



Fortran Library Reference

Forte Developer 6 update 2
(Sun WorkShop 6 update 2)

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303
U.S.A. 650-960-1300

Part No. 806-7985-10
July 2001, Revision A

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

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 USA. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or 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 other countries, exclusively licensed through X/Open Company, Ltd. For Netscape™, Netscape Navigator™, and the Netscape Communications Corporation logo™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook2, Solaris, SunOS, JavaScript, SunExpress, Sun WorkShop, Sun WorkShop Professional, Sun Performance Library, Sun Performance WorkShop, Sun Visual WorkShop, and Forte are trademarks, registered trademarks, or service marks 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 an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

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

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 2001 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est 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. La notice suivante est applicable à Netscape™, Netscape Navigator™, et the Netscape Communications Corporation logo™: Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook2, Solaris, SunOS, JavaScript, SunExpress, Sun WorkShop, Sun WorkShop Professional, Sun Performance Library, Sun Performance WorkShop, Sun Visual WorkShop, et Forte sont des marques de fabrique ou des marques déposées, ou marques de service, 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.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

Sun f90/f95 est dérivé de CRAY CF90™, un produit de Cray Inc.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPENDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Important Note on New Product Names

As part of Sun's new developer product strategy, we have changed the names of our development tools from Sun WorkShop™ to Forte™ Developer products. The products, as you can see, are the same high-quality products you have come to expect from Sun; the only thing that has changed is the name.

We believe that the Forte™ name blends the traditional quality and focus of Sun's core programming tools with the multi-platform, business application deployment focus of the Forte tools, such as Forte Fusion™ and Forte™ for Java™. The new Forte organization delivers a complete array of tools for end-to-end application development and deployment.

For users of the Sun WorkShop tools, the following is a simple mapping of the old product names in WorkShop 5.0 to the new names in Forte Developer 6.

Old Product Name	New Product Name
Sun Visual WorkShop™ C++	Forte™ C++ Enterprise Edition 6
Sun Visual WorkShop™ C++ Personal Edition	Forte™ C++ Personal Edition 6
Sun Performance WorkShop™ Fortran	Forte™ for High Performance Computing 6
Sun Performance WorkShop™ Fortran Personal Edition	Forte™ Fortran Desktop Edition 6
Sun WorkShop Professional™ C	Forte™ C 6
Sun WorkShop™ University Edition	Forte™ Developer University Edition 6

In addition to the name changes, there have been major changes to two of the products.

- Forte for High Performance Computing contains all the tools formerly found in Sun Performance WorkShop Fortran and now includes the C++ compiler, so High Performance Computing users need to purchase only one product for all their development needs.
- Forte Fortran Desktop Edition is identical to the former Sun Performance WorkShop Personal Edition, except that the Fortran compilers in that product no longer support the creation of automatically parallelized or explicit, directive-based parallel code. This capability is still supported in the Fortran compilers in Forte for High Performance Computing.

We appreciate your continued use of our development products and hope that we can continue to fulfill your needs into the future.

Contents

Before You Begin	1
Typographic Conventions	2
Shell Prompts	3
Supported Platforms	3
Accessing Sun WorkShop Development Tools and Man Pages	3
Accessing Sun WorkShop Documentation	5
Accessing Related Documentation	6
Ordering Sun Documentation	7
Sending Your Comments	7
1. Fortran Library Routines	9
Data Type Considerations	9
64-Bit Environments	10
Fortran Math Functions	11
abort: Terminate and Write Core File	20
access: Check File Permissions or Existence	20
alarm: Call Subroutine After a Specified Time	21
bit: Bit Functions: and, or, ..., bit, setbit, ...	22
chdir: Change Default Directory	26
chmod: Change the Mode of a File	27

