
character*24 fdate
write(*,*) fdate()
end
```

See also: `ctime(3)`, `time(3F)`, and `idate(3F)`.

1.15 `flush`: Flush Output to a Logical Unit

The function is called by:

INTEGER*4 <code>flush</code> <code>n = flush(lunit)</code>			
<i>lunit</i>	INTEGER*4	Input	Logical unit
Return value	INTEGER*4	Output	<i>n</i> = 0 no error <i>n</i> > 0 error number

The `flush` function flushes the contents of the buffer for the logical unit, `lunit`, to the associated file. This is most useful for logical units 0 and 6 when they are both associated with the console terminal. The function returns a positive error number if an error was encountered; zero otherwise.

See also `fclose(3S)`.

1.16 fork: Create a Copy of the Current Process

The function is called by:

INTEGER*4 fork <i>n</i> = fork()			
Return value	INTEGER*4	Output	<i>n</i> >0: <i>n</i> =Process ID of copy <i>n</i> <0, <i>n</i> =System error code

The `fork` function creates a copy of the calling process. The only distinction between the two processes is that the value returned to one of them, referred to as the *parent* process, will be the process ID of the copy. The copy is usually referred to as the *child* process. The value returned to the child process will be zero.

All logical units open for writing are flushed before the `fork` to avoid duplication of the contents of I/O buffers in the external files.

Example: `fork()`:

```
INTEGER*4 fork, pid
pid = fork()
if(pid.lt.0) stop 'fork error'
if(pid.gt.0) then
    print *, 'I am the parent'
else
    print *, 'I am the child'
endif
```

A corresponding `exec` routine has not been provided because there is no satisfactory way to retain open logical units across the `exec` routine. However, the usual function of `fork/exec` can be performed using `system(3F)`. See also: `fork(2)`, `wait(3F)`, `kill(3F)`, `system(3F)`, and `perror(3F)`.

1.17 fseek, ftell: Determine Position and Reposition a File

`fseek` and `ftell` are routines that permit repositioning of a file. `ftell` returns a file's current position as an offset of so many bytes from the beginning of the file. At some later point in the program, `fseek` can use this saved offset value to reposition the file to that same place for reading.

1.17.1 fseek: Reposition a File on a Logical Unit

The function is called by:

INTEGER*4 fseek <i>n = fseek(lunit, offset, from)</i>			
<i>lunit</i>	INTEGER*4	Input	Open logical unit
<i>offset</i>	INTEGER*4 <i>or</i> INTEGER*8	Input	Offset in bytes relative to position specified by <i>from</i>
	An INTEGER*8 offset value is required when compiled for a 64-bit environment, such as Solaris 7 or 8, with -xarch=v9. If a literal constant is supplied, it must be a 64-bit constant, for example: 100_8		
<i>from</i>	INTEGER*4	Input	0=Beginning of file 1=Current position 2=End of file
Return value	INTEGER*4	Output	<i>n</i> =0: OK; <i>n</i> >0: System error code

Note – On sequential files, following a call to `fseek` by an output operation (for example, `WRITE`) causes all data records following the `fseek` position to be deleted and replaced by the new data record (and an end-of-file mark). Rewriting a record in place can only be done with direct access files.

Example: `fseek()`—Reposition `MyFile` to two bytes from the beginning:

```
INTEGER*4 fseek, lunit/1/, offset/2/, from/0/, n
open( UNIT=lunit, FILE='MyFile' )
n = fseek( lunit, offset, from )
if ( n .gt. 0 ) stop 'fseek error'
end
```

Example: Same example in a 64-bit environment and compiled with `-xarch=v9`:

```
INTEGER*4 fseek, lunit/1/, from/0/, n
INTEGER*8 offset/2/
open( UNIT=lunit, FILE='MyFile' )
n = fseek( lunit, offset, from )
if ( n .gt. 0 ) stop 'fseek error'
end
```

1.17.2 `ftell`: Return Current Position of File

The function is called by:

INTEGER*4 ftell <i>n</i> = ftell(<i>lunit</i>)			
<i>lunit</i>	INTEGER*4	Input	Open logical unit
Return value	INTEGER*4 or INTEGER*8	Output	<i>n</i> >=0: <i>n</i> =Offset in bytes from start of file <i>n</i> <0: <i>n</i> =System error code
	An INTEGER*8 offset value is returned when compiling for a 64-bit environment with -xarch=v9. ftell and variables receiving this return value should be declared INTEGER*8.		

Example: `ftell()`:

```
INTEGER*4 ftell, lunit/1/, n
open( UNIT=lunit, FILE='MyFile' )
...
n = ftell( lunit )
if ( n .lt. 0 ) stop 'ftell error'
...
```

Example: Same example in a 64-bit environment and compiled with `-xarch=v9`:

```
INTEGER*4 lunit/1/
INTEGER*8 ftell, n
open( UNIT=lunit, FILE='MyFile' )
...
n = ftell( lunit )
if ( n .lt. 0 ) stop 'ftell error'
...
```

See also `fseek(3S)` and `perror(3F)`; also `fseeko64(3F)` `ftello64(3F)`.

1.18 `fseeko64`, `ftello64`: Determine Position and Reposition a Large File

`fseeko64` and `ftello64` are "large file" versions of `fseek` and `ftell`. They take and return `INTEGER*8` file position offsets. (A "large file" is larger than 2 Gigabytes and therefore a byte-position must be represented by a 64-bit integer.) Use these versions to determine and/or reposition large files.

1.18.1 `fseeko64`: Reposition a File on a Logical Unit

The function is called by:

<code>INTEGER fseeko64</code> <code>n = fseeko64(lunit, offset64, from)</code>			
<i>lunit</i>	<code>INTEGER*4</code>	Input	Open logical unit
<i>offset64</i>	<code>INTEGER*8</code>	Input	64-bit offset in bytes relative to position specified by <i>from</i>
<i>from</i>	<code>INTEGER*4</code>	Input	0=Beginning of file 1=Current position 2=End of file
Return value	<code>INTEGER*4</code>	Output	<i>n</i> =0: OK; <i>n</i> >0: System error code

Note – On sequential files, following a call to `fseeko64` by an output operation (for example, `WRITE`) causes all data records following the `fseek` position to be deleted and replaced by the new data record (and an end-of-file mark). Rewriting a record in place can only be done with direct access files.

Example: `fseeko64( )`—Reposition `MyFile` to two bytes from the beginning:

```

INTEGER fseeko64, lunit/1/, from/0/, n
INTEGER*8 offset/200/
open( UNIT=lunit, FILE='MyFile' )
n = fseeko64( lunit, offset, from )
if ( n .gt. 0 ) stop 'fseek error'
end

```

1.18.2 `ftello64`: Return Current Position of File

The function is called by:

INTEGER*8 <code>ftello64</code> <code>n = ftello64(lunit)</code>			
<i>lunit</i>	INTEGER*4	Input	Open logical unit
Return value	INTEGER*8	Output	$n \geq 0$: n =Offset in bytes from start of file $n < 0$: n =System error code

Example: `ftello64( )`:

```

INTEGER*8 ftello64, lunit/1/, n
open( UNIT=lunit, FILE='MyFile' )
...
n = ftello64( lunit )
if ( n .lt. 0 ) stop 'ftell error'
...

```

1.19 `getarg`, `iargc`: Get Command-Line Arguments

`getarg` and `iargc` access arguments on the command line (after expansion by the command-line preprocessor).

1.19.1 `getarg`: Get a Command-Line Argument

The subroutine is called by:

call <code>getarg(k, arg)</code>			
<i>k</i>	INTEGER*4	Input	Index of argument (0=first=command name)
<i>arg</i>	character*n	Output	<i>k</i> th argument
<i>n</i>	INTEGER*4	Size of <i>arg</i>	Large enough to hold longest argument

1.19.2 `iargc`: Get the Number of Command-Line Arguments

The function is called by:

<i>m</i> = <code>iargc()</code>			
Return value	INTEGER*4	Output	Number of arguments on the command line

Example: `iargc` and `getarg`, get argument count and each argument:

```
demo% cat yarg.f
      character argv*10
      INTEGER*4 i, iargc, n
      n = iargc()
      do 1 i = 1, n
        call getarg( i, argv )
1     write( *, '( i2, 1x, a )' ) i, argv
      end
demo% f95 yarg.f
demo% a.out *.f
1 first.f
2 yarg.f
```

See also `execve(2)` and `getenv(3F)`.

1.20 getc, fgetc: Get Next Character

`getc` and `fgetc` get the next character from the input stream. Do not mix calls to these routines with normal Fortran I/O on the same logical unit.

1.20.1 getc: Get Next Character From stdin

The function is called by:

INTEGER*4 getc <i>status</i> = getc(<i>char</i>)			
<i>char</i>	character	Output	Next character
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> =-1: End of file <i>status</i> >0: System error code or f77 I/O error code

Example: `getc` gets each character from the keyboard; note the Control-D (^D):

```
character char
INTEGER*4 getc, status
status = 0
do while ( status .eq. 0 )
    status = getc( char )
    write(*, '(i3, o4.3)') status, char
end do
end
```

After compiling, a sample run of the above source is:

```
demo% a.out
ab          Program reads letters typed in
0 141       Program outputs status and octal value of the characters entered
0 142       141 represents 'a', 142 is 'b'
0 012       012 represents the RETURN key
^D          terminated by a CONTROL-D.
-1 377     Next attempt to read returns CONTROL-D
demo%
```

For any logical unit, do not mix normal Fortran input with `getc()`.

1.20.2 fgetc: Get Next Character From Specified Logical Unit

The function is called by:

INTEGER*4 fgetc status = fgetc(lunit, char)			
lunit	INTEGER*4	Input	Logical unit
char	character	Output	Next character
Return value	INTEGER*4	Output	status=-1: End of File status>0: System error code or f77 I/O error code

Example: `fgetc` gets each character from `tfgetc.data`; note the linefeeds (Octal 012):

```
character char
INTEGER*4 fgetc, status
open( unit=1, file='tfgetc.data' )
status = 0
do while ( status .eq. 0 )
    status = fgetc( 1, char )
    write(*, '(i3, o4.3)') status, char
end do
end
```

After compiling, a sample run of the above source is:

```
demo% cat tfgetc.data
ab
yz
demo% a.out
0 141      'a' read
0 142      'b' read
0 012      linefeed read
0 171      'y' read
0 172      'z' read
0 012      linefeed read
-1 012     CONTROL-D read
demo%
```

For any logical unit, do not mix normal Fortran input with `fgetc()`.

See also: `getc(3S)`, `intro(2)`, and `perror(3F)`.

1.21 `getcwd`: Get Path of Current Working Directory

The function is called by:

<code>INTEGER*4 getcwd</code> <code>status = getcwd(dirname)</code>			
<i>dirname</i>	<code>character*n</code>	Output The path of the current directory is returned	Path name of the current working directory. <i>n</i> must be large enough for longest path name
Return value	<code>INTEGER*4</code>	Output	<code>status=0</code> : OK <code>status>0</code> : Error code

Example: `getcwd`:

```
INTEGER*4 getcwd, status
character*64 dirname
status = getcwd( dirname )
if ( status .ne. 0 ) stop 'getcwd: error'
write(*,*) dirname
end
```

See also: `chdir`(3F), `perror`(3F), and `getwd`(3).

Note: the path names cannot be longer than `MAXPATHLEN` as defined in `<sys/param.h>`.

1.22 `getenv`: Get Value of Environment Variables

The subroutine is called by:

call <code>getenv(<i>ename</i>, <i>eval</i>ue)</code>			
<i>ename</i>	character*n	Input	Name of the environment variable sought
<i>eval</i> ue	character*n	Output	Value of the environment variable found; blanks if not successful

The size of *ename* and *eval*ue must be large enough to hold their respective character strings.

The `getenv` subroutine searches the environment list for a string of the form *ename*=*eval*ue and returns the value in *eval*ue if such a string is present; otherwise, it fills *eval*ue with blanks.

Example: Use `getenv()` to print the value of `$SHELL`:

```
character*18 eval
```

```
call getenv( 'SHELL', eval )
```

```
write(*,*) "'", eval, "'"
```

```
end
```

See also: `execve`(2) and `environ`(5).

1.23 getfd: Get File Descriptor for External Unit Number

The function is called by:

INTEGER*4 getfd <i>fil</i> des = getfd(<i>unit</i> n)			
<i>unit</i> n	INTEGER*4	Input	External unit number
Return value	INTEGER*4 -or- INTEGER*8	Output	File descriptor if file is connected; -1 if file is not connected An INTEGER*8 result is returned when compiling for 64-bit environments

Example: getfd():

```
INTEGER*4 fildes, getfd, unitn/1/  
open( unitn, file='tgetfd.data' )  
fildes = getfd( unitn )  
if ( fildes .eq. -1 ) stop 'getfd: file not connected'  
write(*,*) 'file descriptor = ', fildes  
end
```

See also open(2).

1.24 `getfilep`: Get File Pointer for External Unit Number

The function is:

<code>irtn = c_read(getfilep(unitn), inbyte, 1)</code>			
<code>c_read</code>	C function	Input	User's own C function. See example.
<code>unitn</code>	INTEGER*4	Input	External unit number.
<code>getfilep</code>	INTEGER*4 -or- INTEGER*8	Return value	File pointer if the file is connected; -1 if the file is not connected. An INTEGER*8 value is returned when compiling for 64-bit environments

This function is used for mixing standard Fortran I/O with C I/O. Such a mix is nonportable, and is not guaranteed for subsequent releases of the operating system or Fortran. Use of this function is not recommended, and no direct interface is provided. You must create your own C routine to use the value returned by `getfilep`. A sample C routine is shown below.

Example: Fortran uses `getfilep` by passing it to a C function:

```
tgetfilepF.f:

character*1 inbyte
integer*4   c_read, getfilep, unitn / 5 /
external    getfilep
write(*,'(a,$)') 'What is the digit? '

irtn = c_read( getfilep( unitn ), inbyte, 1 )

write(*,9) inbyte
9 format('The digit read by C is ', a )
end
```

Sample C function actually using `getfilep`:

```
tgetfilepC.c:

#include <stdio.h>
int c_read_ ( fd, buf, nbytes, buf_len )
FILE **fd ;
char *buf ;
int *nbytes, buf_len ;
{
    return fread( buf, 1, *nbytes, *fd ) ;
}
```

A sample compile-build-run is:

```
demo 11% cc -c tgetfilepC.c
demo 12% f95 tgetfilepC.o tgetfilepF.f
demo 13% a.out
What is the digit? 3
The digit read by C is 3
demo 14%
```

For more information, read the chapter on the C-Fortran interface in the *Fortran Programming Guide*. See also `open(2)`.



1.25 getlog: Get User’s Login Name

The subroutine is called by:

call getlog(name)			
name	character*n	Output	User’s login name, or all blanks if the process is running detached from a terminal. <i>n</i> should be large enough to hold the longest name.

Example: getlog:

```
character*18 name
call getlog( name )
write(*,*) "'", name, "'"
end
```

See also getlogin(3).

1.26 getpid: Get Process ID

The function is called by:

```
INTEGER*4 getpid
```

```
pid = getpid()
```

Return value	INTEGER*4	Output	Process ID of the current process
--------------	-----------	--------	-----------------------------------

Example: getpid:

```
INTEGER*4 getpid, pid
pid = getpid()
write(*,*) 'process id = ', pid
end
```

See also getpid(2).

1.27 getuid, getgid: Get User or Group ID of Process

getuid and getgid get the user or group ID of the process, respectively.

1.27.1 `getuid`: Get User ID of the Process

The function is called by:

INTEGER*4 <code>getuid</code> <code>uid = getuid()</code>			
Return value	INTEGER*4	Output	User ID of the process

1.27.2 `getgid`: Get Group ID of the Process

The function is called by:

INTEGER*4 <code>getgid</code> <code>gid = getgid()</code>			
Return value	INTEGER*4	Output	Group ID of the process

Example: `getuid()` and `getpid()`:

```
INTEGER*4 getuid, getgid, gid, uid
uid = getuid()
gid = getgid()
write(*,*) uid, gid
end
```

See also: `getuid(2)`.

1.28 hostnm: Get Name of Current Host

The function is called by:

INTEGER*4 hostnm <i>status</i> = hostnm(<i>name</i>)			
<i>name</i>	character* <i>n</i>	Output	Name of current host system. <i>n</i> must be large enough to hold the host name.
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> >0: Error

Example: hostnm():

```
INTEGER*4 hostnm, status
character*8 name
status = hostnm( name )
write(*,*) 'host name = ', name, ' '
end
```

See also gethostname(2).

1.29 idate: Return Current Date

idate puts the current system date into one integer array: day, month, and year.

The subroutine is called by:

call idate(<i>iarray</i>) <i>Standard Version</i>			
<i>iarray</i>	INTEGER*4	Output	Three-element array: day, month, year.

Example: `idate` (standard version):

```
demo% cat tidate.f
      INTEGER*4 iarray(3)
      call idate( iarray )
      write(*, "(' The date is: ',3i5)" ) iarray
end
demo% f95 tidate.f
demo% a.out
      The date is: 10 8 1998
demo%
```

1.30 `ieee_flags`, `ieee_handler`, `sigfpe`: IEEE Arithmetic

These subprograms provide modes and status required to fully exploit ANSI/IEEE Standard 754-1985 arithmetic in a Fortran program. They correspond closely to the functions `ieee_flags(3M)`, `ieee_handler(3M)`, and `sigfpe(3)`.

Here is a summary:

TABLE 1-5 IEEE Arithmetic Support Routines

ieeeer = ieee_flags(action,mode,in,out)		
ieeeer = ieee_handler(action,exception,hdl)		
ieeeer = sigfpe(code, hdl)		
action	character	Input
code	sigfpe_code_type	Input
mode	character	Input
in	character	Input
exception	character	Input
hdl	sigfpe_handler_type	Input
out	character	Output
Return value	INTEGER*4	Output

See the Sun *Numerical Computation Guide* for details on how these functions can be used strategically.

If you use `sigfpe`, you must do your own setting of the corresponding trap-enable-mask bits in the floating-point status register. The details are in the SPARC architecture manual. The `libm` function `ieee_handler` sets these trap-enable-mask bits for you.

The character keywords accepted for *mode* and *exception* depend on the value of *action*.

TABLE 1-6 `ieee_flags(action,mode,in,out)` Parameters and Actions

<i>action</i> = 'clearall'	<i>mode</i> , <i>in</i> , <i>out</i> , unused; returns 0		
<i>action</i> = 'clear'	<i>mode</i> = 'direction'		
clear <i>mode</i> , <i>in</i> <i>out</i> is unused; returns 0	<i>mode</i> = 'exception'	<i>in</i> = 'inexact'	or
		'division'	or
		'underflow'	or
		'overflow'	or
		'invalid'	or
		'all'	or
		'common'	
<i>action</i> = 'set'	<i>mode</i> = 'direction'	<i>in</i> = 'nearest'	or
set floating-point <i>mode</i> , <i>in</i> <i>out</i> is unused; returns 0		'tozero'	or
		'positive'	or
		'negative'	
	<i>mode</i> = 'exception'	<i>in</i> = 'inexact'	or
		'division'	or
		'underflow'	or
		'overflow'	or
		'invalid'	or
		'all'	or
		'common'	
<i>action</i> = 'get'	<i>mode</i> = 'direction'	<i>out</i> = 'nearest'	or
test <i>mode</i> settings <i>in</i> , <i>out</i> may be blank or one of the settings to test returns the current setting depending on <i>mode</i> , or 'not available' The function returns 0 or the current exception flags if <i>mode</i> = 'exception'		'tozero'	or
		'positive'	or
		'negative'	
	<i>mode</i> = 'exception'	<i>out</i> = 'inexact'	or
		'division'	or
		'underflow'	or
		'overflow'	or
		'invalid'	or
		'all'	or
		'common'	

TABLE 1-7 `ieee_handler(action,in,out)` Parameters

<i>action</i> = 'clear'	<i>in</i> = 'inexact'	<i>or</i>
clear user exception handing of <i>in</i> ; <i>out</i> is unused	'division'	<i>or</i>
	'underflow'	<i>or</i>
	'overflow'	<i>or</i>
	'invalid'	<i>or</i>
	'all'	<i>or</i>
	'common'	
<i>action</i> = 'set'	<i>in</i> = 'inexact'	<i>or</i>
set user exception handing of <i>in</i> ; <i>out</i> is address of handler routine, or SIGFPE_DEFAULT, or SIGFPE_ABORT, or SIGFPE_IGNORE defined in floating point.h	'division'	<i>or</i>
	'underflow'	<i>or</i>
	'overflow'	<i>or</i>
	'invalid'	<i>or</i>
	'all'	<i>or</i>
	'common'	

Example 1: Set rounding direction to round toward zero, unless the hardware does not support directed rounding modes:

```
INTEGER*4 ieeeer
character*1 mode, out, in
ieeeer = ieee_flags( 'set', 'direction', 'tozero', out )
```

Example 2: Clear rounding direction to default (round toward nearest):

```
character*1 out, in
ieeeer = ieee_flags('clear','direction', in, out )
```

Example 3: Clear all accrued exception-occurred bits:

```
character*18 out
ieeeer = ieee_flags( 'clear', 'exception', 'all', out )
```

Example 4: Detect overflow exception as follows:

```
character*18 out
ieeeer = ieee_flags( 'get', 'exception', 'overflow', out )
if (out .eq. 'overflow' ) stop 'overflow'
```

The above code sets out to overflow and ieeeer to 25 (this value is platform dependent). Similar coding detects exceptions, such as invalid or inexact.

Example 5: hand1.f, write and use a signal handler:

```
external hand
real r / 14.2 /, s / 0.0 /
i = ieee_handler( 'set', 'division', hand )
t = r/s
end

INTEGER*4 function hand ( sig, sip, uap )
INTEGER*4 sig, address
structure /fault/
    INTEGER*4 address
end structure
structure /siginfo/
    INTEGER*4 si_signo
    INTEGER*4 si_code
    INTEGER*4 si_errno
    record /fault/ fault
end structure
record /siginfo/ sip
address = sip.fault.address
write (*,10) address
10  format('Exception at hex address ', z8 )
end
```

Change the declarations for address and function hand to INTEGER*8 to enable Example 5 in a 64-bit, SPARC V9 environment (-xarch=v9)

See the *Numerical Computation Guide*. See also: floatingpoint(3), signal(3), sigfpe(3), floatingpoint(3F), ieee_flags(3M), and ieee_handler(3M).

1.30.1 floatingpoint.h: Fortran IEEE Definitions

The header file floatingpoint.h defines constants and types used to implement standard floating-point according to ANSI/IEEE Std 754-1985.

Include the file in a Fortran 95 source program as follows:

```
#include "floatingpoint.h"
```

Use of this include file requires preprocessing prior to Fortran compilation. The source file referencing this include file will automatically be preprocessed if the name has a `.F`, `.F90` or `.F95` extension.

IEEE Rounding Mode:

<code>fp_direction_type</code>	The type of the IEEE rounding direction mode. The order of enumeration varies according to hardware.
--------------------------------	--

SIGFPE Handling:

<code>sigfpe_code_type</code>	The type of a SIGFPE code.
<code>sigfpe_handler_type</code>	The type of a user-definable SIGFPE exception handler called to handle a particular SIGFPE code.
<code>SIGFPE_DEFAULT</code>	A macro indicating default SIGFPE exception handling: IEEE exceptions to continue with a default result and to abort for other SIGFPE codes.
<code>SIGFPE_IGNORE</code>	A macro indicating an alternate SIGFPE exception handling, namely to ignore and continue execution.
<code>SIGFPE_ABORT</code>	A macro indicating an alternate SIGFPE exception handling, namely to abort with a core dump.

IEEE Exception Handling:

<code>N_IEEE_EXCEPTION</code>	The number of distinct IEEE floating-point exceptions.
<code>fp_exception_type</code>	The type of the <code>N_IEEE_EXCEPTION</code> exceptions. Each exception is given a bit number.
<code>fp_exception_field_type</code>	The type intended to hold at least <code>N_IEEE_EXCEPTION</code> bits corresponding to the IEEE exceptions numbered by <code>fp_exception_type</code> . Thus, <code>fp_inexact</code> corresponds to the least significant bit and <code>fp_invalid</code> to the fifth least significant bit. Some operations can set more than one exception.

IEEE Classification:

<code>fp_class_type</code>	A list of the classes of IEEE floating-point values and symbols.
----------------------------	--

Refer to the *Numerical Computation Guide*. See also `ieee_environment(3F)`.

1.31 `index`, `rindex`, `lnblnk`: Index or Length of Substring

These functions search through a character string:

<code>index(<i>a1</i>, <i>a2</i>)</code>	Index of first occurrence of string <i>a2</i> in string <i>a1</i>
<code>rindex(<i>a1</i>, <i>a2</i>)</code>	Index of last occurrence of string <i>a2</i> in string <i>a1</i>
<code>lnblnk(<i>a1</i>)</code>	Index of last nonblank in string <i>a1</i>

`index` has the following forms:

1.31.1 `index`: First Occurrence of a Substring in a String

The `index` is an intrinsic function called by:

<code>n = index(<i>a1</i>, <i>a2</i>)</code>			
<i>a1</i>	character	Input	Main string
<i>a2</i>	character	Input	Substring
Return value	INTEGER	Output	<i>n</i> >0: Index of first occurrence of <i>a2</i> in <i>a1</i> <i>n</i> =0: <i>a2</i> does not occur in <i>a1</i> .

If declared `INTEGER*8`, `index()` will return an `INTEGER*8` value when compiled for a 64-bit environment and character variable *a1* is a very large character string (greater than 2 Gigabytes).

1.31.2 rindex: Last Occurrence of a Substring in a String

The function is called by:

INTEGER*4 rindex <i>n</i> = rindex(<i>a1</i> , <i>a2</i>)			
<i>a1</i>	character	Input	Main string
<i>a2</i>	character	Input	Substring
Return value	INTEGER*4 or INTEGER*8	Output	<i>n</i> >0: Index of last occurrence of <i>a2</i> in <i>a1</i> <i>n</i> =0: <i>a2</i> does not occur in <i>a1</i> INTEGER*8 returned in 64-bit environments

1.31.3 lnblnk: Last Nonblank in a String

The function is called by:

<i>n</i> = lnblnk(<i>a1</i>)			
<i>a1</i>	character	Input	String
Return value	INTEGER*4 or INTEGER*8	Output	<i>n</i> >0: Index of last nonblank in <i>a1</i> <i>n</i> =0: <i>a1</i> is all nonblank INTEGER*8 returned in 64-bit environments

Example: index(), rindex(), lnblnk():

```
*
                                123456789012345678901
character s*24 / 'abcPDQxyz...abcPDQxyz' /
INTEGER*4 declen, index, first, last, len, lnblnk, rindex
declen = len( s )
first = index( s, 'abc' )
last = rindex( s, 'abc' )
lastnb = lnblnk( s )
write(*,*) declen, lastnb
write(*,*) first, last
end
demo% f95 tindex.f
demo% a.out
24 21      <- declen is 24 because intrinsic len() returns the declared length of s
1 13
```

Note – Programs compiled to run in a 64-bit environment must declare `index`, `rindex` and `lnblnk` (and their receiving variables) `INTEGER*8` to handle very large character strings.

1.32 `inmax`: Return Maximum Positive Integer

The function is called by:

<code>m = inmax()</code>			
Return value	<code>INTEGER*4</code>	Output	The maximum positive integer

Example: `inmax`:

```
INTEGER*4 inmax, m
m = inmax()
write(*,*) m
end
demo% f95 tinmax.f
demo% a.out
      2147483647
demo%
```

See also `libm_single(3F)` and `libm_double(3F)`. See also the non-standard FORTRAN 77 intrinsic function `ephuge( )` described in Chapter 3.

1.33 `itime`: Current Time

`itime` puts the current system time into an integer array: hour, minute, and second. The subroutine is called by:

call <code>itime(iarray)</code>			
<i>iarray</i>	INTEGER*4	Output	3-element array: <i>iarray</i> (1) = hour <i>iarray</i> (2) = minute <i>iarray</i> (3) = second

Example: `itime`:

```
demo% cat titime.f
      INTEGER*4 iarray(3)
      call itime( iarray )
      write ( *, "( ' The time is: ',3i5)" ) iarray
      end
demo% f95 titime.f
demo% a.out
      The time is: 15 42 35
```

See also `time(3F)`, `ctime(3F)`, and `fdate(3F)`.

1.34 `kill`: Send a Signal to a Process

The function is called by:

<code>status = kill(pid, signum)</code>			
<i>pid</i>	INTEGER*4	Input	Process ID of one of the user's processes
<i>signum</i>	INTEGER*4	Input	Valid signal number. See <i>signal(3)</i> .
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> >0: Error code

Example (fragment): Send a message using `kill()`:

```
INTEGER*4 kill, pid, signum
*
...
status = kill( pid, signum )
if ( status .ne. 0 ) stop 'kill: error'
write(*,*) 'Sent signal ', signum, ' to process ', pid
end
```

The function sends signal *signum*, and integer signal number, to the process *pid*. Valid signal numbers are listed in the C include file `/usr/include/sys/signal.h`

See also: `kill(2)`, `signal(3)`, `signal(3F)`, `fork(3F)`, and `perror(3F)`.

1.35 `link`, `symlink`: Make a Link to an Existing File

`link` creates a link to an existing file. `symlink` creates a symbolic link to an existing file.

The functions are called by:

<code>status = link(name1, name2)</code>			
INTEGER*4 <code>symlink</code> <code>status = symlink(name1, name2)</code>			
<i>name1</i>	character*n	Input	Path name of an existing file
<i>name2</i>	character*n	Input	Path name to be linked to the file, <i>name1</i> . <i>name2</i> must not already exist.
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> >0: System error code

1.35.1 link: Create a Link to an Existing File

Example 1: link: Create a link named data1 to the file, tlink.db.data.1:

```
demo% cat tlink.f
      character*34 name1/'tlink.db.data.1'/, name2/'data1'/
      integer*4 link, status
      status = link( name1, name2 )
      if ( status .ne. 0 ) stop 'link: error'
      end
demo% f95 tlink.f
demo% ls -l data1
data1 not found
demo% a.out
demo% ls -l data1
-rw-rw-r-- 2 generic 2 Aug 11 08:50 data1
demo%
```

1.35.2 symlink: Create a Symbolic Link to an Existing File

Example 2: symlink: Create a symbolic link named data1 to the file, tlink.db.data.1:

```
demo% cat tsymlnk.f
      character*34 name1/'tlink.db.data.1'/, name2/'data1'/
      INTEGER*4 status, symlnk
      status = symlnk( name1, name2 )
      if ( status .ne. 0 ) stop 'symlnk: error'
      end
demo% f95 tsymlnk.f
demo% ls -l data1
data1 not found
demo% a.out
demo% ls -l data1
lrwxrwxrwx 1 generic 15 Aug 11 11:09 data1 -> tlink.db.data.1
demo%
```

See also: link(2), symlink(2), perror(3F), and unlink(3F).

Note: the path names cannot be longer than MAXPATHLEN as defined in <sys/param.h>.

1.36 `loc`: Return the Address of an Object

This intrinsic function is called by:

$k = \text{loc}(\text{arg})$			
arg	Any type	Input	Variable or array
Return value	INTEGER*4 -or- INTEGER*8	Output	Address of arg
	Returns an INTEGER*8 pointer when compiled to run in a 64-bit environment with -xarch=v9. See Note below.		

Example: `loc`:

```
INTEGER*4 k, loc
real arg / 9.0 /
k = loc( arg )
write(*,*) k
end
```

Note – Programs compiled to run in a 64-bit environment should declare `INTEGER*8` the variable receiving output from the `loc( )` function.

1.37 `long`, `short`: Integer Object Conversion

`long` and `short` handle integer object conversions between `INTEGER*4` and `INTEGER*2`, and is especially useful in subprogram call lists.

1.37.1 long: Convert a Short Integer to a Long Integer

The function is called by:

call <i>ExpecLong</i> (long(<i>int2</i>))		
<i>int2</i>	INTEGER*2	Input
Return value	INTEGER*4	Output

1.37.2 short: Convert a Long Integer to a Short Integer

The function is:

INTEGER*2 short call <i>ExpecShort</i> (short(<i>int4</i>))		
<i>int4</i>	INTEGER*4	Input
Return value	INTEGER*2	Output

Example (fragment): long() and short():

```
integer*4 int4/8/, long
integer*2 int2/8/, short
call ExpecLong( long(int2) )
call ExpecShort( short(int4) )
...
end
```

ExpecLong is some subroutine called by the user program that expects a *long* (INTEGER*4) integer argument. Similarly, *ExpecShort* expects a *short* (INTEGER*2) integer argument.

long is useful if constants are used in calls to library routines and the code is compiled with the `-i2` option.

short is useful in similar context when an otherwise long object must be passed as a short integer. Passing an integer to *short* that is too large in magnitude does not cause an error, but will result in unexpected behavior.

1.38 longjmp, issetjmp: Return to Location Set by issetjmp

issetjmp sets a location for longjmp; longjmp returns to that location.

1.38.1 issetjmp: Set the Location for longjmp

This intrinsic function is called by:

<i>ival</i> = issetjmp(<i>env</i>)			
<i>env</i>	INTEGER*4	Output	<i>env</i> is a 12-element integer array. In 64-bit environments it must be declared INTEGER*8
Return value	INTEGER*4	Output	<i>ival</i> = 0 if issetjmp is called explicitly <i>ival</i> ≠ 0 if issetjmp is called through longjmp

1.38.2 longjmp: Return to the Location Set by issetjmp

The subroutine is called by:

call longjmp(<i>env</i> , <i>ival</i>)			
<i>env</i>	INTEGER*4	Input	<i>env</i> is the 12-word integer array initialized by issetjmp. In 64-bit environments it must be declared INTEGER*8
<i>ival</i>	INTEGER*4	Output	<i>ival</i> = 0 if issetjmp is called explicitly <i>ival</i> ≠ 0 if issetjmp is called through longjmp

1.38.3 Description

The `isetjmp` and `longjmp` routines are used to deal with errors and interrupts encountered in a low-level routine of a program. They are f95 intrinsics.

These routines should be used only as a last resort. They require discipline, and are not portable. Read the man page, `setjmp(3V)`, for bugs and other details.

`isetjmp` saves the stack environment in *env*. It also saves the register environment.

`longjmp` restores the environment saved by the last call to `isetjmp`, and returns in such a way that execution continues as if the call to `isetjmp` had just returned the value *ival*.

The integer expression *ival* returned from `isetjmp` is zero if `longjmp` is not called, and nonzero if `longjmp` is called.

Example: Code fragment using `isetjmp` and `longjmp`:

```
INTEGER*4 env(12)
common /jmpblk/ env
j = isetjmp( env )
if ( j .eq. 0 ) then
  call sbrtnA
else
  call error_processor
end if
end
subroutine sbrtnA
  INTEGER*4 env(12)
  common /jmpblk/ env
  call longjmp( env, ival )
  return
end
```

1.38.4 Restrictions

- You must invoke `isetjmp` before calling `longjmp`.
- The *env* integer array argument to `isetjmp` and `longjmp` must be at least 12 elements long.
- You must pass the *env* variable from the routine that calls `isetjmp` to the routine that calls `longjmp`, either by common or as an argument.
- `longjmp` attempts to clean up the stack. `longjmp` must be called from a lower call-level than `isetjmp`.

- Passing `isetjmp` as an argument that is a procedure name does not work.

See `setjmp(3V)`.

1.39 `malloc`, `malloc64`, `realloc`, `free`: Allocate/Reallocate/Deallocate Memory

The functions `malloc()`, `malloc64()`, and `realloc()` allocate blocks of memory and return the starting address of the block. The return value can be used to set an `INTEGER` or Cray-style `POINTER` variable. `realloc()` reallocates an existing memory block with a new size. `free()` deallocates memory blocks allocated by `malloc()`, `malloc64()`, or `realloc()`.

Note – These routines are implemented as intrinsic functions in f95, but were external functions in f77. They should not appear on type declarations in Fortran 95 programs, or on `EXTERNAL` statements unless you wish to use your own versions. The `realloc()` routine is only implemented for f95.

Standard-conforming Fortran 95 programs should use `ALLOCATE` and `DEALLOCATE` statements on `ALLOCATABLE` arrays to perform dynamic memory management, and not make direct calls to `malloc/realloc/free`.

Legacy Fortran 77 programs could use `malloc()/malloc64()` to assign values to Cray-style `POINTER` variables, which have the same data representation as `INTEGER` variables. Cray-style `POINTER` variables are implemented in f95 to support portability from Fortran 77.

1.39.1 Allocate Memory: malloc, malloc64

The `malloc()` function is called by:

$k = \text{malloc}(n)$			
n	INTEGER	Input	Number of bytes of memory
Return value	INTEGER (Cray POINTER)	Output	$k > 0$: k = address of the start of the block of memory allocated $k = 0$: Error
	An INTEGER*8 pointer value is returned when compiled for a 64-bit environment with <code>-xarch=v9</code> . See Note below.		

Note – This function is intrinsic in Fortran 95 and was external in Fortran 77. Fortran 77 programs compiled to run in 64-bit environments would declare the `malloc()` function and the variables receiving its output as `INTEGER*8`. The function `malloc64(3F)` was provided to make programs portable between 32-bit and 64-bit environments.

$k = \text{malloc64}(n)$			
n	INTEGER*8	Input	Number of bytes of memory
Return value	INTEGER*8 (Cray POINTER)	Output	$k > 0$: k =address of <i>the</i> start of the block of memory allocated $k = 0$: Error

These functions allocate an area of memory and return the address of the start of that area. (In a 64-bit environment, this returned byte address may be outside the `INTEGER*4` numerical range—the receiving variables must be declared `INTEGER*8` to avoid truncation of the memory address.) The region of memory is not initialized in any way, and it should not be assumed to be preset to anything, especially zero!

Example: Code fragment using `malloc()`:

```

parameter (NX=1000)
integer ( p2X, X )
real*4 X(1)
...
p2X = malloc( NX*4 )
if ( p2X .eq. 0 ) stop 'malloc: cannot allocate'
do 11 i=1,NX
11  X(i) = 0.
...
end

```

In the above example, we acquire 4,000 bytes of memory, pointed to by `p2X`, and initialize it to zero.

1.39.2 Reallocate Memory: `realloc`

The `realloc()` f95 intrinsic function is called by:

$k = \text{realloc}(\text{ptr}, n)$			
<i>ptr</i>	INTEGER	Input	Pointer to existing memory block. (Value returned from a previous <code>malloc()</code> or <code>realloc()</code> call).
<i>n</i>	INTEGER	Input	Requested new size of block, in bytes.
Return value	INTEGER (<i>Cray</i> POINTER)	Output	$k > 0$: k =address of <i>the</i> start of the new block of memory allocated $k = 0$: Error
	An <code>INTEGER*8</code> pointer value is returned when compiled for a 64-bit environment with <code>-xarch=v9</code> . See Note below.		

The `realloc()` function changes the size of the memory block pointed to by `ptr` to `n` bytes and returns a pointer to the (possibly moved) new block. The contents of the memory block will be unchanged up to the lesser of the new and old sizes.

If `ptr` is zero, `realloc()` behaves the same as `malloc()` and allocates a new memory block of size `n` bytes.

If `n` is zero and `ptr` is not zero, the memory block pointed to is made available for further allocation and is returned to the system only upon termination of the application.

Example: Using `malloc()` and `realloc()` and Cray-style `POINTER` variables:

```
PARAMETER (nsize=100001)
POINTER (p2space, space)
REAL*4 space(1)

p2space = malloc(4*nsize)
if(p2space .eq. 0) STOP 'malloc: cannot allocate space'
...
p2space = realloc(p2space, 9*4*nsize)
if(p2space .eq. 0) STOP 'realloc: cannot reallocate space'
...
CALL free(p2space)
...
```

Note that `realloc()` is only implemented for f95.

1.39.3 `free`: Deallocate Memory Allocated by Malloc

The subroutine is called by:

call free (<i>ptr</i>)		
<i>ptr</i>	Cray <code>POINTER</code>	Input

`free` deallocates a region of memory previously allocated by `malloc` and `realloc()`. The region of memory is returned to the memory manager; it is no longer available to the user's program.

Example: `free()`:

```
real x
pointer ( ptr, x )
ptr = malloc ( 10000 )
call free ( ptr )
end
```

1.40 mvbits: Move a Bit Field

The subroutine is called by:

call mvbits(src, ini1, nbits, des, ini2)			
src	INTEGER*4	Input	Source
ini1	INTEGER*4	Input	Initial bit position in the source
nbits	INTEGER*4	Input	Number of bits to move
des	INTEGER*4	Output	Destination
ini2	INTEGER*4	Input	Initial bit position in the destination

Example: mvbits:

```
demo% cat mvb1.f
* mvb1.f -- From src, initial bit 0, move 3 bits to des, initial
*           bit 3.
*   src     des
* 543210 543210 ← Bit numbers
* 000111 000001 ← Values before move
* 000111 111001 ← Values after move
      INTEGER*4 src, ini1, nbits, des, ini2
      data src, ini1, nbits, des, ini2
      1 / 7, 0, 3, 1, 3 /
      call mvbits ( src, ini1, nbits, des, ini2 )
      write (*,"(5o3)") src, ini1, nbits, des, ini2
      end
demo% f95 mvb1.f
demo% a.out
  7 0 3 71 3
demo%
```

Note the following:

- Bits are numbered 0 to 31, from least significant to most significant.
- mvbits changes only bits *ini2* through *ini2+nbits-1* of the *des* location, and no bits of the *src* location.
- The restrictions are:
 - $ini1 + nbits \geq 32$
 - $ini2 + nbits \leq 32$

1.41 perror, gerror, ierrno: Get System Error Messages

These routines perform the following functions:

perror	Print a message to Fortran logical unit 0, stderr.
gerror	Get a system error message (of the last detected system error)
ierrno	Get the error number of the last detected system error.

1.41.1 perror: Print Message to Logical Unit 0, stderr

The subroutine is called by:

call perror(<i>string</i>)			
<i>string</i>	character*n	Input	The message. It is written preceding the standard error message for the last detected system error.

Example 1:

```
call perror( "file is for formatted I/O" )
```

1.41.2 gerror: Get Message for Last Detected System Error

The subroutine or function is called by:

call gerror(<i>string</i>)			
<i>string</i>	character*n	Output	Message for the last detected system error

Example 2: `gerror()` as a subroutine:

```
character string*30
...
call gerror ( string )
write(*,*) string
```

Example 3: `gerror()` as a function; *string* not used:

```
character gerror*30, z*30
...
z = gerror( )
write(*,*) z
```

1.41.3 `ierrno`: Get Number for Last Detected System Error

The function is called by:

<code>n = ierrno()</code>			
Return value	INTEGER*4	Output	Number of last detected system error

This number is updated only when an error actually occurs. Most routines and I/O statements that might generate such errors return an error code after the call; that value is a more reliable indicator of what caused the error condition.

Example 4: `ierrno()`:

```
INTEGER*4 ierrno, n
...
n = ierrno()
write(*,*) n
```

See also `intro(2)` and `perror(3)`.

Note:

- *string* in the call to `perror` cannot be longer than 127 characters.
- The length of the string returned by `gerror` is determined by the calling program.

- Runtime I/O error codes for f95 are listed in the *Fortran User's Guide*.

1.42 `putc, fputc`: Write a Character to a Logical Unit

`putc` writes to logical unit 6, normally the control terminal output.

`fputc` writes to a logical unit.

These functions write a character to the file associated with a Fortran logical unit bypassing normal Fortran I/O.

Do not mix normal Fortran output with output by these functions on the same unit.

Note that to write any of the special `\` escape characters, such as `'\n'` newline, requires compiling with `-f77=backslash` FORTRAN 77 compatibility option.

1.42.1 `putc`: Write to Logical Unit 6

The function is called by:

<pre>INTEGER*4 putc status = putc(char)</pre>			
<i>char</i>	character	Input	The character to write to the unit
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> >0: System error code

Example: `putc()`:

```
character char, s*10 / 'OK by putc' /
INTEGER*4 putc, status
do i = 1, 10
  char = s(i:i)
  status = putc( char )
end do
status = putc( '\n' )
end
demo% f95 -f77=backslash tputc.f
demo% a.out
OK by putc
demo%
```

1.42.2 `fputc`: Write to Specified Logical Unit

The function is called by:

INTEGER*4 <code>fputc</code> <code>status = fputc(lunit, char)</code>			
<i>lunit</i>	INTEGER*4	Input	The unit to write to
<i>char</i>	character	Input	The character to write to the unit
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> >0: System error code

Example: `fputc()`:

```
character char, s*11 / 'OK by fputc' /
INTEGER*4 fputc, status
open( 1, file='tfputc.data')
do i = 1, 11
    char = s(i:i)
    status = fputc( 1, char )
end do
status = fputc( 1, '\n' )
end
demo% f95 -f77=backslash tfputc.f
demo% a.out
demo% cat tfputc.data
OK by fputc
demo%
```

See also `putc(3S)`, `intro(2)`, and `perror(3F)`.

1.43 `qsort`, `qsort64`: Sort the Elements of a One-Dimensional Array

The subroutine is called by:

```
call qsort( array, len, isize, compar )      (Continued)
call qsort64( array, len8, isize8, compar )
```

<i>array</i>	array	Input	Contains the elements to be sorted
<i>len</i>	INTEGER*4	Input	Number of elements in the array.
<i>len8</i>	INTEGER*8	Input	Number of elements in the array

<i>isize</i>	INTEGER*4	Input	Size of an element, typically: 4 for integer or real 8 for double precision or complex 16 for double complex Length of character object for character arrays
<i>isize8</i>	INTEGER*8	Input	Size of an element, typically: 4_8 for integer or real 8_8 for double precision or complex 16_8 for double complex Length of character object for character arrays
<i>compar</i>	function name	Input	Name of a user-supplied INTEGER*2 function. Determines sorting order: <code>compar(arg1,arg2)</code>

Use `qsort64` in 64-bit environments with arrays larger than 2 Gbytes. Be sure to specify the array length, *len8*, and the element size, *isize8*, as INTEGER*8 data. Use the Fortran 95 style constants to explicitly specify INTEGER*8 constants.

The `compar(arg1, arg2)` arguments are elements of *array*, returning:

Negative	If <i>arg1</i> is considered to precede <i>arg2</i>
Zero	If <i>arg1</i> is equivalent to <i>arg2</i>
Positive	If <i>arg1</i> is considered to follow <i>arg2</i>

For example:

```
demo% cat tqsort.f
      external compar
      integer*2 compar
      INTEGER*4 array(10)/5,1,9,0,8,7,3,4,6,2/,len/10/,
1      isize/4/
      call qsort( array, len, isize, compar )
      write(*,'(10i3)') array
      end
      integer*2 function compar( a, b )
      INTEGER*4 a, b
      if ( a .lt. b ) compar = -1
      if ( a .eq. b ) compar = 0
      if ( a .gt. b ) compar = 1
      return
      end
demo% f95 tqsort.f
demo% a.out
      0 1 2 3 4 5 6 7 8 9
```

1.44 ran: Generate a Random Number Between 0 and 1

Repeated calls to ran generate a sequence of random numbers with a uniform distribution. See lcrans(3m).

<i>r</i> = ran(<i>i</i>)			
<i>i</i>	INTEGER*4	Input	Variable or array element
<i>r</i>	REAL	Output	Variable or array element

Example: ran:

```
demo% cat ran1.f
* ran1.f -- Generate random numbers.
      INTEGER*4 i, n
      real r(10)
      i = 760013
      do n = 1, 10
        r(n) = ran ( i )
      end do
      write ( *, "( 5 f11.6 )" ) r
      end
demo% f95 ran1.f
demo% a.out
      0.222058 0.299851 0.390777 0.607055 0.653188
      0.060174 0.149466 0.444353 0.002982 0.976519
demo%
```

Note the following:

- The range includes 0.0 and excludes 1.0.
- The algorithm is a multiplicative, congruential type, general random number generator.
- In general, the value of *i* is set *once* during execution of the calling program.
- The initial value of *i* should be a large odd integer.
- Each call to RAN gets the next random number in the sequence.
- To get a different sequence of random numbers each time you run the program, you must set the argument to a different initial value for each run.
- The argument is used by RAN to store a value for the calculation of the next random number according to the following algorithm:

$$\text{SEED} = 6909 * \text{SEED} + 1 \pmod{2^{32}}$$

- SEED contains a 32-bit number, and the high-order 24 bits are converted to floating point, and that value is returned.

1.45 rand, drand, irand: Return Random Values

rand returns real values in the range 0.0 through 1.0.

drand returns double precision values in the range 0.0 through 1.0.

irand returns positive integers in the range 0 through 2147483647.

These functions use random(3) to generate sequences of random numbers. The three functions share the same 256 byte state array. The only advantage of these functions is that they are widely available on UNIX systems. For better random number generators, compare lcrans, addrans, and shufrans. See also random(3), and the *Numerical Computation Guide*

<pre>i = irand(k) r = rand(k) d = drand(k)</pre>			
k	INTEGER*4	Input	k=0: Get next random number in the sequence k=1: Restart sequence, return first number k>0: Use as a seed for new sequence, return first number
rand	REAL*4	Output	
drand	REAL*8	Output	
irand	INTEGER*4	Output	

Example: irand():

```
integer*4 v(5), iflag/0/
do i = 1, 5
  v(i) = irand( iflag )
end do
write(*,*) v
end
demo% f95 trand.f
demo% a.out
      2078917053 143302914 1027100827 1953210302 755253631
demo%
```

1.46 rename: Rename a File

The function is called by:

INTEGER*4 rename <i>status</i> = rename(<i>from</i> , <i>to</i>)			
<i>from</i>	character*n	Input	Path name of an existing file
<i>to</i>	character*n	Input	New path name for the file
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> >0: System error code

If the file specified by *to* exists, then both *from* and *to* must be the same type of file, and must reside on the same file system. If *to* exists, it is removed first.

Example: rename()—Rename file `trename.old` to `trename.new`

```
demo% cat trename.f
      INTEGER*4 rename, status
      character*18 from/'trename.old'/, to/'trename.new'/
      status = rename( from, to )
      if ( status .ne. 0 ) stop 'rename: error'
      end
demo% f95 trename.f
demo% ls trename*
trename.f trename.old
demo% a.out
demo% ls trename*
trename.f trename.new
demo%
```

See also `rename(2)` and `perror(3F)`.

Note: the path names cannot be longer than `MAXPATHLEN` as defined in `<sys/param.h>`.

1.47 secnds: Get System Time in Seconds, Minus Argument

$t = \text{secnds}(t_0)$			
t_0	REAL	Input	Constant, variable, or array element
Return Value	REAL	Output	Number of seconds since midnight, minus t_0

Example: secnds:

```
demo% cat sec1.f
      real elapsed, t0, t1, x, y
      t0 = 0.0
      t1 = secnds( t0 )
      y = 0.1
      do i = 1, 10000
        x = asin( y )
      end do
      elapsed = secnds( t1 )
      write ( *, 1 ) elapsed
1    format ( ' 10000 arcsines: ', f12.6, ' sec' )
      end
demo% f95 sec1.f
demo% a.out
10000 arcsines:      0.009064 sec
demo%
```

Note that:

- The returned value from SECNDS is accurate to 0.01 second.
- The value is the system time, as the number of seconds from midnight, and it correctly spans midnight.
- Some precision may be lost for small time intervals near the end of the day.

1.48 sh: Fast Execution of an sh Command

The function is called by:

INTEGER*4 sh <i>status</i> = sh(<i>string</i>)			
<i>string</i>	character*n	Input	String containing command to do
Return value	INTEGER*4	Output	Exit status of the shell executed. See <i>wait(2)</i> for an explanation of this value.

Example: sh():