date: Get Current Date as a Character String 27
 date_and_time: Get Date and Time 28
 dtime, etime: Elapsed Execution Time 31
 exit: Terminate a Process and Set the Status 33
 fdate: Return Date and Time in an ASCII String 34
 flush: Flush Output to a Logical Unit 35
 fork: Create a Copy of the Current Process 36
 fseek, ftell: Determine Position and Reposition a File 37
 fseeko64, ftello64: Determine Position and Reposition a Large File 39
 getarg, iargc: Get Command-Line Arguments 41
 getc, fgetc: Get Next Character 42
 getcwd: Get Path of Current Working Directory 45
 getenv: Get Value of Environment Variables 46
 getfd: Get File Descriptor for External Unit Number 47
 getlog: Get User's Login Name 49
 getpid: Get Process ID 50
 getuid, getgid: Get User or Group ID of Process 50
 hostnm: Get Name of Current Host 52
 index, rindex, lnb1nk: Index or Length of Substring 60
 inmax: Return Maximum Positive Integer 63
 ioinit: Initialize I/O Properties 63
 itime: Current Time 68
 kill: Send a Signal to a Process 69
 link, symlink: Make a Link to an Existing File 69
 loc: Return the Address of an Object 71
 long, short: Integer Object Conversion 72
 longjmp, setjmp: Return to Location Set by setjmp 73
 malloc, malloc64, realloc, free: Allocate/Reallocate/Deallocate
 Memory 76

mvbits: Move a Bit Field	79
perror, gerror, ierrno: Get System Error Messages	80
putc, fputc: Write a Character to a Logical Unit	82
qsort, qsort64: Sort the Elements of a One-Dimensional Array	85
ran: Generate a Random Number Between 0 and 1	86
rand, drand, irand: Return Random Values	88
rename: Rename a File	89
secnds: Get System Time in Seconds, Minus Argument	90
sh: Fast Execution of an sh Command	91
signal: Change the Action for a Signal	92
sleep: Suspend Execution for an Interval	93
stat64, lstat64, fstat64: Get File Status	97
system: Execute a System Command	97
time, ctime, ltime, gmtime: Get System Time	98
topen, tclose, tread, ..., tstate: Tape I/O	103
ttynam, isatty: Get Name of a Terminal Port	113
unlink: Remove a File	114
wait: Wait for a Process to Terminate	115
Index	117

Tables

TABLE 1-1	Library Routines for 64-bit Environments	11
TABLE 1-2	Single-Precision Math Functions	13
TABLE 1-3	Double Precision Math Functions	16
TABLE 1-4	Quadruple-Precision <code>libm</code> Functions	19
TABLE 1-5	IEEE Arithmetic Support Routines	55
TABLE 1-6	<code>ieee_flags(action,mode,in,out)</code> Parameters and Actions	56
TABLE 1-7	<code>ieee_handler(action,in,out)</code> Parameters	57

Before You Begin

The *Fortran Library Reference* describes the routines in the Sun WorkShop™ 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 Sun WorkShop *Fortran Programming Guide*.

Other Fortran manuals in this collection include the *Fortran User's Guide*, and the *FORTTRAN 77 Language Reference*. See "Accessing Related Documentation" on page 6.

Typographic Conventions

The following table and notes describe the typographical conventions used in the manual.

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 <code>rm filename</code> .

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

$\Delta\Delta 36.001$

- FORTRAN 77 examples appear in tab format, while Fortran 95 examples appear in free format. Examples common to both Fortran 77 and 95 use tab format except where indicated.
- The FORTRAN 77 standard uses an older convention of spelling the name "FORTRAN" capitalized. Sun 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 GETENV will appear as `getenv(3F)`, implying that the man command to access this page would be: `man -s 3F getenv`
- System Administrators may install the Sun WorkShop Fortran compilers and supporting material at: `<install_point>/SUNWspr0/` where `<install_point>` is usually `/opt` for a standard install. This is the location assumed in this book.