```
character*18 string / 'ls > MyOwnFile.names' /  
INTEGER*4 status, sh  
status = sh( string )  
if ( status .ne. 0 ) stop 'sh: error'  
...  
end
```

The function sh passes *string* to the sh shell as input, as if the string had been typed as a command.

The current process waits until the command terminates.

The forked process flushes all open files:

- For output files, the buffer is flushed to the actual file.
- For input files, the position of the pointer is unpredictable.

The sh() function is not MT-safe. Do not call it from multithreaded or parallelized programs.

See also: `execve(2)`, `wait(2)`, and `system(3)`.

Note: *string* cannot be longer than 1,024 characters.

1.49 signal: Change the Action for a Signal

The function is called by:

[illegible]

If *proc* is called, it is passed the signal number as an integer argument.

If a process incurs a signal, the default action is usually to clean up and abort. A signal handling routine provides the capability of catching specific exceptions or interrupts for special processing.

The returned value can be used in subsequent calls to `signal` to restore a previous action definition.

You can get a negative return value even though there is no error. In fact, if you pass a *valid* signal number to `signal()`, and you get a return value less than -1, then it is OK.

floatingpoint.h defines *proc* values SIGFPE_DEFAULT, SIGFPE_IGNORE, and SIGFPE_ABORT. See “floatingpoint.h: Fortran IEEE Definitions” on page 46.

In 64-bit environments, `signal` must be declared `INTEGER*8`, along with the variables receiving its output, to avoid truncation of the address that may be returned.

See also `kill(1)`, `signal(3)`, and `kill(3F)`, and *Numerical Computation Guide*.

1.50 `sleep`: Suspend Execution for an Interval

The subroutine is called by:

call <code>sleep(itime)</code>			
<i>itime</i>	INTEGER*4	Input	Number of seconds to sleep

The actual time can be up to 1 second less than *itime* due to granularity in system timekeeping.

Example: `sleep()`:

```
INTEGER*4 time / 5 /  
write(*,*) 'Start'  
call sleep( time )  
write(*,*) 'End'  
end
```

See also `sleep(3)`.

1.51 `stat`, `lstat`, `fstat`: Get File Status

These functions return the following information:

- device,
- inode number,
- protection,
- number of hard links,
- user ID,
- group ID,

- device type,
- size,
- access time,
- modify time,
- status change time,
- optimal blocksize,
- blocks allocated

Both `stat` and `lstat` query by file name. `fstat` queries by logical unit.

1.51.1 `stat`: Get Status for File, by File Name

The function is called by:

INTEGER*4 stat <i>ierr = stat (name, statb)</i>			
<i>name</i>	character*n	Input	Name of the file
<i>statb</i>	INTEGER*4	Output	Status structure for the file, 13-element array
Return value	INTEGER*4	Output	<i>ierr</i> =0: OK <i>ierr</i> >0: Error code

Example 1: `stat()`:

```

character name*18 /'MyFile'/
INTEGER*4 ierr, stat, lunit/1/, statb(13)
open( unit=lunit, file=name )
ierr = stat ( name, statb )
if ( ierr .ne. 0 ) stop 'stat: error'
write(*,*)'UID of owner = ',statb(5),', blocks = ',statb(13)
end

```

1.51.2 fstat: Get Status for File, by Logical Unit

The function

INTEGER*4 fstat <i>ierr</i> = fstat (<i>lunit</i> , <i>statb</i>)			
<i>lunit</i>	INTEGER*4	Input	Logical unit number
<i>statb</i>	INTEGER*4	Output	Status for the file: 13-element array
Return value	INTEGER*4	Output	<i>ierr</i> =0: OK <i>ierr</i> >0: Error code

is called by:

Example 2: fstat():

```
character name*18 /'MyFile'/
INTEGER*4 fstat, lunit/1/, statb(13)
open( unit=lunit, file=name )
ierr = fstat ( lunit, statb )
if ( ierr .ne. 0 ) stop 'fstat: error'
write(*,*)'UID of owner = ',statb(5),',', blocks = ',statb(13)
end
```

1.51.3 lstat: Get Status for File, by File Name

The function is called by:

<i>ierr</i> = lstat (<i>name</i> , <i>statb</i>)			
<i>name</i>	character*n	Input	File name
<i>statb</i>	INTEGER*4	Output	Status array of file, 13 elements
Return value	INTEGER*4	Output	<i>ierr</i> =0: OK <i>ierr</i> >0: Error code

Example 3: `lstat()`:

```
character name*18 /'MyFile'/  
INTEGER*4 lstat, lunit/1/, statb(13)  
open( unit=lunit, file=name )  
ierr = lstat ( name, statb )  
if ( ierr .ne. 0 ) stop 'lstat: error'  
write(*,*)'UID of owner = ',statb(5),', blocks = ',statb(13)  
end
```

1.51.4 Detail of Status Array for Files

The meaning of the information returned in the `INTEGER*4` array *statb* is as described for the structure *stat* under `stat(2)`.

Spare values are not included. The order is shown in the following table:

<code>statb(1)</code>	Device inode resides on
<code>statb(2)</code>	This inode's number
<code>statb(3)</code>	Protection
<code>statb(4)</code>	Number of hard links to the file
<code>statb(5)</code>	User ID of owner
<code>statb(6)</code>	Group ID of owner
<code>statb(7)</code>	Device type, for inode that is device
<code>statb(8)</code>	Total size of file
<code>statb(9)</code>	File last access time
<code>statb(10)</code>	File last modify time
<code>statb(11)</code>	File last status change time
<code>statb(12)</code>	Optimal blocksize for file system I/O ops
<code>statb(13)</code>	Actual number of blocks allocated

See also `stat(2)`, `access(3F)`, `perror(3F)`, and `time(3F)`.

Note: the path names can be no longer than `MAXPATHLEN` as defined in `<sys/param.h>`.

1.52 `stat64`, `lstat64`, `fstat64`: Get File Status

64-bit "long file" versions of `stat`, `lstat`, `fstat`. These routines are identical to the non-64-bit routines, except that the 13-element array *statb* must be declared `INTEGER*8`.

1.53 `system`: Execute a System Command

The function is called by:

<code>INTEGER*4 system</code> <code>status = system(string)</code>			
<i>string</i>	<code>character*n</code>	Input	String containing command to do
Return value	<code>INTEGER*4</code>	Output	Exit status of the shell executed. See <i>wait(2)</i> for an explanation of this value.

Example: `system()`:

```
character*8 string / 'ls s*' /
INTEGER*4 status, system
status = system( string )
if ( status .ne. 0 ) stop 'system: error'
end
```

The function `system` passes *string* to your shell as input, as if the string had been typed as a command. Note: *string* cannot be longer than 1024 characters.

If `system` can find the environment variable `SHELL`, then `system` uses the value of `SHELL` as the command interpreter (shell); otherwise, it uses `sh(1)`.

The current process waits until the command terminates.

Historically, `cc` developed with different assumptions:

- If `cc` calls `system`, the shell is always the Bourne shell.

- For output files, the buffer is flushed to the actual file.
- For input files, the position of the pointer is unpredictable.

The `system()` function is not MT-safe. Do not call it from multithreaded or parallelized programs.

These routines have the following functions:

<code>time</code>	Standard version: Get system time as integer (seconds since 0 GMT 1/1/70) VMS Version: Get the system time as character (hh:mm:ss)
<code>ctime</code>	Convert a system time to an ASCII string.
<code>ltime</code>	Dissect a system time into month, day, and so forth, local time.
<code>gmtime</code>	Dissect a system time into month, day, and so forth, GMT.

The *time()* function is called by:

INTEGER*4 time or INTEGER*8			
n = time()		Standard Version	
Return value	INTEGER*4	Output	Time, in seconds, since 0:0:0, GMT, 1/1/70
	INTEGER*8	Output	In 64-bit environments, time returns an INTEGER*8 value

The function `time()` returns an integer with the time since 00:00:00 GMT, January 1, 1970, measured in seconds. This is the value of the operating system clock.

Example: `time()`, version standard with the operating system:

```
      INTEGER*4  n, time
      n = time()
      write(*,*) 'Seconds since 0 1/1/70 GMT = ', n
      end
demo% f95 ttime.f
demo% a.out
      Seconds since 0 1/1/70 GMT =    913240205
demo%
```

1.54.2 `ctime`: Convert System Time to Character

The function `ctime` converts a system time, *stime*, and returns it as a 24-character ASCII string.

The function is called by:

CHARACTER <code>ctime</code> *24 <code>string = ctime(stime)</code>			
<i>stime</i>	INTEGER*4	Input	System time from <code>time()</code> (standard version)
Return value	character*24	Output	System time as character string. Declare <code>ctime</code> and <i>string</i> as character*24.

The format of the `ctime` returned value is shown in the following example. It is described in the man page `ctime(3C)`.

Example: `ctime()`:

```
      character*24 ctime, string
      INTEGER*4  n, time
      n = time()
      string = ctime( n )
      write(*,*) 'ctime: ', string
      end
demo% f95 tctime.f
demo% a.out
      ctime: Wed Dec  9 13:50:05 1998
demo%
```

1.54.3 `ltime`: Split System Time to Month, Day,... (Local)

This routine dissects a system time into month, day, and so forth, for the local time zone.

The subroutine is called by:

call <code>ltime(stime, tarray)</code>			
<i>stime</i>	INTEGER*4	Input	System time from <code>time()</code> (standard version)
<i>tarray</i>	INTEGER*4(9)	Output	System time, local, as day, month, year, ...

For the meaning of the elements in `tarray`, see the next section.

Example: `ltime()`:

```
integer*4 stime, tarray(9), time
stime = time()
call ltime( stime, tarray )
write(*,*) 'ltime: ', tarray
end
demo% f95 tltime.f
demo% a.out
ltime: 25 49 10 12 7 91 1 223 1
demo%
```

1.54.4 `gmtime`: Split System Time to Month, Day, ... (GMT)

This routine dissects a system time into month, day, and so on, for GMT.

The subroutine is:

call <code>gmtime(stime, tarray)</code>			
<i>stime</i>	INTEGER*4	Input	System time from <code>time()</code> (standard version)
<i>tarray</i>	INTEGER*4(9)	Output	System time, GMT, as day, month, year, ...

Example: `gmtime`:

```
integer*4 stime, tarray(9), time
stime = time()
call gmtime( stime, tarray )
write(*,*) 'gmtime: ', tarray
end
demo% f95t tgmtime.f
demo% a.out
gmtime:  12  44  19  18  5  94  6  168  0
demo%
```

Here are the `tarray()` values for `ltime` and `gmtime`: index, units, and range:

1	Seconds (0 - 61)	6	Year - 1900
2	Minutes (0 - 59)	7	Day of week (Sunday = 0)
3	Hours (0 - 23)	8	Day of year (0 - 365)
4	Day of month (1 - 31)	9	Daylight Saving Time,
5	Months since January (0 - 11)		1 if DST in effect

These values are defined by the C library routine `ctime(3C)`, which explains why the system may return a count of seconds *greater than* 59. See also: `idate(3F)`, and `fdate(3F)`.

1.54.5 `ctime64`, `gmtime64`, `ltime64`: System Time Routines for 64-bit Environments

These are versions of the corresponding routines `ctime`, `gmtime`, and `ltime`, to provide portability on 64-bit environments. They are identical to these routines except that the input variable `stime` must be `INTEGER*8`.

When used in a 32-bit environment with an `INTEGER*8 stime`, if the value of `stime` is beyond the `INTEGER*4` range `ctime64` returns all asterisks, while `gmtime` and `ltime` fill the `tarray` array with -1.

1.55 `ttynam`, `isatty`: Get Name of a Terminal Port

`ttynam` and `isatty` handle terminal port names.

1.55.1 `ttynam`: Get Name of a Terminal Port

The function `ttynam` returns a blank padded path name of the terminal device associated with logical unit *lunit*.

The function is called by:

CHARACTER <code>ttynam</code> *24 <code>name</code> = <code>ttynam</code> (<i>lunit</i>)			
<i>lunit</i>	INTEGER*4	Input	Logical unit
Return value	character*n	Output	If nonblank returned: <i>name</i> =path name of device on <i>lunit</i> . Size <i>n</i> must be large enough for the longest path name. If empty string (all blanks) returned: <i>lunit</i> is not associated with a terminal device in the directory, /dev

1.55.2 `isatty`: Is this Unit a Terminal?

The function `isatty` returns true or false depending on whether or not logical unit *lunit* is a terminal device or not.

The function is called by:

<code>terminal</code> = <code>isatty</code> (<i>lunit</i>)			
<i>lunit</i>	INTEGER*4	Input	Logical unit
Return value	LOGICAL*4	Output	<i>terminal</i> =true: It is a terminal device <i>terminal</i> =false: It is <i>not</i> a terminal device

Example: Determine if *lunit* is a tty:

```
character*12 name, ttynam
INTEGER*4 lunit /5/
logical*4 isatty, terminal
terminal = isatty( lunit )
name = ttynam( lunit )
write(*,*) 'terminal = ', terminal, ', name = "', name, '"'
end
```

The output is:

```
terminal = T, name = "/dev/ttyp1 "
```

1.56 unlink: Remove a File

The function is called by:

INTEGER*4 unlink <i>n</i> = unlink (<i>patnam</i>)			
<i>patnam</i>	character* <i>n</i>	Input	File name
Return value	INTEGER*4	Output	<i>n</i> =0: OK <i>n</i> >0: Error

The function `unlink` removes the file specified by path name *patnam*. If this is the last link to the file, the contents of the file are lost.

Example: `unlink()`—Remove the `tunlink.data` file:

```
call unlink( 'tunlink.data' )
end
demo% f95 tunlink.f
demo% ls tunl*
tunlink.f tunlink.data
demo% a.out
demo% ls tunl*
tunlink.f
demo%
```

See also: `unlink(2)`, `link(3F)`, and `perror(3F)`. Note: the path names cannot be longer than `MAXPATHLEN` as defined in `<sys/param.h>`.

1.57 wait: Wait for a Process to Terminate

The function is:

INTEGER*4 wait <i>n</i> = wait(<i>status</i>)			
<i>status</i>	INTEGER*4	Output	Termination status of the child process
Return value	INTEGER*4	Output	<i>n</i> >0: Process ID of the child process <i>n</i> <0: <i>n</i> =System error code; see <code>wait(2)</code> .

`wait` suspends the caller until a signal is received, or one of its child processes terminates. If any child has terminated since the last `wait`, return is immediate. If there are no children, return is immediate with an error code.

Example: Code fragment using `wait()`:

```
INTEGER*4 n, status, wait
...
n = wait( status )
if ( n .lt. 0 ) stop 'wait: error'
...
end
```

See also: `wait(2)`, `signal(3F)`, `kill(3F)`, and `perror(3F)`.

Fortran 95 Intrinsic Functions

This chapter lists the intrinsic function names recognized by the `f95` compiler.

2.1 Standard Fortran 95 Generic Intrinsic Functions

The generic Fortran 95 intrinsic functions are grouped in this section by functionality as they appear in the Fortran 95 standard.

The arguments shown are the names that can be used as argument keywords when using the keyword form, as in `cmplx(Y=B, KIND=M, X=A)`.

Consult the Fortran 95 standard for the detailed specifications of these generic intrinsic procedures.

2.1.1 Argument Presence Inquiry Function

Generic Intrinsic Name	Description
PRESENCE	Argument presence

2.1.2 Numeric Functions

Generic Intrinsic Name	Description
ABS (A)	Absolute value
AIMAG (Z)	Imaginary part of a complex number
AINT (A [, KIND])	Truncation to whole number
ANINT (A [, KIND])	Nearest whole number
CEILING (A [, KIND])	Least integer greater than or equal to number
CMPLX (X [, Y, KIND])	Conversion to complex type
CONJG (Z)	Conjugate of a complex number
DBLE (A)	Conversion to double precision real type
DIM (X, Y)	Positive difference
DPROD (X, Y)	Double precision real product
FLOOR (A [, KIND])	Greatest integer less than or equal to number
INT (A [, KIND])	Conversion to integer type
MAX (A1, A2 [, A3,...])	Maximum value
MIN (A1, A2 [, A3,...])	Minimum value
MOD (A, P)	Remainder function
MODULO (A, P)	Modulo function
NINT (A [, KIND])	Nearest integer
REAL (A [, KIND])	Conversion to real type
SIGN (A, B)	Transfer of sign

2.1.3 Mathematical Functions

Generic Intrinsic Name	Description
ACOS (X)	Arccosine
ASIN (X)	Arcsine
ATAN (X)	Arctangent
ATAN2 (Y, X)	Arctangent

Generic Intrinsic Name	Description
COS (X)	Cosine
COSH (X)	Hyperbolic cosine
EXP (X)	Exponential
LOG (X)	Natural logarithm
LOG10 (X)	Common logarithm (base 10)
SIN (X)	Sine
SINH (X)	Hyperbolic sine
SQRT (X)	Square root
TAN (X)	Tangent
TANH (X)	Hyperbolic tangent

2.1.4 Character Functions

Generic Intrinsic Name	Description
ACHAR (I)	Character in given position in ASCII collating sequence
ADJUSTL (STRING)	Adjust left
ADJUSTR (STRING)	Adjust right
CHAR (I [, KIND])	Character in given position in processor collating sequence
IACHAR (C)	Position of a character in ASCII collating sequence
ICHAR (C)	Position of a character in processor collating sequence
INDEX (STRING, SUBSTRING [, BACK])	Starting position of a substring
LEN_TRIM (STRING)	Length without trailing blank characters
LGE (STRING_A, STRING_B)	Lexically greater than or equal
LGT (STRING_A, STRING_B)	Lexically greater than
LLE (STRING_A, STRING_B)	Lexically less than or equal
LLT (STRING_A, STRING_B)	Lexically less than
REPEAT (STRING, NCOPIES)	Repeated concatenation

Generic Intrinsic Name	Description
SCAN (STRING, SET [, BACK])	Scan a string for a character in a set
TRIM (STRING)	Remove trailing blank characters
VERIFY (STRING, SET [, BACK])	Verify the set of characters in a string

2.1.5 Character Inquiry Function

Generic Intrinsic Name	Description
LEN (STRING)	Length of a character entity

2.1.6 Kind Functions

Generic Intrinsic Name	Description
KIND (X)	Kind type parameter value
SELECTED_INT_KIND (R)	Integer kind type parameter value, given range
SELECTED_REAL_KIND ([P, R])	Real kind type parameter value, given precision and range

2.1.7 Logical Function

Generic Intrinsic Name	Description
LOGICAL (L [, KIND])	Convert between objects of type logical with different kind type parameters

2.1.8 Numeric Inquiry Functions

Generic Intrinsic Name	Description
DIGITS (X)	Number of significant digits of the model
EPSILON (X)	Number that is almost negligible compared to one
HUGE (X)	Largest number of the model
MAXEXPONENT (X)	Maximum exponent of the model
MINEXPONENT (X)	Minimum exponent of the model
PRECISION (X)	Decimal precision
RADIX (X)	Base of the model
RANGE (X)	Decimal exponent range
TINY (X)	Smallest positive number of the model

2.1.9 Bit Inquiry Function

Generic Intrinsic Name	Description
BIT_SIZE (I)	Number of bits of the model

2.1.10 Bit Manipulation Functions

Generic Intrinsic Name	Description
BTEST (I, POS)	Bit testing
IAND (I, J)	Logical AND
IBCLR (I, POS)	Clear bit
IBITS (I, POS, LEN)	Bit extraction
IBSET (I, POS)	Set bit
IEOR (I, J)	Exclusive OR
IOR (I, J)	Inclusive OR

Generic Intrinsic Name	Description
ISHFT (I, SHIFT)	Logical shift
ISHFTC (I, SHIFT [, SIZE])	Circular shift
NOT (I)	Logical complement

2.1.11 Transfer Function

Generic Intrinsic Name	Description
TRANSFER (SOURCE, MOLD [, SIZE])	Treat first argument as if of type of second argument

2.1.12 Floating-Point Manipulation Functions

Generic Intrinsic Name	Description
EXPONENT (X)	Exponent part of a model number
FRACTION (X)	Fractional part of a number
NEAREST (X, S)	Nearest different processor number in given direction
RRSPACING (X)	Reciprocal of the relative spacing of model numbers near given number
SCALE (X, I)	Multiply a real by its base to an integer power
SET_EXPONENT (X, I)	Set exponent part of a number
SPACING (X)	Absolute spacing of model numbers near given number

2.1.13 Vector and Matrix Multiply Functions

Generic Intrinsic Name	Description
DOT_PRODUCT (VECTOR_A, VECTOR_B)	Dot product of two rank-one arrays
MATMUL (MATRIX_A, MATRIX_B)	Matrix multiplication

2.1.14 Array Reduction Functions

Generic Intrinsic Name	Description
ALL (MASK [, DIM])	True if all values are true
ANY (MASK [, DIM])	True if any value is true
COUNT (MASK [, DIM])	Number of true elements in an array
MAXVAL (ARRAY, DIM [, MASK]) or MAXVAL (ARRAY [, MASK])	Maximum value in an array
MINVAL (ARRAY, DIM [, MASK]) or MINVAL (ARRAY [, MASK])	Minimum value in an array
PRODUCT (ARRAY, DIM [, MASK]) or PRODUCT (ARRAY [, MASK])	Product of array elements
SUM (ARRAY, DIM [, MASK]) or SUM (ARRAY [, MASK])	Sum of array elements

2.1.15 Array Inquiry Functions

Generic Intrinsic Name	Description
ALLOCATED (ARRAY)	Array allocation status
LBOUND (ARRAY [, DIM])	Lower dimension bounds of an array
SHAPE (SOURCE)	Shape of an array or scalar
SIZE (ARRAY [, DIM])	Total number of elements in an array
UBOUND (ARRAY [, DIM])	Upper dimension bounds of an array

2.1.16 Array Construction Functions

Generic Intrinsic Name	Description
MERGE (TSOURCE, FSOURCE, MASK)	Merge under mask
PACK (ARRAY, MASK [, VECTOR])	Pack an array into an array of rank one under a mask
SPREAD (SOURCE, DIM, NCOPIES)	Replicates array by adding a dimension
UNPACK (VECTOR, MASK, FIELD)	Unpack an array of rank one into an array under a mask

2.1.17 Array Reshape Function

Generic Intrinsic Name	Description
RESHAPE (SOURCE, SHAPE[, PAD, ORDER])	Reshape an array

2.1.18 Array Manipulation Functions

Generic Intrinsic Name	Description
CSHIFT (ARRAY, SHIFT [, DIM])	Circular shift
EOSHIFT (ARRAY, SHIFT [, BOUNDARY, DIM])	End-off shift
TRANSPOSE (MATRIX)	Transpose of an array of rank two

2.1.19 Array Location Functions

Generic Intrinsic Name	Description
MAXLOC (ARRAY, DIM [, MASK]) or MAXLOC (ARRAY [, MASK])	Location of a maximum value in an array
MINLOC (ARRAY, DIM [, MASK]) or MINLOC (ARRAY [, MASK])	Location of a minimum value in an array

2.1.20 Pointer Association Status Functions

Generic Intrinsic Name	Description
ASSOCIATED (POINTER [, TARGET])	Association status inquiry or comparison
NULL ([MOLD])	Returns disassociated pointer

2.1.21 Intrinsic Subroutines

Generic Intrinsic Name	Description
CPU_TIME (TIME)	Obtain processor time
DATE_AND_TIME ([DATE, TIME, ZONE, VALUES])	Obtain date and time
MVBITS (FROM, FROMPOS, LEN, TO, TOPOS)	Copies bits from one integer to another
RANDOM_NUMBER (HARVEST)	Returns pseudorandom number
RANDOM_SEED ([SIZE, PUT, GET])	Initializes or restarts the pseudorandom number generator
SYSTEM_CLOCK ([COUNT, COUNT_RATE, COUNT_MAX])	Obtain data from the system clock

2.1.22 Specific names for intrinsic functions

TABLE 2-1 Specific and Generic Names for Fortran 95 Intrinsic Functions

Specific Name	Generic Name	Argument Type
ABS (A)	ABS (A)	default real
ACOS (X)	ACOS (X)	default real
AIMAG (Z)	AIMAG (Z)	default complex
AINT (A)	AINT (A)	default real
ALOG (X)	LOG (X)	default real
ALOG10 (X)	LOG10 (X)	default real
• AMAX0 (A1, A2 [, A3,...])	REAL (MAX (A1, A2 [, A3,...]))	default integer
• AMAX1 (A1, A2 [, A3,...])	MAX (A1, A2 [, A3,...])	default real
• AMIN0 (A1, A2 [, A3,...])	REAL (MIN (A1, A2 [, A3,...]))	default integer
• AMIN1 (A1, A2 [, A3,...])	MIN (A1, A2 [, A3,...])	default real
AMOD (A, P)	MOD (A, P)	default real
ANINT (A)	ANINT (A)	default real
ASIN (X)	ASIN (X)	default real
ATAN (X)	ATAN (X)	default real
ATAN2 (Y, X)	ATAN2 (Y, X)	default real
CABS (A)	ABS (A)	default complex
CCOS (X)	COS (X)	default complex
CEXP (X)	EXP (X)	default complex
• CHAR (I)	CHAR (I)	default integer
CLOG (X)	LOG (X)	default complex
CONJG (Z)	CONJG (Z)	default complex
COS (X)	COS (X)	default real
COSH (X)	COSH (X)	default real
CSIN (X)	SIN (X)	default complex
CSQRT (X)	SQRT (X)	default complex
DABS (A)	ABS (A)	double precision
DACOS (X)	ACOS (X)	double precision
DASIN (X)	ASIN (X)	double precision

TABLE 2-1 Specific and Generic Names for Fortran 95 Intrinsic Functions *(Continued)*

Specific Name	Generic Name	Argument Type
DATAN (X)	ATAN (X)	double precision
DATAN2 (Y, X)	ATAN2 (Y, X)	double precision
DCOS (X)	COS (X)	double precision
DCOSH (X)	COSH (X)	double precision
DDIM (X, Y)	DIM (X, Y)	double precision
DEXP (X)	EXP (X)	double precision
DIM (X, Y)	DIM (X, Y)	default real
DINT (A)	AINT (A)	double precision
DLOG (X)	LOG (X)	double precision
DLOG10 (X)	LOG10 (X)	double precision
• DMAX1 (A1, A2 [, A3,...])	MAX (A1, A2 [, A3,...])	double precision
• DMIN1 (A1, A2 [, A3,...])	MIN (A1, A2 [, A3,...])	double precision
DMOD (A, P)	MOD (A, P)	double precision
DNINT (A)	ANINT (A)	double precision
DPROD (X, Y)	DPROD (X, Y)	default real
DSIGN (A, B)	SIGN (A, B)	double precision
DSIN (X)	SIN (X)	double precision
DSINH (X)	SINH (X)	double precision
DSQRT (X)	SQRT (X)	double precision
DTAN (X)	TAN (X)	double precision
DTANH (X)	TANH (X)	double precision
EXP (X)	EXP (X)	default real
• FLOAT (A)	REAL (A)	default integer
IABS (A)	ABS (A)	default integer
• ICHAR (C)	ICHAR (C)	default character
IDIM (X, Y)	DIM (X, Y)	default integer
• IDINT (A)	INT (A)	double precision
IDNINT (A)	NINT (A)	double precision
• IFIX (A)	INT (A)	default real
INDEX (STRING, SUBSTRING)	INDEX (STRING, SUBSTRING)	default character

TABLE 2-1 Specific and Generic Names for Fortran 95 Intrinsic Functions *(Continued)*

Specific Name	Generic Name	Argument Type
• INT (A)	INT (A)	default real
ISIGN (A, B)	SIGN (A, B)	default integer
LEN (STRING)	LEN (STRING)	default character
• LGE (STRING_A, STRING_B)	LGE (STRING_A, STRING_B)	default character
• LGT (STRING_A, STRING_B)	LGT (STRING_A, STRING_B)	default character
• LLE (STRING_A, STRING_B)	LLE (STRING_A, STRING_B)	default character
• LLT (STRING_A, STRING_B)	LLT (STRING_A, STRING_B)	default character
• MAX0 (A1, A2 [, A3,...])	MAX (A1, A2 [, A3,...])	default integer
• MAX1 (A1, A2 [, A3,...])	INT (MAX (A1, A2 [, A3,...]))	default real
• MIN0 (A1, A2 [, A3,...])	MIN (A1, A2 [, A3,...])	default integer
• MIN1 (A1, A2 [, A3,...])	INT (MIN (A1, A2 [, A3,...]))	default real
MOD (A, P)	MOD (A, P)	default integer
NINT (A)	NINT (A)	default real
• REAL (A)	REAL (A)	default integer
SIGN (A, B)	SIGN (A, B)	default real
SIN (X)	SIN (X)	default real
SINH (X)	SINH (X)	default real
• SNGL (A)	REAL (A)	double precision
SQRT (X)	SQRT (X)	default real
TAN (X)	TAN (X)	default real
TANH (X)	TANH (X)	default real

Functions marked with • cannot be used as an actual argument. “Double precision” means double precision real.

2.2 Non-Standard Fortran 95 Intrinsic Functions

The following functions are considered intrinsics by the f95 compiler, but are not part of the Fortran 95 standard.

2.2.1 Basic Linear Algebra Functions (BLAS)

When compiling with `-xknown_lib=blas`, the compiler will recognize calls to the following routines as intrinsics and will optimize for and link to the Sun Performance Library implementation. The compiler will ignore user-supplied versions of these routines.

TABLE 2-2 BLAS Intrinsics

Function	Description
CAXPY DAXPY SAXPY ZAXPY	Product of a scalar and a vector plus a vector
CCOPY DCOPY SCOPY ZCOPY	Copy a vector
CDOTC CDOTU DDOT SDOT ZDOTC ZDOTU	Dot product (inner product)
CSCAL DSCAL SSCAL ZSCAL	Scale a vector

See the *Sun Performance Library User's Guide* for more information on these routines.

2.2.2 Interval Arithmetic Intrinsic Functions

The following table lists intrinsic functions that are recognized by the compiler when compiling for interval arithmetic (-x1a). For details, see the *Fortran 95 Interval Arithmetic Programming Reference*.

DINTERVAL	DIVIX	INF	INTERVAL
ISEMPTY	MAG	MID	MIG
NDIGITS	QINTERVAL	SINTERVAL	SUP
VDABS	VDACOS	VDASIN	VDATAN
VDATAN2	VDCEILING	VDCOS	VDCOSH
VDEXP	VDFLOOR	VDINF	VDINT
VDISEMPTY	VDLOG	VDLOG10	VDMAG
VDMID	VDMIG	VDMOD	VDNINT
VDSIGN	VDSIN	VDSINH	VDSQRT
VDSUP	VDTAN	VDTANH	VDWID
VQABS	VQCEILING	VQFLOOR	VQINF
VQINT	VQISEMPTY	VQMAG	VQMID
VQMIG	VQNINT	VQSUP	VQWID
VSABS	VSACOS	VSASIN	VSATAN
VSATAN2	VSCEILING	VSCOS	VSCOSH
VSEXP	VSFLOOR	VSINF	VSINT
VSISEMPTY	VSLOG	VSLOG10	VSMAG
VSMID	VSMIG	VSMOD	VSNINT
VSSIGN	VSSIN	VSSINH	VSSQRT
VSSUP	VSTAN	VSTANH	VSWID
WID			

2.2.3 Other Vendor Intrinsic Functions

The f95 compiler recognizes a variety of legacy intrinsic functions that were defined by Fortran compilers from other vendors, including Cray Research, Inc. These are obsolete and their use should be avoided.

TABLE 2-3 Intrinsic functions from Cray CF90 and other compilers

Function	Arguments	Description
CLOC	([C=]c)	Obtains the address of a character object
COMPL	([I=]i)	Bit-by-bit complement of a word. Use NOT(i) instead
COT	([X=]x)	Generic cotangent. (Also: DCOT, QCOT)
CSMG	([I=]i,[J=]j,[K=]k)	Conditional Scalar Merge
DSHIFTL	([I=]i,[J=]j,[K=]k)	Double-object left shift of i and j by k bits
DSHIFTR	([I=]i,[J=]j,[K=]k)	Double-object right shift of i and j by k bits
EQV	([I=]i,[J=]j)	Logical equivalence. Use IOER(i,j) instead.
FCD	([I=]i,[J=]j)	Constructs a character pointer
GETPOS	([I=]i)	Obtains file position
IBCHNG	([I=]i, [POS=]j)	Generic function to change specified bit in a word.
ISHA	([I=]i, [SHIFT=]j)	Generic arithmetic shift
ISHC	([I=]i, [SHIFT=]j)	Generic circular shift
ISHL	([I=]i, [SHIFT=]j)	Generic left shift
LEADZ	([I=]i)	Counts number of leading 0 bits
LENGTH	([I=]i)	Returns the number of Cray words successfully transferred
LOC	([I=]i)	Returns the address of a variable (Section 1.36, “loc: Return the Address of an Object” on page 1-54)
NEQV	([I=]i,[J=]j)	Logical non-equivalence. Use IOER(i,j) instead.
POPCNT	([I=]i)	Counts number of bits set to 1
POPPAR	([I=]i)	Computes bit population parity
SHIFT	([I=]i,[J=]j)	Shift left circular. Use ISHFT(i,j) or ISHFTC(i,j,k) instead.
SHIFTA	([I=]i,[J=]j)	Arithmetic shift with sign extension.
SHIFTL	([I=]i,[J=]j)	Shift left with zero fill. Use ISHFT(i,j) or ISHFTC(i,j,k) instead.
SHIFTR	([I=]i,[J=]j)	Shift right with zero fill. Use ISHFT(i,j) or ISHFTC(i,j,k) instead.

TABLE 2-3 Intrinsic functions from Cray CF90 and other compilers (*Continued*)

Function	Arguments	Description
TIMEF	()	Returns elapsed time since the first call
UNIT	([I=]i)	Returns status of BUFFERIN or BUFFEROUT
XOR	([I=]i,[J=]j)	Logical exclusive OR. Use IOER(i,j) instead.

See also Chapter 3 for a list of VMS Fortran 77 intrinsics.

2.2.4 Other Extensions

The Fortran 95 compiler recognizes the following additional intrinsic functions:

2.2.4.1 MPI_SIZEOF

MPI_SIZEOF(*x*, *size*, *error*)

Returns the size in bytes of the machine representation of the given variable, *x*. If *x* is an array, it returns the size of the base element and not the size of the whole array

x input; variable or array of arbitrary type

size output; integer; size in bytes of *x*

error output; integer; set to an error code if an error detected, zero otherwise

2.2.4.2 Memory Functions