Shell Prompts

Shell	Prompt
C shell	%
Bourne shell and Korn shell	\$
C shell, Bourne shell, and Korn shell superuser	#

Supported Platforms

This Sun WorkShop™ release of the Fortran compilers supports only versions 2.6, 7, and 8 of the Solaris™ SPARC™ Platform Edition.

Accessing Sun WorkShop Development Tools and Man Pages

The Sun WorkShop product components and man pages are not installed into the standard `/usr/bin/` and `/usr/share/man` directories. To access the Sun WorkShop compilers and tools, you must have the Sun WorkShop component directory in your `PATH` environment variable. To access the Sun WorkShop man pages, you must have the Sun WorkShop 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 release, see the *Sun WorkShop 6 update 2 Installation Guide* or your system administrator.

Note – The information in this section assumes that your Sun WorkShop 6 update 2 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 Sun WorkShop Compilers and Tools

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

To Determine If You Need to Set Your `PATH` Environment Variable

1. **Display the current value of the `PATH` variable by typing:**

```
% echo $PATH
```

2. **Review the output for a string of paths containing `/opt/SUNWspro/bin/`.**

If you find the path, your `PATH` variable is already set to access Sun WorkShop 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 Sun WorkShop 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 Sun WorkShop Man Pages

Use the following steps to determine whether you need to change your `MANPATH` variable to access the Sun WorkShop man pages.

To Determine If You Need to Set Your MANPATH Environment Variable

1. Request the workshop man page by typing:

```
% man workshop
```

2. Review the output, if any.

If the workshop(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 Sun WorkShop 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 Sun WorkShop Documentation

You can access Sun WorkShop 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.

Point your Netscape™ Communicator 4.0 or compatible Netscape version browser to the following 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.

- Manuals are available from the docs.sun.comsm Web site.

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.

Accessing Related Documentation

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

Document Collection	Document Title	Description
Forte™ for High Performance Computing Collection	<i>Fortran Programming Guide</i>	Discusses issues relating to input/output, libraries, program analysis, debugging, and performance.
	<i>Fortran User's Guide</i>	Provides information on command-line options and how to use the compilers.
	<i>FORTTRAN 77 Language Reference</i>	Provides a complete language reference to Sun FORTRAN 77.
Numerical Computation Guide Collection	<i>Numerical Computation Guide</i>	Describes issues regarding the numerical accuracy of floating-point computations.
Solaris 8 Reference Manual Collection	See the titles of man page sections.	Provides information about the Solaris operating environment.
Solaris 8 Software Developer Collection	<i>Linker and Libraries Guide</i>	Describes the operations of the Solaris link-editor and runtime linker.
Solaris 8 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.

Ordering Sun Documentation

You can order product documentation directly from Sun through the `docs.sun.com` Web site or from Fatbrain.com, an Internet bookstore. You can find the Sun Documentation Center on Fatbrain.com at the following URL:

`http://www.fatbrain.com/documentation/sun`

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 alphabetically. See the *FORTRAN 77 Language Reference* for details on Fortran 77 and VMS intrinsic functions. 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.

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 `f77(1)` and `f95(1)` man pages 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 `f77` and `f95` compilers is described in the *Fortran User's Guide*, the *Fortran Programming Guide*, and the `f77(1)` and `f95(1)` man pages.

64-Bit Environments

Compiling a program to run in a 64-bit operating environment (that is, compiling with `-xarch=v9` or `v9a` 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 page 76), 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

Library Routines		
<code>malloc64</code>	Allocate memory and return a pointer	page 76
<code>fseeko64</code>	Reposition a large file	page 39
<code>ftello64</code>	Determine position of a large file	page 39
<code>stat64</code> , <code>fstat64</code> , <code>lstat64</code>	Determine status of a file	page 97
<code>time64</code> , <code>ctime64</code> , <code>gmtime64</code> , <code>ltime64</code>	Get system time, convert to character or dissected	page 98
<code>qsort64</code>	Sort the elements of an array	page 85

Fortran Math Functions

The following functions and subroutines are part of the Fortran math libraries. They are available to all programs compiled with `f77` and `f95`. Some routines are intrinsics and return the same data type (single precision, double precision, or quad precision) as their argument. The rest are non-intrinsics that take a specific data type as an argument and return the same. These 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_quaduple(3F)`, for more information about these libraries.

Intrinsic Math Functions

Here is a list of *intrinsic* math functions. You need not put them in a type statement. These functions take single, double, or quad precision data as arguments and return the same.

<code>sqrt(x)</code>	<code>asin(x)</code>	<code>cosd(x)</code>
<code>log(x)</code>	<code>acos(x)</code>	<code>asind(x)</code>
<code>log10(x)</code>	<code>atan(x)</code>	<code>acosd(x)</code>
<code>exp(x)</code>	<code>atan2(x,y)</code>	<code>atand(x)</code>
<code>x**y</code>	<code>sinh(x)</code>	<code>atan2d(x,y)</code>
<code>sin(x)</code>	<code>cosh(x)</code>	<code>aint(x)</code>
<code>cos(x)</code>	<code>tanh(x)</code>	<code>anint(x)</code>
<code>tan(x)</code>	<code>sind(x)</code>	<code>nint(x)</code>

The functions `sind(x)`, `cosd(x)`, `asind(x)`, `acosd(x)`, `atand(x)`, `atan2d(x,y)` are not part of the Fortran standard.

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

<code>r_acos(x)</code>	REAL	Function	arc cosine
<code>r_acosd(x)</code>	REAL	Function	--
<code>r_acosh(x)</code>	REAL	Function	arc cosh
<code>r_acosp(x)</code>	REAL	Function	--
<code>r_acospi(x)</code>	REAL	Function	--
<code>r_atan(x)</code>	REAL	Function	arc tangent
<code>r_atand(x)</code>	REAL	Function	--
<code>r_atanh(x)</code>	REAL	Function	arc tanh
<code>r_atanp(x)</code>	REAL	Function	--
<code>r_atanpi(x)</code>	REAL	Function	--
<code>r_asin(x)</code>	REAL	Function	arc sine
<code>r_asind(x)</code>	REAL	Function	--
<code>r_asinh(x)</code>	REAL	Function	arc sinh
<code>r_asinp(x)</code>	REAL	Function	--
<code>r_asinpi(x)</code>	REAL	Function	--
<code>r_atan2((y, x)</code>	REAL	Function	arc tangent
<code>r_atan2d(y, x)</code>	REAL	Function	--
<code>r_atan2pi(y, x)</code>	REAL	Function	--
<code>r_cbrt(x)</code>	REAL	Function	cube root
<code>r_ceil(x)</code>	REAL	Function	ceiling
<code>r_copysign(x, y)</code>	REAL	Function	--
<code>r_cos(x)</code>	REAL	Function	cosine
<code>r_cosd(x)</code>	REAL	Function	--
<code>r_cosh(x)</code>	REAL	Function	hyperb cos
<code>r_cosp(x)</code>	REAL	Function	--
<code>r_cospi(x)</code>	REAL	Function	--
<code>r_erf(x)</code>	REAL	Function	err function
<code>r_erfc(x)</code>	REAL	Function	--
<code>r_expml(x)</code>	REAL	Function	$(e^*x)-1$
<code>r_floor(x)</code>	REAL	Function	floor
<code>r_hypot(x, y)</code>	REAL	Function	hypotenuse
<code>r_infinity()</code>	REAL	Function	--
<code>r_j0(x)</code>	REAL	Function	Bessel
<code>r_j1(x)</code>	REAL	Function	--
<code>r_jn(x)</code>	REAL	Function	--