Memory allocation, reallocation, and deallocation functions `malloc()`, `realloc()`, and `free()` are implemented as f95 intrinsics. See “`malloc`, `malloc64`, `realloc`, `free`: Allocate/Reallocate/Deallocate Memory” on page 58 for details.

FORTRAN 77 and VMS Intrinsic Functions

This chapter lists the set of FORTRAN 77 intrinsic functions accepted by f95 and is provided to aid migration of programs from legacy FORTRAN 77 to Fortran 95.

f95 recognizes as intrinsic functions all the FORTRAN 77 and VMS functions listed in this chapter, along with all the Fortran 95 functions listed in the previous chapter. As an aid to migrating legacy FORTRAN 77 programs to f95, compiling with -f77=intrinsics causes the compiler to recognize only FORTRAN 77 and VMS functions as intrinsics, and not the Fortran 95 intrinsics.

Intrinsic functions that are Sun extensions of the ANSI FORTRAN 77 standard are marked with *. Programs using non-standard intrinsics and library functions may not be portable to other platforms.

Intrinsic functions have *generic* and *specific* names when they accept arguments of more than one data type. In general, the *generic* name returns a value with the same data type as its argument. However, there are exceptions such as the type conversion functions (TABLE 3-2) and the inquiry functions (TABLE 3-7). The function may also be called by one of its *specific* names to handle a specific argument data type.

With functions that work on more than one data item (e.g. `sign(a1, a2)`), all the data arguments must be the same type.

In the following tables, the FORTRAN 77 intrinsic functions are listed by:

- Intrinsic Function –description of what the function does
- Definition – a mathematical definition
- No. of Args. – number of arguments the function accepts
- Generic Name – the function's generic name
- Specific Names – the function's specific names
- Argument Type – data type associated with each specific name
- Function Type – data type returned for specific argument data type

Note – Compiler option `-xtypemap` changes the default sizes of variables and has an affect on intrinsic references. See “Remarks” on page 14, and the discussion of default sizes and alignment in the *Fortran User’s Guide*.

3.1 Arithmetic and Mathematical Functions

This section details arithmetic, type conversion, trigonometric, and other functions. “a” stands for a function’s single argument, “a1” and “a2” for the first and second arguments of a two argument function, and “ar” and “ai” for the real and imaginary parts of a function’s complex argument.

3.1.1 Arithmetic Functions

TABLE 3-1 Fortran 77 Arithmetic Functions

Intrinsic Function	Definition	No. of Args	Generic Name	Specific Names	Argument Type	Function Type
Absolute value <i>See Note (6).</i>	$ a = (a^2 + ai^2)^{1/2}$	1	ABS	IABS ABS DABS CABS QABS ✕ ZABS ✕ CDABS ✕ CQABS ✕	INTEGER REAL DOUBLE COMPLEX REAL *16 DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32	INTEGER REAL DOUBLE REAL REAL *16 DOUBLE DOUBLE REAL *16
Truncation <i>See Note (1).</i>	int(a)	1	AINT	AINT DINT QINT ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Nearest whole number	int(a+.5) if a ≥ 0 int(a-.5) if a < 0	1	ANINT	ANINT DNINT QNINT ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Nearest integer	int(a+.5) if a ≥ 0 int(a-.5) if a < 0	1	NINT	NINT IDNINT IQNINT ✕	REAL DOUBLE REAL *16	INTEGER INTEGER INTEGER
Remainder <i>See Note (1).</i>	a1-int(a1/a2)*a2	2	MOD	MOD AMOD DMOD QMOD ✕	INTEGER REAL DOUBLE REAL *16	INTEGER REAL DOUBLE REAL *16
Transfer of sign	a1 if a2 ≥ 0 - a1 if a2 < 0	2	SIGN	ISIGN SIGN DSIGN QSIGN ✕	INTEGER REAL DOUBLE REAL *16	INTEGER REAL DOUBLE REAL *16
Positive difference	a1-a2 if a1 > a2 0 if a1 ≤ a2	2	DIM	IDIM DIM DDIM QDIM ✕	INTEGER REAL DOUBLE REAL *16	INTEGER REAL DOUBLE REAL *16
Double and quad products	a1 * a2	2	-	DPROD QPROD ✕	REAL DOUBLE	DOUBLE REAL *16

TABLE 3-1 Fortran 77 Arithmetic Functions (*Continued*)

Intrinsic Function	Definition	No. of Args .	Generic Name	Specific Names	Argument Type	Function Type
Choosing largest value	max(a1, a2, ...)	≥ 2	MAX	MAX0	INTEGER	INTEGER
				AMAX1	REAL	REAL
				DMAX1	DOUBLE	DOUBLE
				QMAX1 $\neq$	REAL *16	REAL *16
				AMAX0	AMAX0	INTEGER
Choosing smallest value	min(a1, a2, ...)	≥ 2	MIN	MIN0	INTEGER	INTEGER
				AMIN1	REAL	REAL
				DMIN1	DOUBLE	DOUBLE
				QMIN1 $\neq$	REAL *16	REAL *16
				AMIN0	AMIN0	INTEGER
			MIN1	MIN1	REAL	INTEGER

3.1.2 Type Conversion Functions

TABLE 3-2 Fortran 77 Type Conversion Functions

Conversion to	No. of Args	Generic Name	Specific Names	Argument Type	Function Type
INTEGER <i>See Note (1).</i>	1	INT	-	INTEGER	INTEGER
			INT	REAL	INTEGER
			IFIX	REAL	INTEGER
			IDINT	DOUBLE	INTEGER
			-	COMPLEX	INTEGER
			-	COMPLEX*16	INTEGER
			-	COMPLEX*32	INTEGER
			IQINT ✕	REAL*16	INTEGER
REAL <i>See Note (2).</i>	1	REAL	REAL	INTEGER	REAL
			FLOAT	INTEGER	REAL
			-	REAL	REAL
			SNGL	DOUBLE	REAL
			SNGLQ ✕	REAL*16	REAL
			-	COMPLEX	REAL
			-	COMPLEX*16	REAL
			-	COMPLEX*32	REAL
DOUBLE <i>See Note (3).</i>	1	DBLE	FLOATK	INTEGER*8	REAL*4
			DBLE	INTEGER	DOUBLE PRECISION
			DFLOAT	INTEGER	DOUBLE PRECISION
			DFLOATK	INTEGER*8	DOUBLE PRECISION
			DREAL ✕	REAL	DOUBLE PRECISION
			-	DOUBLE	DOUBLE PRECISION
			-	COMPLEX	DOUBLE PRECISION
			-	COMPLEX*16	DOUBLE PRECISION
			-	REAL*16	DOUBLE PRECISION
			-	COMPLEX*32	DOUBLE PRECISION
			DBLEQ ✕	REAL*16	DOUBLE PRECISION
			-	COMPLEX*32	DOUBLE PRECISION
REAL*16 <i>See Note (3').</i>	1	QREAL ✕	QREAL ✕	INTEGER	REAL*16
			QFLOAT ✕	INTEGER	REAL*16
		QEXT ✕	-	REAL	REAL*16
			QEXT ✕	INTEGER	REAL*16
		QEXTD ✕	QEXTD ✕	DOUBLE	REAL*16
			-	REAL*16	REAL*16
		-	-	COMPLEX	REAL*16
			-	COMPLEX*16	REAL*16
			-	COMPLEX*32	REAL*16
			-	COMPLEX*32	REAL*16

TABLE 3-2 Fortran 77 Type Conversion Functions (*Continued*)

Conversion to	No. of Args	Generic Name	Specific Names	Argument Type	Function Type
COMPLEX <i>See Notes (4) and (8).</i>	1 or 2	CMPLX	-	INTEGER	COMPLEX
			-	REAL	COMPLEX
			-	DOUBLE	COMPLEX
			-	REAL*16	COMPLEX
			-	COMPLEX	COMPLEX
			-	COMPLEX*16	COMPLEX
			-	COMPLEX*32	COMPLEX
DOUBLE COMPLEX <i>See Note (8).</i>	1 or 2	DCMPLX⌘	-	INTEGER	DOUBLE COMPLEX
			-	REAL	DOUBLE COMPLEX
			-	DOUBLE	DOUBLE COMPLEX
			-	REAL*16	DOUBLE COMPLEX
			-	COMPLEX	DOUBLE COMPLEX
			-	COMPLEX*16	DOUBLE COMPLEX
			-	COMPLEX*32	DOUBLE COMPLEX
COMPLEX*32 <i>See Note (8).</i>	1 or 2	QCMPLX⌘	-	INTEGER	COMPLEX*32
			-	REAL	COMPLEX*32
			-	DOUBLE	COMPLEX*32
			-	REAL*16	COMPLEX*32
			-	COMPLEX	COMPLEX*32
			-	COMPLEX*16	COMPLEX*32
			-	COMPLEX*32	COMPLEX*32
INTEGER <i>See Note (5).</i>	1	-	ICHAR	CHARACTER	INTEGER
		-	IACHAR ⌘		
CHARACTER <i>See Note (5).</i>	1	-	CHAR	INTEGER	CHARACTER
		-	ACHAR ⌘		

On an ASCII platform, including Sun systems:

- ACHAR is a nonstandard synonym for CHAR
- IACHAR is a nonstandard synonym for ICHAR

On a non-ASCII platform, ACHAR and IACHAR were intended to provide a way to deal directly with ASCII.

3.1.3 Trigonometric Functions

TABLE 3-3 Fortran 77 Trigonometric Functions

Intrinsic Function	Definition	Args	Generic Name	Specific Names	Argument Type	Function Type
Sine <i>See Note (7).</i>	$\sin(a)$	1	SIN	SIN DSIN QSIN ✖ CSIN ZSIN ✖ CDSIN ✖ CQSIN ✖	REAL DOUBLE REAL *16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32	REAL DOUBLE REAL *16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32
Sine (degrees) <i>See Note (7).</i>	$\sin(a)$	1	SIND ✖	SIND ✖ DSIND ✖ QSIND ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Cosine <i>See Note (7).</i>	$\cos(a)$	1	COS	COS DCOS QCOS ✖ CCOS ZCOS ✖ CDCOS ✖ CQCOS ✖	REAL DOUBLE REAL *16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32	REAL DOUBLE REAL *16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32
Cosine (degrees) <i>See Note (7).</i>	$\cos(a)$	1	COSD ✖	COSD ✖ DCOSD ✖ QCOSD ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Tangent <i>See Note (7).</i>	$\tan(a)$	1	TAN	TAN DTAN QTAN ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Tangent (degrees) <i>See Note (7).</i>	$\tan(a)$	1	TAND ✖	TAND ✖ DTAND ✖ QTAND ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Arcsine <i>See Note (7).</i>	$\arcsin(a)$	1	ASIN	ASIN DASIN QASIN ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Arcsine (degrees) <i>See Note (7).</i>	$\arcsin(a)$	1	ASIND ✖	ASIND ✖ DASIND ✖ QASIND ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Arccosine <i>See Note (7).</i>	$\arccos(a)$	1	ACOS	ACOS DACOS QACOS ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16

TABLE 3-3 Fortran 77 Trigonometric Functions (*Continued*)

Intrinsic Function	Definition	Args	Generic Name	Specific Names	Argument Type	Function Type
Arccosine (degrees) <i>See Note (7).</i>	arccos(a)	1	ACOSD ✕	ACOSD ✕ DACOSD ✕ QACOSD ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Arctangent <i>See Note (7).</i>	arctan(a)	1	ATAN	ATAN DATAN QATAN ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
	arctan (a1/a2)	2	ATAN2	ATAN2 DATAN2 QATAN2 ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Arctangent (degrees) <i>See Note (7).</i>	arctan(a)	1	ATAND ✕	ATAND ✕ DATAND ✕ QATAND ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
	arctan (a1/a2)	2	ATAN2D✕	ATAN2D ✕ DATAN2D ✕ QATAN2D ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Hyperbolic Sine <i>See Note (7).</i>	sinh(a)	1	SINH	SINH DSINH QSINH ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Hyperbolic Cosine <i>See Note (7).</i>	cosh(a)	1	COSH	COSH DCOSH QCOSH ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Hyperbolic Tangent <i>See Note (7).</i>	tanh(a)	1	TANH	TANH DTANH QTANH ✕	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16

3.1.4 Other Mathematical Functions

TABLE 3-4 Other Fortran 77 Mathematical Functions

Intrinsic Function	Definition	No. of Args.	Generic Name	Specific Names	Argument Type	Function Type
Imaginary part of a complex number <i>See Note (6).</i>	ai	1	IMAG	AIMAG DIMAG ✕ QIMAG ✕	COMPLEX DOUBLE COMPLEX COMPLEX*32	REAL DOUBLE REAL*16
Conjugate of a complex number <i>See Note (6).</i>	(ar, -ai)	1	CONJG	CONJG DCONJG ✕ QCONJG ✕	COMPLEX DOUBLE COMPLEX COMPLEX*32	COMPLEX DOUBLE COMPLEX COMPLEX*32
Square root	a**(1/2)	1	SQRT	SQRT DSQRT QSQRT ✕ CSQRT ZSQRT ✕ CDSQRT ✕ CQSQRT ✕	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32
Cube root <i>See Note(8').</i>	a**(1/3)	1	CBRT	CBRT ✕ DCBRT ✕ QCBRT ✕ CCBRT ✕ ZCBRT ✕ CDCBRT ✕ CQCBRT ✕	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32
Exponential	e**a	1	EXP	EXP DEXP QEXP ✕ CEXP ZEXP ✕ CDEXP ✕ CQEXP ✕	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32
Natural logarithm	log(a)	1	LOG	ALOG DLOG QLOG ✕ CLOG ZLOG ✕ CDLOG ✕ CQLOG ✕	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32	REAL DOUBLE REAL*16 COMPLEX DOUBLE COMPLEX DOUBLE COMPLEX COMPLEX*32

TABLE 3-4 Other Fortran 77 Mathematical Functions *(Continued)*

Intrinsic Function	Definition	No. of Args.	Generic Name	Specific Names	Argument Type	Function Type
Common logarithm	$\log_{10}(a)$	1	LOG10	ALOG10 DLOG10 QLOG10 ✖	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Error function (See note below)	$\text{erf}(a)$	1	ERF	ERF ✖ DERF ✖	REAL DOUBLE	REAL DOUBLE
Error function	$1.0 - \text{erf}(a)$	1	ERFC	ERFC ✖ DERFC ✖	REAL DOUBLE	REAL DOUBLE

- The error function: $2/\sqrt{\pi} \times \text{integral from } 0 \text{ to } a \text{ of } \exp(-t^2) dt$

3.2 Character Functions

TABLE 3-5 Fortran 77 Character Functions

Intrinsic Function	Definition	No. of Args.	Specific Names	Argument Type	Function Type
Conversion <i>See Note (5).</i>	Conversion to character	1	CHAR	INTEGER	CHARACTER
	Conversion to integer <i>See also:</i> TABLE 3-2.	1	ACHAR ✕ ICHAR IACHAR ✕	CHARACTER	INTEGER
Index of a substring	Location of substring a2 in string a1 <i>See Note (10).</i>	2	INDEX	CHARACTER	INTEGER
Length	Length of character entity <i>See Note (11).</i>	1	LEN	CHARACTER	INTEGER
Lexically greater than or equal	$a1 \geq a2$ <i>See Note (12).</i>	2	LGE	CHARACTER	LOGICAL
Lexically greater than	$a1 > a2$ <i>See Note (12).</i>	2	LGT	CHARACTER	LOGICAL
Lexically less than or equal	$a1 \leq a2$ <i>See Note (12).</i>	2	LLE	CHARACTER	LOGICAL
Lexically less than	$a1 < a2$ <i>See Note (12).</i>	2	LLT	CHARACTER	LOGICAL

On an ASCII platform (including Sun systems):

- ACHAR is a nonstandard synonym for CHAR
- IACHAR is a nonstandard synonym for ICHAR

On a non-ASCII platform, ACHAR and IACHAR were intended to provide a way to deal directly with ASCII.

3.3 Miscellaneous Functions

Other miscellaneous functions include bitwise functions, environmental inquiry functions, and memory allocation and deallocation functions.

3.3.1 Bit Manipulation ✖

None of these functions are part of the FORTRAN 77 Standard.

TABLE 3-6 Fortran 77 Bitwise Functions

Bitwise Operations	No. of Args.	Specific Name	Argument Type	Function Type
Complement	1	NOT	INTEGER	INTEGER
And	2	AND	INTEGER	INTEGER
	2	IAND	INTEGER	INTEGER
Inclusive or	2	OR	INTEGER	INTEGER
	2	IOR	INTEGER	INTEGER
Exclusive or	2	XOR	INTEGER	INTEGER
	2	IEOR	INTEGER	INTEGER
Shift <i>See Note (14).</i>	2	ISHFT	INTEGER	INTEGER
Left shift <i>See Note (14).</i>	2	LSHIFT	INTEGER	INTEGER
Right shift <i>See Note (14).</i>	2	RSHIFT	INTEGER	INTEGER
Logical right shift <i>See Note (14).</i>	2	LRSHT	INTEGER	INTEGER
Circular shift	3	ISHFTC	INTEGER	INTEGER
Bit extraction	3	IBITS	INTEGER	INTEGER
Bit set	2	IBSET	INTEGER	INTEGER
Bit test	2	BTEST	INTEGER	LOGICAL
Bit clear	2	IBCLR	INTEGER	INTEGER

The above functions are available as intrinsic or extrinsic functions. See also the discussion of the library bit manipulation routines in the *Fortran Library Reference* manual.

3.3.2 Environmental Inquiry Functions ✱

None of these functions are part of the FORTRAN 77 Standard.

TABLE 3-7 Fortran 77 Environmental Inquiry Functions

Definition	No. of Args.	Generic Name	Argument Type	Function Type
Base of Number System	1	EPBASE	INTEGER REAL DOUBLE REAL *16	INTEGER INTEGER INTEGER INTEGER
Number of Significant Bits	1	EPPREC	INTEGER REAL DOUBLE REAL *16	INTEGER INTEGER INTEGER INTEGER
Minimum Exponent	1	EPEMIN	REAL DOUBLE REAL *16	INTEGER INTEGER INTEGER
Maximum Exponent	1	EPEMAX	REAL DOUBLE REAL *16	INTEGER INTEGER INTEGER
Least Nonzero Number	1	EPTINY	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16
Largest Number Representable	1	EPHUGE	INTEGER REAL DOUBLE REAL *16	INTEGER REAL DOUBLE REAL *16
Epsilon <i>See Note (16).</i>	1	EPMRSP	REAL DOUBLE REAL *16	REAL DOUBLE REAL *16

3.3.3 Memory ☞

None of these functions are part of the FORTRAN 77 Standard.

TABLE 3-8 Fortran 77 Memory Functions

Intrinsic Function	Definition	No. of Args	Specific Name	Argument Type	Function Type
Location	Address of <i>See Note (17).</i>	1	LOC	<i>Any</i>	INTEGER*4 INTEGER*8
Allocate	Allocate memory and return address. <i>See Note (17).</i>	1	MALLOC MALLOC64	INTEGER*4 INTEGER*8	INTEGER INTEGER*8
Deallocate	Deallocate memory allocated by MALLOC. <i>See Note (17).</i>	1	FREE	<i>Any</i>	-
Size	Return the size of the argument in bytes. <i>See Note (18).</i>	1	SIZEOF	<i>Any expression</i>	INTEGER

3.4 Remarks

The following remarks apply to all of the intrinsic function tables in this chapter.

- The abbreviation **DOUBLE** stands for **DOUBLE PRECISION**.
- An intrinsic that takes **INTEGER** arguments accepts **INTEGER*2**, **INTEGER*4**, or **INTEGER*8**.
- **INTEGER** intrinsics that take **INTEGER** arguments return values of **INTEGER** type determined as follows. Note that the **-xtypemap** option may alter the default sizes of actual arguments:
 - **mod sign dim max min** and **iand** or **ior xor ieor** — size of the value returned is the largest of the sizes of the arguments.
 - **abs ishft lshift rshift lrshft ibset btest ivclr ishftc ibits** — size of the value returned is the size of the first argument.
 - **int epbase epprec** — size of the value returned is the size of default **INTEGER**.
 - **ephuge** — size of the value returned is the size of the default **INTEGER**, or the size of the argument, whichever is largest.

- Options that change the default data sizes also change the way some intrinsics are used. For example, with `-dbl` in effect, a call to `ZCOS` with a `DOUBLE COMPLEX` argument will automatically become a call to `CQCOS` because the argument has been promoted to `COMPLEX*32`. The following functions have this capability:

```
aimag alog amod cabs cbrt ccos cdabs cdcbrt cdcos cdexp
cdlog cdsin cdsqrt cexp clog csin csqrt dabs dacos dacosd
dasin dasind datan datand dcbdt dconjg dcos dcosh ddim
derf derfc dexp dimag dint dlog dmod dnint dprod dsign dsin
dsind dsinh dsqrt dtan dtand dtanh idnint iidnnt jidnnt zabs
zcbdt zcos zexp zlog zsin zsqr
```

- The following functions permit arguments of an integer or logical type of any size:

```
and iand ieor iiand iieor iior inot ior jiand jieor jior jnot
lrshft lshift not or rshft xor
```

- An intrinsic that is shown to return a *default* `REAL`, `DOUBLE PRECISION`, `COMPLEX`, or `DOUBLE COMPLEX` value will return the prevailing type depending on certain compilation options. For example, if compiled with `-xtypemap=real:64,double:64`:
 - A call to a `REAL` function returns `REAL*8`
 - A call to a `DOUBLE PRECISION` function returns `REAL*8`
 - A call to a `COMPLEX` function returns `COMPLEX*16`
 - A call to a `DOUBLE COMPLEX` function returns `COMPLEX*16`

Other options that alter the data sizes of default data types are `-r8` and `-dbl`, which also promote `DOUBLE` to `QUAD`. The `-xtypemap=` option provides more flexibility than these older compiler options and is preferred.

- A function with a generic name returns a value with the same type as the argument—except for type conversion functions, the nearest integer function, the absolute value of a complex argument, and others. If there is more than one argument, they must all be of the same type.
- If a function name is used as an actual argument, then it must be a specific name.
- If a function name is used as a dummy argument, then it does not identify an intrinsic function in the subprogram, and it has a data type according to the same rules as for variables and arrays.

3.4.1 Notes on Functions

Tables and notes 1 through 12 are based on the “Table of Intrinsic Functions,” from *ANSI X3.9-1978 Programming Language FORTRAN*, with the Fortran extensions added.

(1) INT

If A is type integer, then $\text{INT}(A)$ is A .

If A is type real or double precision, then:

if $|A| < 1$, then $\text{INT}(A)$ is 0

if $|A| \geq 1$, then $\text{INT}(A)$ is the greatest integer that does not exceed the magnitude of A , and whose sign is the same as the sign of A . (Such a mathematical integer value may be too large to fit in the computer integer type.)

If A is type complex or double complex, then apply the above rule to the real part of A .

If A is type real, then $\text{IFIX}(A)$ is the same as $\text{INT}(A)$.

(2) REAL

If A is type real, then $\text{REAL}(A)$ is A .

If A is type integer or double precision, then $\text{REAL}(A)$ is as much precision of the significant part of A as a real datum can contain.

If A is type complex, then $\text{REAL}(A)$ is the real part of A .

If A is type double complex, then $\text{REAL}(A)$ is as much precision of the significant part of the real part of A as a real datum can contain.

(3) DBLE

If A is type double precision, then $\text{DBLE}(A)$ is A .

If A is type integer or real, then $\text{DBLE}(A)$ is as much precision of the significant part of A as a double precision datum can contain.

If A is type complex, then $\text{DBLE}(A)$ is as much precision of the significant part of the real part of A as a double precision datum can contain.

If A is type COMPLEX^*16 , then $\text{DBLE}(A)$ is the real part of A .

(3') QREAL

If A is type REAL^*16 , then $\text{QREAL}(A)$ is A .

If A is type integer, real, or double precision, then $\text{QREAL}(A)$ is as much precision of the significant part of A as a REAL^*16 datum can contain.

If A is type complex or double complex, then $\text{QREAL}(A)$ is as much precision of the significant part of the real part of A as a REAL^*16 datum can contain.

If A is type COMPLEX^*16 or COMPLEX^*32 , then $\text{QREAL}(A)$ is the real part of A .

(4) CMPLX

If A is type complex, then `CMPLX(A)` is A.

If A is type integer, real, or double precision, then `CMPLX(A)` is `REAL(A) + 0i`.

If A1 and A2 are type integer, real, or double precision, then `CMPLX(A1, A2)` is `REAL(A1) + REAL(A2)*i`.

If A is type double complex, then `CMPLX(A)` is `REAL( DBLE(A) ) + i*REAL( DIMAG(A) )`.

If `CMPLX` has two arguments, then they must be of the same type, and they may be one of integer, real, or double precision.

If `CMPLX` has one argument, then it may be one of integer, real, double precision, complex, `COMPLEX*16`, or `COMPLEX*32`.

(4') `DCMPLX`

If A is type `COMPLEX*16`, then `DCMPLX(A)` is A.

If A is type integer, real, or double precision, then `DCMPLX(A)` is `DBLE(A) + 0i`.

If A1 and A2 are type integer, real, or double precision, then `DCMPLX(A1, A2)` is `DBLE(A1) + DBLE(A2)*i`.

If `DCMPLX` has two arguments, then they must be of the same type, and they may be one of integer, real, or double precision.

If `DCMPLX` has one argument, then it may be one of integer, real, double precision, complex, `COMPLEX*16`, or `COMPLEX*32`.

(5) `ICHAR`

`ICHAR(A)` is the position of A in the collating sequence.

The first position is 0, the last is N-1, $0 \leq \text{ICHAR}(A) \leq N-1$, where N is the number of characters in the collating sequence, and A is of type character of length one.

`CHAR` and `ICHAR` are inverses in the following sense:

- `ICHAR(CHAR(I)) = I`, for $0 \leq I \leq N-1$
- `CHAR(ICHAR(C)) = C`, for any character C capable of representation in the processor

(6) `COMPLEX`

A `COMPLEX` value is expressed as an ordered pair of reals, `(ar, ai)`, where `ar` is the real part, and `ai` is the imaginary part.

(7) Radians

All angles are expressed in radians, unless the "Intrinsic Function" column includes the "(degrees)" remark.

(8) COMPLEX Function

The result of a function of type COMPLEX is the principal value.

(8') CBRT

If *a* is of COMPLEX type, CBRT results in COMPLEX RT1=(*A*, *B*), where:
 $A \geq 0.0$, and $-60 \text{ degrees} \leq \arctan (B/A) < +60 \text{ degrees}$.

Other two possible results can be evaluated as follows:

- RT2 = RT1 * (-0.5, square_root (0.75))
- RT3 = RT1 * (-0.5, square_root (0.75))

(9) Argument types

All arguments in an intrinsic function reference must be of the same type.

(10) INDEX

INDEX(*X*, *Y*) is the place in *X* where *Y* starts. That is, it is the starting position within character string *X* of the first occurrence of character string *Y*.

If *Y* does not occur in *X*, then INDEX(*X*, *Y*) is 0.

If $\text{LEN}(X) < \text{LEN}(Y)$, then INDEX(*X*, *Y*) is 0.

INDEX returns default INTEGER*4 data. If compiling for a 64-bit environment, the compiler will issue a warning if the result overflows the INTEGER*4 data range. To use INDEX in a 64-bit environment with character strings larger than the INTEGER*4 limit (2 Gbytes), the INDEX function and the variables receiving the result must be declared INTEGER*8.

(11) LEN

LEN returns the declared length of the CHARACTER argument variable. The actual value of the argument is of no importance.

LEN returns default INTEGER*4 data. If compiling for a 64-bit environment, the compiler will issue a warning if the result overflows the INTEGER*4 data range. To use LEN in a 64-bit environment with character variables larger than the INTEGER*4 limit (2 Gbytes), the LEN function and the variables receiving the result must be declared INTEGER*8.

(12) Lexical Compare

LGE(*X*, *Y*) is true if *X*=*Y*, or if *X* follows *Y* in the collating sequence; otherwise, it is false.

LGT(*X*, *Y*) is true if *X* follows *Y* in the collating sequence; otherwise, it is false.

LLE(*X*, *Y*) is true if $X=Y$, or if *X* precedes *Y* in the collating sequence; otherwise, it is false.

LLT(*X*, *Y*) is true if *X* precedes *Y* in the collating sequence; otherwise, it is false.

If the operands for LGE, LGT, LLE, and LLT are of unequal length, the shorter operand is considered as if it were extended on the right with blanks.

(13) Bit Functions

There are other bitwise operations in VMS Fortran, but these are not implemented.

(14) Shift

LSHIFT shifts *a1* logically *left* by *a2* bits (inline code).

LRSHT shifts *a1* logically *right* by *a2* bits (inline code).

RSHIFT shifts *a1* arithmetically *right* by *a2* bits.

ISHFT shifts *a1* logically *left* if $a2 > 0$ and *right* if $a2 < 0$.

The LSHIFT and RSHIFT functions are the Fortran analogs of the C << and >> operators. As in C, the semantics depend on the hardware.

The behavior of the shift functions with an out of range shift count is hardware dependent and generally unpredictable. In this release, shift counts larger than 31 result in hardware dependent behavior.

(15) Environmental inquiries

Only the type of the argument is significant.

(16) Epsilon

Epsilon is the least e , such that $1.0 + e \neq 1.0$.

(17) LOC, MALLOC, and FREE

The LOC function returns the address of a variable or of an external procedure. The function call MALLOC(*n*) allocates a block of at least *n* bytes, and returns the address of that block.

LOC returns default INTEGER*4 in 32-bit environments, INTEGER*8 in 64-bit environments.

MALLOC is a library function and not an intrinsic in FORTRAN 77. It too returns default INTEGER*4 in 32-bit environments, INTEGER*8 in 64-bit environments. However, MALLOC must be explicitly declared INTEGER*8 when compiling for 64-bit environments.

The value returned by `LOC` or `MALLOC` should be stored in variables typed `POINTER`, `INTEGER*4`, or `INTEGER*8` in 64-bit environments. The argument to `FREE` must be the value returned by a previous call to `MALLOC` and hence should have data type `POINTER`, `INTEGER*4`, or `INTEGER*8`.

`MALLOC64` always takes an `INTEGER*8` argument (size of memory request in bytes) and always returns an `INTEGER*8` value. Use this routine rather than `MALLOC` when compiling programs that must run in both 32-bit and 64-bit environments. The receiving variable must be declared either `POINTER` or `INTEGER*8`.

(18) `SIZEOF`

The `SIZEOF` intrinsic cannot be applied to arrays of an assumed size, characters of a length that is passed, or subroutine calls or names. `SIZEOF` returns default `INTEGER*4` data. If compiling for a 64-bit environment, the compiler will issue a warning if the result overflows the `INTEGER*4` data range. To use `SIZEOF` in a 64-bit environment with arrays larger than the `INTEGER*4` limit (2 Gbytes), the `SIZEOF` function and the variables receiving the result must be declared `INTEGER*8`.

3.5 VMS Intrinsic Functions

This section lists VMS FORTRAN intrinsic routines recognized by f95. They are, of course, nonstandard. ✱

3.5.1 VMS Double-Precision Complex

TABLE 3-9 VMS Double-Precision Complex Functions

Generic Name	Specific Names	Function	Argument Type	Result Type
	CDABS	Absolute value	COMPLEX*16	REAL*8
	CDEXP	Exponential, e^{**a}	COMPLEX*16	COMPLEX*16
	CDLOG	Natural log	COMPLEX*16	COMPLEX*16
	CDSQRT	Square root	COMPLEX*16	COMPLEX*16
	CDSIN	Sine	COMPLEX*16	COMPLEX*16
	CDCOS	Cosine	COMPLEX*16	COMPLEX*16
DCMPLX		Convert to DOUBLE COMPLEX	<i>Any numeric</i>	COMPLEX*16
	DCONJG	Complex conjugate	COMPLEX*16	COMPLEX*16
	DIMAG	Imaginary part of complex	COMPLEX*16	REAL*8
	DREAL	Real part of complex	COMPLEX*16	REAL*8

3.5.2

VMS Degree-Based Trigonometric

TABLE 3-10 VMS Degree-Based Trigonometric Functions

Generic Name	Specific Names	Function	Argument Type	Result Type
SIND		Sine	-	-
	SIND		REAL *4	REAL *4
	DSIND		REAL *8	REAL *8
	QSIND		REAL *16	REAL *16
COSD		Cosine	-	-
	COSD		REAL *4	REAL *4
	DCOSD		REAL *8	REAL *8
	QCOSD		REAL *16	REAL *16
TAND		Tangent	-	-
	TAND		REAL *4	REAL *4
	DTAND		REAL *8	REAL *8
	QTAND		REAL *16	REAL *16
ASIND		Arc sine	-	-
	ASIND		REAL *4	REAL *4
	DASIND		REAL *8	REAL *8
	QASIND		REAL *16	REAL *16
ACOSD		Arc cosine	-	-
	ACOSD		REAL *4	REAL *4
	DACOSD		REAL *8	REAL *8
	QACOSD		REAL *16	REAL *16
ATAND		Arc tangent	-	-
	ATAND		REAL *4	REAL *4
	DATAND		REAL *8	REAL *8
	QATAND		REAL *16	REAL *16
ATAN2D		Arc tangent of a1/a2	-	-
	ATAN2D		REAL *4	REAL *4
	DATAN2D		REAL *8	REAL *8
	QATAN2D		REAL *16	REAL *16

3.5.3 VMS Bit-Manipulation

TABLE 3-11 VMS Bit-Manipulation Functions

Generic Name	Specific Names	Function	Argument Type	Result Type
IBITS		From a1, initial bit a2, extract a3 bits	-	-
	IIBITS		INTEGER*2	INTEGER*2
	JIBITS		INTEGER*4	INTEGER*4
	KIBITS		INTEGER*8	INTEGER*8
ISHFT		Shift a1 logically by a2 bits; if a2 positive shift left, if a2 negative shift right	-	-
	IISHFT		INTEGER*2	INTEGER*2
	JISHFT		INTEGER*4	INTEGER*4
	KISHFT		INTEGER*8	INTEGER*8
ISHFTC		In a1, circular shift by a2 places, of right a3 bits	-	-
	IISHFTC		INTEGER*2	INTEGER*2
	JISHFTC		INTEGER*4	INTEGER*4
IAND		Bitwise AND of a1, a2	-	-
	IIAND		INTEGER*2	INTEGER*2
	JIAND		INTEGER*4	INTEGER*4
IOR		Bitwise OR of a1, a2	-	-
	IIOR		INTEGER*2	INTEGER*2
	JIOR		INTEGER*4	INTEGER*4
	KIOR		INTEGER*8	INTEGER*8
IEOR		Bitwise exclusive OR of a1, a2	-	-
	IIEOR		INTEGER*2	INTEGER*2
	JIEOR		INTEGER*4	INTEGER*4
	KIEOR		INTEGER*8	INTEGER*8
NOT		Bitwise complement	-	-
	INOT		INTEGER*2	INTEGER*2
	JNOT		INTEGER*4	INTEGER*4
	KNOT		INTEGER*8	INTEGER*8

TABLE 3-11 VMS Bit-Manipulation Functions (*Continued*)

Generic Name	Specific Names	Function	Argument Type	Result Type
IBSET		In a1, set bit a2 to 1; return new a1	-	-
	IIBSET		INTEGER*2	INTEGER*2
	JIBSET		INTEGER*4	INTEGER*4
	KIBSET		INTEGER*8	INTEGER*8
BTEST		If bit a2 of a1 is 1, return .TRUE.	-	-
	BITEST		INTEGER*2	INTEGER*2
	BJTEST		INTEGER*4	INTEGER*4
	BKTEST		INTEGER*8	INTEGER*8
IBCLR		In a1, set bit a2 to 0; return new a1	-	-
	IIBCLR		INTEGER*2	INTEGER*2
	JIBCLR		INTEGER*4	INTEGER*4
	KIBCLR		INTEGER*8	INTEGER*8

3.5.4 VMS Multiple Integer Types

The possibility of multiple integer types is not addressed by the Fortran Standard. The compiler copes with their existence by treating a specific `INTEGER`-to-`INTEGER` function name (`IABS`, and so forth) as a special sort of generic. The argument type is used to select the appropriate runtime routine name, which is not accessible to the programmer.

VMS Fortran takes a similar approach, but makes the specific names available.

TABLE 3-12 VMS Integer Functions

Specific Names	Function	Argument Type	Result Type
IIABS	Absolute value	INTEGER*2	INTEGER*2
JIABS		INTEGER*4	INTEGER*4
KIABS		INTEGER*8	INTEGER*8
IMAX0	Maximum ¹	INTEGER*2	INTEGER*2
JMAX0		INTEGER*4	INTEGER*4
IMIN0	Minimum ¹	INTEGER*2	INTEGER*2
JMIN0		INTEGER*4	INTEGER*4

TABLE 3-12 VMS Integer Functions *(Continued)*

Specific Names	Function	Argument Type	Result Type
IIDIM	Positive difference ²	INTEGER*2	INTEGER*2
JIDIM		INTEGER*4	INTEGER*4
KIDIM		INTEGER*8	INTEGER*8
IMOD	Remainder of a1/a2	INTEGER*2	INTEGER*2
JMOD		INTEGER*4	INTEGER*4
IISIGN	Transfer of sign, a1 * sign(a2)	INTEGER*2	INTEGER*2
JISIGN		INTEGER*4	INTEGER*4
KISIGN		INTEGER*8	INTEGER*8

1. There must be at least two arguments.

2. The positive difference is: a1-min(a1,a2))

Index

SYMBOLS

(e**x)-1, 1-5, 1-7

A

abort, 1-11
access, 1-11
accessible documentation, xiv
alarm, 1-12
and, 1-13
arc
 cosh, 1-4, 1-7
 cosine, 1-4
 sine, 1-4
 sinh, 1-4
 tangent, 1-4
 tanh, 1-7
arguments
 command line, getarg, 1-31
array functions
 Fortran 95 intrinsics, 2-7, 2-8

B

Bessel functions, 1-5, 1-6, 1-8, 1-9
bic, 1-14
bis, 1-14
bit
 functions, 1-14
 move bits, mvbits, 1-62
bit, 1-14

bit manipulation functions
 FORTRAN 77 intrinsics, 3-12
 Fortran 95 intrinsics, 2-5
 VMS intrinsics, 3-23

bitwise

 and, 1-13
 complement, 1-13
 exclusive or, 1-13
 inclusive or, 1-13

BLAS (Basic Linear Algebra Subroutines), 2-13

C

ceiling, 1-4
change
 default directory, chdir, 1-16
 mode of a file, chmod, 1-17
character
 get a character getc, fgetc, 1-32
 put a character, putc, fputc, 1-65
character functions
 FORTRAN 77 intrinsics, 3-11
 Fortran 95 intrinsics, 2-3
chdir, 1-16
chmod, 1-17
command-line argument, getarg, 1-31
compilers, accessing, xi
conversion by long, short, 1-54
conversion functions
 FORTRAN 77 intrinsics, 3-5
ctime, convert system time to character, 1-81, 1-82
ctime64, 1-84
cube root, 1-4

current working directory, `getcwd`, 1-35

D

`d_acos(x)`, 1-7
`d_acosd(x)`, 1-7
`d_acosh(x)`, 1-7
`d_acosp(x)`, 1-7
`d_acospi(x)`, 1-7
`d_addran()`, 1-8
`d_addrans()`, 1-8
`d_asin(x)`, 1-7
`d_asind(x)`, 1-7
`d_asinh(x)`, 1-7
`d_asinp(x)`, 1-7
`d_asinpi(x)`, 1-7
`d_atan(x)`, 1-7
`d_atan2(x)`, 1-7
`d_atan2d(x)`, 1-7
`d_atan2pi(x)`, 1-7
`d_atand(x)`, 1-7
`d_atanh(x)`, 1-7
`d_atanp(x)`, 1-7
`d_atanpi(x)`, 1-7
`d_cbrt(x)`, 1-7
`d_ceil(x)`, 1-7
`d_erf(x)`, 1-7
`d_erfc(x)`, 1-7
`d_expml(x)`, 1-7
`d_floor(x)`, 1-7
`d_hypot(x)`, 1-7
`d_infinity()`, 1-7
`d_j0(x)`, 1-8
`d_j1(x)`, 1-8
`d_jn(n, x)`, 1-8
`d_lcran()`, 1-8
`d_lcrans()`, 1-8
`d_lgamma(x)`, 1-8
`d_log1p(x)`, 1-8
`d_log2(x)`, 1-8
`d_logb(x)`, 1-8
`d_max_normal()`, 1-8
`d_max_subnormal()`, 1-8
`d_min_normal()`, 1-8
`d_min_subnormal()`, 1-8
`d_nextafter(x, y)`, 1-8
`d_quiet_nan(n)`, 1-8
`d_remainder(x, y)`, 1-8

`d_rint(x)`, 1-8
`d_scalbn(x, n)`, 1-8
`d_shuftrans()`, 1-8
`d_signaling_nan(n)`, 1-8
`d_significand(x)`, 1-8
`d_sin(x)`, 1-9
`d_sincos(x, s, c)`, 1-9
`d_sincosd(x, s, c)`, 1-9
`d_sincosp(x, s, c)`, 1-9
`d_sincospi(x, s, c)`, 1-9
`d_sind(x)`, 1-9
`d_sinh(x)`, 1-9
`d_sinp(x)`, 1-9
`d_sinpi(x)`, 1-9
`d_tan(x)`, 1-9
`d_tand(x)`, 1-9
`d_tanh(x)`, 1-9
`d_tanp(x)`, 1-9
`d_tanpi(x)`, 1-9
`d_y0(x)`, *bessel*, 1-9
`d_y1(x)`, *bessel*, 1-9
`d_yn(n, x)`, 1-9
data types, 1-1
date
 and time, as characters, `fdate`, 1-24
 current date, `date`, 1-18
 `date_and_time`, 1-19
`date_and_time`, 1-19
deallocate memory by `free`, 1-61
delay execution, `alarm`, 1-12
descriptor, get file, `getfd`, 1-37
directory
 default change, `chdir`, 1-16
 get current working directory, `getcwd`, 1-35
documentation index, xiii
documentation, accessing, xiii to xv
double-precision libm functions, 1-6
`drand`, 1-71
`dtime`, 1-20

E

elapsed time, 1-20
environment variables, `getenv`, 1-36
error
 function, 1-4
 messages, `perror`, `gerror`, `ierrno`, 1-63
errors and interrupts, `longjmp`, 1-57

etime, 1-20
exception handling, 1-43, 1-47
execute an OS command, `system`, 1-74, 1-80
execution time, 1-20
existence of file, access, 1-11
exit, 1-23

F

fdate, 1-24
fgetc, 1-33
file
 descriptor, `get`, `getfd`, 1-37
 get file pointer, `getfilep`, 1-38
 mode, access, 1-11
 permissions, access, 1-11
 remove, `unlink`, 1-86
 rename, 1-72
 status, `stat`, 1-76
 status, `stat64`, 1-80
find substring, index, 1-48
floating-point
 IEEE definitions, 1-46
 IEEE exception handling, 1-43
floating-point functions
 Fortran 95 intrinsics, 2-6
floatingpoint.h header file, 1-46
floor, 1-5
flush, 1-25
fork, 1-26
FORTRAN 77
 intrinsic functions, 3-1
Fortran 95
 non-standard intrinsic functions, 2-13
 standard generic intrinsic functions, 2-1
fputc, 1-65
free, 1-61
fseek, 1-27
fseeko64, 1-29
fstat, 1-76
fstat64, 1-80
ftell, 1-27
ftello64, 1-29

G

gerror, 1-63

get
 character `getc`, `fgetc`, 1-32
 current working directory, `getcwd`, 1-35
 environment variables, `getenv`, 1-36
 file descriptor, `getfd`, 1-37
 file pointer, `getfilep`, 1-38
 group id, `getgid`, 1-41
 login name, `getlog`, 1-40
 process id, `getpid`, 1-40
 user id, `getuid`, 1-41
`getarg`, 1-31
`getc`, 1-32
`getcwd`, 1-35
`getenv`, 1-36
`getfd`, 1-37
`getfilep`, 1-38
`getgid`, 1-41
`getlog`, 1-40
`getpid`, 1-40
`getuid`, 1-41
`gmtime`, 1-81
`gmtime`, GMT, 1-84
`gmtime64`, 1-84
Greenwich Mean Time, `gmtime`, 1-81
group ID, `get`, `getgid`, 1-41

H

host name, `get`, `hostnm`, 1-42
`hostnm`, 1-42
hyperbolic cos, 1-4
hyperbolic tan, 1-6, 1-9
hypotenuse, 1-5

I

`iargc`, 1-31
`id`, process, `get`, `getpid`, 1-40
`id_finite(x)`, 1-8
`id_fp_class(x)`, 1-8
`id_rint(x)`, 1-8
`id_isinf(x)`, 1-8
`id_isnan(x)`, 1-8
`id_isnormal(x)`, 1-8
`id_issubnormal(x)`, 1-8
`id_iszero(x)`, 1-8
`id_logb(x)`, 1-8

- id_signbit(x), 1-8
- IEEE arithmetic, 1-43
- IEEE environment, 1-46
 - exception handling, 1-47
 - rounding mode, 1-47
- ieee_flags, 1-43
- ieee_handler, 1-43
- ierrno, 1-63
- IMPLICIT, 1-1
- index, 1-48
- inmax, 1-50
- inode, 1-76
- inquiry functions
 - FORTRAN 77 intrinsics, 3-13
 - Fortran 95 intrinsics, 2-1, 2-4, 2-5, 2-7
- integer
 - conversion by long, short, 1-54
- interrupts and errors, longjmp, 1-57
- intrinsic functions, 2-1, 3-1
 - FORTRAN 77, 3-1
 - Fortran 95 non-standard, 2-13
 - Fortran 95 standard, 2-1
 - interval arithmetic, 2-14
 - MPI_SIZEOF, 2-16
 - other vendor functions, 2-15
 - VMS Fortran, 3-21
- iq_finite(x), 1-10
- iq_fp_class(x), 1-10
- iq_isinf(x), 1-10
- iq_isnan(x), 1-10
- iq_isnormal(x), 1-10
- iq_issubnormal(x), 1-10
- iq_iszero(x), 1-10
- iq_logb(x), 1-10
- iq_signbit(x), 1-10
- ir_finite(x), 1-5
- ir_fp_class(x), 1-5
- ir_rint(x), 1-5
- ir_isinf(x), 1-5
- ir_isnan(x), 1-5
- ir_isnormal(x), 1-5
- ir_issubnormal(x), 1-5
- ir_iszero(x), 1-5
- ir_logb(x), 1-5
- ir_signbit(x), 1-5
- irand, 1-71
- isatty, 1-85
- isetjmp, 1-56

J

- jump, longjmp, issetjmp, 1-57

K

- kill, send signal, 1-51
- kind functions
 - Fortran 95 intrinsics, 2-4

L

- left shift, lshift, 1-13
- libm_double, 1-6
- libm_quadruple, 1-9
- libm_single, 1-3
- link, 1-52
- link to an existing file, link, 1-52
- lnblk, 1-49
- local time zone, localtime(), 1-83
- location of
 - a variable loc, 1-54
- log gamma, 1-5
- login name, get getlog, 1-39
- long, 1-54
- longjmp, 1-56
- lshift, 1-13
- lstat, 1-76
- lstat64, 1-80
- ltime, 1-81
- ltime, local time zone, 1-83
- ltime64, 1-84

M

- malloc, 1-58
- man pages, accessing, xii
- MANPATH environment variable, setting, xiii
- mathematical functions
 - FORTRAN 77 intrinsics, 3-2, 3-9
 - Fortran 95 intrinsics, 2-2
 - VMS intrinsics, 3-21
- matrix functions
 - Fortran 95 intrinsics, 2-6
- maximum
 - positive integer, inmax, 1-50

- memory
 - deallocate by free, 1-61
- memory allocation
 - FORTTRAN 77 intrinsics, 3-14
- memory dump, 1-11
- mode
 - of file, access, 1-11
- MPI_SIZEOF, 2-16
- mvbits, move bits, 1-62

N

- name
 - login, get, getlog, 1-39
 - terminal port, ttynam, 1-85
- not, 1-13
- numeric functions
 - Fortran 95 intrinsics, 2-2

O

- or, 1-13
- OS command, execute, system, 1-74, 1-80
- overflow, 1-44

P

- PATH environment variable, setting, xii
- permissions
 - access function, 1-11
- perror, 1-63
- pid, process id, getpid, 1-40
- pointer
 - get file pointer, getfilep, 1-38
- position file
 - fseek, ftell, 1-27
 - fseeko64, ftello64, 1-29
- process
 - create copy with fork function, 1-26
 - id, get, getpid, 1-40
 - send signal to, kill, 1-51
 - wait for termination, wait, 1-87
- put a character, putc, fputc, 1-65
- putc, 1-65

Q

- q_copysign(x), 1-10
- q_fabs(x), 1-10
- q_fmod(x), 1-10
- q_infinity(), 1-10
- q_max_normal(), 1-10
- q_max_subnormal(), 1-10
- q_min_normal(), 1-10
- q_min_subnormal(), 1-10
- q_nextafter(x,y), 1-10
- q_quiet_nan(n), 1-10
- q_remainder(x,y), 1-10
- q_scalbn(x,n), 1-10
- q_signaling_nan(n), 1-10
- qsort, qsort64, 1-67
- quadruple-precision libm functions, 1-9
- quick sort, qsort, 1-67

R

- r_acos(x), 1-4
- r_acosd(x), 1-4
- r_acosh(x), 1-4
- r_acosp(x), 1-4
- r_acospi(x), 1-4
- r_addran(), 1-5
- r_addrans(), 1-5
- r_asin(x), 1-4
- r_asind(x), 1-4
- r_asinh(x), 1-4
- r_asinp(x), 1-4
- r_asinpi(x), 1-4
- r_atan(x), 1-4
- r_atan2(x), 1-4
- r_atan2d(x), 1-4
- r_atan2pi(x), 1-4
- r_atand(x), 1-4
- r_atanh(x), 1-4
- r_atanp(x), 1-4
- r_atanpi(x), 1-4
- r_cbrt(x), 1-4
- r_ceil(x), 1-4
- r_erf(x), 1-4
- r_erfc(x), 1-4
- r_expml(x), 1-5
- r_floor(x), 1-5
- r_hypot(x), 1-5
- r_infinity(), 1-5

- r_j0(x), 1-5
- r_j1(x), 1-5
- r_jn(n,x), 1-5
- r_lcran(), 1-5
- r_lcrans(), 1-5
- r_lgamma(x), 1-5
- r_log1p(x), 1-5
- r_log2(x), 1-5
- r_logb(x), 1-5
- r_max_normal(), 1-5
- r_max_subnormal(), 1-5
- r_min_normal(), 1-5
- r_min_subnormal(), 1-5
- r_nextafter(x,y), 1-5
- r_quiet_nan(n), 1-5
- r_remainder(x,y), 1-5
- r_rint(x), 1-5
- r_scalbn(x,n), 1-5
- r_shufans(), 1-5
- r_signaling_nan(n), 1-5
- r_significand(x), 1-5
- r_sin(x), 1-6
- r_sincos(x,s,c), 1-6
- r_sincosd(x,s,c), 1-6
- r_sincosp(x,s,c), 1-6
- r_sincospi(x,s,c), 1-6
- r_sind(x), 1-6
- r_sinh(x), 1-6
- r_sinp(x), 1-6
- r_sinpi(x), 1-6
- r_tan(x), 1-6
- r_tand(x), 1-6
- r_tanh(x), 1-6
- r_tanp(x), 1-6
- r_tanpi(x), 1-6
- r_y0(x),bessel, 1-6
- r_y1(x),bessel, 1-6
- r_yn(n,x),bessel, 1-6
- rand, 1-71
- random
 - number, 1-5
 - values, rand, 1-71
- read
 - character getc, fgetc, 1-32
- remove a file, unlink, 1-86
- reposition file
 - fseek, ftell, 1-27
 - fseek064, ftello64, 1-29
- right shift, rshift, 1-14

- rindex, 1-49
- rounding direction, 1-44
- rshift, 1-14

S

- secnds, system time, 1-73
- send signal to process, kill, 1-51
- setbit, 1-14
- setjmp, *See* isetjmp
- shell prompts, xi
- short, 1-54
- sigfpe, 1-43
- SIGFPE handling, 1-47
- signal, 1-75
- signal a process, kill, 1-51
- sine, 1-6
- single-precision libm functions, 1-4
- 64-bit environments, 1-2
- sleep, 1-76
- sort quick, qsort, 1-67
- specific names
 - Fortran 95 intrinsics, 2-10
- stat, 1-76
- stat64, 1-80
- status
 - file, stat, 1-76
 - file, stat64, 1-80
 - termination, exit, 1-23
- substring
 - find, index, 1-48
- suspend execution for an interval, sleep, 1-76
- symbolic
 - link to an existing file, symlink, 1-52
- symlink, 1-52
- system, 1-74, 1-80
- system time
 - secnds, 1-73
 - time, 1-81
- system.inc include file, 1-2

T

- tangent, 1-6
- tarray() values for various time routines, 1-84
- terminal
 - port name, ttynam, 1-85

- terminate
 - wait for process to terminate, `wait`, 1-87
 - with status, `exit`, 1-23
 - write memory to core file, 1-11
- time, 1-20
 - `secnds`, 1-73
- `time`
 - standard version, 1-81
- `time`, get system time, 1-81
- trapping, floating-point, 1-43
- trigonometric functions
 - FORTRAN 77 intrinsics, 3-7
 - VMS intrinsics, 3-22
- `ttynam`, 1-85
- typographic conventions, x

U

- underflow, 1-44
- `unlink`, 1-86
- user ID, `get`, `getuid`, 1-41

V

- vector functions
 - Fortran 95 intrinsics, 2-6
- VMS Fortran
 - intrinsic functions, 3-21

W

- `wait`, 1-87
- write a character `putc`, `fputc`, 1-65

X

- `xknown_lib=blas`, 2-13
- `xor`, 1-13