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

ir_finite(x)	INTEGER	Function	--
ir_fp_class(x)	INTEGER	Function	--
ir_ilogb(x)	INTEGER	Function	--
ir_rint(x)	INTEGER	Function	--
ir_isinf(x)	INTEGER	Function	--
ir_isnan(x)	INTEGER	Function	--
ir_isnormal(x)	INTEGER	Function	--
ir_issubnormal(x)	INTEGER	Function	--
ir_iszero(x)	INTEGER	Function	--
ir_signbit(x)	INTEGER	Function	--
r_addran()	REAL	Function	random
r_addrans(x, p, l, u)	n/a	Subroutine	number
r_lcran()	REAL	Function	generators
r_lcrans(x, p, l, u)	n/a	Subroutine	
r_shufrans(x, p, l, u)	n/a	Subroutine	
r_lgamma(x)	REAL	Function	log gamma
r_logb(x)	REAL	Function	--
r_loglp(x)	REAL	Function	--
r_log2(x)	REAL	Function	--
r_max_normal()	REAL	Function	
r_max_subnormal()	REAL	Function	
r_min_normal()	REAL	Function	
r_min_subnormal()	REAL	Function	
r_nextafter(x, y)	REAL	Function	
r_quiet_nan(n)	REAL	Function	
r_remainder(x, y)	REAL	Function	
r_rint(x)	REAL	Function	
r_scalb(x, y)	REAL	Function	
r_scalbn(x, n)	REAL	Function	
r_signaling_nan(n)	REAL	Function	
r_significand(x)	REAL	Function	
r_sin(x)	REAL	Function	sine
r_sind(x)	REAL	Function	--
r_sinh(x)	REAL	Function	hyperb sin
r_sinp(x)	REAL	Function	--
r_sinpi(x)	REAL	Function	--

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

r_sincos(x, s, c)	n/a	Subroutine	sine & cosine
r_sincosd(x, s, c)	n/a	Subroutine	--
r_sincosp(x, s, c)	n/a	Subroutine	--
r_sincospi(x, s, c)	n/a	Subroutine	--
r_tan(x)	REAL	Function	tangent
r_tand(x)	REAL	Function	--
r_tanh(x)	REAL	Function	hyperb tan
r_tanp(x)	REAL	Function	--
r_tanpi(x)	REAL	Function	--
r_y0(x)	REAL	Function	bessel
r_y1(x)	REAL	Function	--
r_yn(n, x)	REAL	Function	--

- 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*.

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

<code>d_acos(x)</code>	DOUBLE PRECISION	Function	arc cosine
<code>d_acosd(x)</code>	DOUBLE PRECISION	Function	--
<code>d_acosh(x)</code>	DOUBLE PRECISION	Function	arc cosh
<code>d_acosp(x)</code>	DOUBLE PRECISION	Function	--
<code>d_acospi(x)</code>	DOUBLE PRECISION	Function	--
<code>d_atan(x)</code>	DOUBLE PRECISION	Function	arc tangent
<code>d_atand(x)</code>	DOUBLE PRECISION	Function	--
<code>d_atanh(x)</code>	DOUBLE PRECISION	Function	arc tanh
<code>d_atanp(x)</code>	DOUBLE PRECISION	Function	--
<code>d_atanpi(x)</code>	DOUBLE PRECISION	Function	--
<code>d_asin(x)</code>	DOUBLE PRECISION	Function	arc sine
<code>d_asind(x)</code>	DOUBLE PRECISION	Function	--
<code>d_asinh(x)</code>	DOUBLE PRECISION	Function	arc sinh
<code>d_asinp(x)</code>	DOUBLE PRECISION	Function	--
<code>d_asinpi(x)</code>	DOUBLE PRECISION	Function	--
<code>d_atan2((y, x)</code>	DOUBLE PRECISION	Function	arc tangent
<code>d_atan2d(y, x)</code>	DOUBLE PRECISION	Function	--
<code>d_atan2pi(y, x)</code>	DOUBLE PRECISION	Function	--
<code>d_cbrt(x)</code>	DOUBLE PRECISION	Function	cube root
<code>d_ceil(x)</code>	DOUBLE PRECISION	Function	ceiling
<code>d_copysign(x, x)</code>	DOUBLE PRECISION	Function	--
<code>d_cos(x)</code>	DOUBLE PRECISION	Function	cosine
<code>d_cosd(x)</code>	DOUBLE PRECISION	Function	--
<code>d_cosh(x)</code>	DOUBLE PRECISION	Function	hyperb cos
<code>d_cosp(x)</code>	DOUBLE PRECISION	Function	--
<code>d_cospi(x)</code>	DOUBLE PRECISION	Function	--
<code>d_erf(x)</code>	DOUBLE PRECISION	Function	error func
<code>d_erfc(x)</code>	DOUBLE PRECISION	Function	--
<code>d_expml(x)</code>	DOUBLE PRECISION	Function	$(e^*x)-1$
<code>d_floor(x)</code>	DOUBLE PRECISION	Function	floor
<code>d_hypot(x, y)</code>	DOUBLE PRECISION	Function	hypotenuse
<code>d_infinity()</code>	DOUBLE PRECISION	Function	--
<code>d_j0(x)</code>	DOUBLE PRECISION	Function	Bessel
<code>d_j1(x)</code>	DOUBLE PRECISION	Function	--
<code>d_jn(x)</code>	DOUBLE PRECISION	Function	--

TABLE 1-3 Double Precision Math Functions (Continued)

<code>id_finite(x)</code>	INTEGER	Function	
<code>id_fp_class(x)</code>	INTEGER	Function	
<code>id_ilogb(x)</code>	INTEGER	Function	
<code>id_rint(x)</code>	INTEGER	Function	
<code>id_isinf(x)</code>	INTEGER	Function	
<code>id_isnan(x)</code>	INTEGER	Function	
<code>id_isnormal(x)</code>	INTEGER	Function	
<code>id_issubnormal(x)</code>	INTEGER	Function	
<code>id_iszero(x)</code>	INTEGER	Function	
<code>id_signbit(x)</code>	INTEGER	Function	
<code>d_addran()</code>	DOUBLE PRECISION	Function	random
<code>d_addrans(x, p, l, u)</code>	n/a	Subroutine	number
<code>d_lcran()</code>	DOUBLE PRECISION	Function	generators
<code>d_lcrans(x, p, l, u)</code>	n/a	Subroutine	
<code>d_shufrans(x, p, l,u)</code>	n/a	Subroutine	
<code>d_lgamma(x)</code>	DOUBLE PRECISION	Function	log gamma
<code>d_logb(x)</code>	DOUBLE PRECISION	Function	--
<code>d_loglp(x)</code>	DOUBLE PRECISION	Function	--
<code>d_log2(x)</code>	DOUBLE PRECISION	Function	--
<code>d_max_normal()</code>	DOUBLE PRECISION	Function	
<code>d_max_subnormal()</code>	DOUBLE PRECISION	Function	
<code>d_min_normal()</code>	DOUBLE PRECISION	Function	
<code>d_min_subnormal()</code>	DOUBLE PRECISION	Function	
<code>d_nextafter(x, y)</code>	DOUBLE PRECISION	Function	
<code>d_quiet_nan(n)</code>	DOUBLE PRECISION	Function	
<code>d_remainder(x, y)</code>	DOUBLE PRECISION	Function	
<code>d_rint(x)</code>	DOUBLE PRECISION	Function	
<code>d_scalb(x, y)</code>	DOUBLE PRECISION	Function	
<code>d_scalbn(x, n)</code>	DOUBLE PRECISION	Function	
<code>d_signaling_nan(n)</code>	DOUBLE PRECISION	Function	
<code>d_significand(x)</code>	DOUBLE PRECISION	Function	
<code>d_sin(x)</code>	DOUBLE PRECISION	Function	sine
<code>d_sind(x)</code>	DOUBLE PRECISION	Function	--
<code>d_sinh(x)</code>	DOUBLE PRECISION	Function	hyper sine
<code>d_sinp(x)</code>	DOUBLE PRECISION	Function	--
<code>d_sinpi(x)</code>	DOUBLE PRECISION	Function	--

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

d_sincos(x, s, c)	n/a	Subroutine	sine & cosine
d_sincosd(x, s, c)	n/a	Subroutine	--
d_sincosp(x, s, c)	n/a	Subroutine	--
d_sincospi(x, s, c)	n/a	Subroutine	
d_tan(x)	DOUBLE PRECISION	Function	tangent
d_tand(x)	DOUBLE PRECISION	Function	--
d_tanh(x)	DOUBLE PRECISION	Function	hyperb tan
d_tanp(x)	DOUBLE PRECISION	Function	--
d_tanpi(x)	DOUBLE PRECISION	Function	--
d_y0(x)	DOUBLE PRECISION	Function	bessel
d_y1(x)	DOUBLE PRECISION	Function	--
d_yn(n, x)	DOUBLE PRECISION	Function	--

- 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*.

Quad-Precision Functions

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

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

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

- 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*.

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.

access: Check File Permissions or Existence

The function is called by:

INTEGER*4 access status = access (name, mode)			
<i>name</i>	character	Input	File name
<i>mode</i>	character	Input	Permissions
Return value	INTEGER*4	Output	<i>status</i> =0: OK <i>status</i> >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
```

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.

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>iand(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 79, and the chapter on Intrinsic Functions in the *FORTRAN 77 Reference Manual*.

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)',
&          6x 'not(4)'/4o12.11)
      end
demo% f77 -silent tandornot.f
demo% a.out
      and(7,4)      or(7,4)      xor(7,4)      not(4)
000000000004 000000000007 000000000003 37777777773
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% f77 -silent tlrshift.f
demo% a.out
lshift(7,1) rshift(4,1)
000000000016 00000000002
demo%

```

Usage: bic, bis, bit, setbit

```
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
```

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.

chmod: Change the Mode of a File

The function is called by:

INTEGER*4 chmod <i>n</i> = chmod(<i>name</i> , <i>mode</i>)			
<i>name</i>	character	Input	Path name
<i>mode</i>	character	Input	Anything recognized by <i>chmod</i> (1), such as <i>o-w</i> , 444, 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.

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% f77 -silent dat1.f
"dat.f", line 2: Warning: Subroutine "date" is not safe after
year 2000; use "date_and_time" instead
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-Jul-98
demo%
```

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

date_and_time: Get Date and Time

This is a FORTRAN 77 version of the 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 <code>date_and_time(date, time, zone, values)</code>			
<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
```

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.

The versions of `dtime` and `etime` used by Fortran 77 return times produced by the runtime system's high resolution clock. The actual resolution depends on the system platform. The resolutions of the clocks on current platforms range between one nanosecond and one microsecond.

Versions of `dtime` 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.

dtime: Elapsed Time Since the Last dtime Call

For `dtime`, the elapsed time is:

- First call: elapsed time since start of execution
- Subsequent calls: elapsed time since the last call to `dtime`
- 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 `dtime` 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:

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

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% f77 -silent tdmtime.f
demo% a.out
elapsed:  0., user:  0., sys:  0.
elapsed:  0.180000, user:  6.00000E-02, sys:  0.120000
demo%
```

etime: Elapsed Time Since Start of Execution

For `etime`, the elapsed time is:

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

Here is how Fortran decides single processor or multiple processor:

For a parallelized Fortran program linked with `libF77_mt`, if the environment variable `PARALLEL` is:

- Undefined, the current run is single processor.
- Defined and in the range 1, 2, 3, ..., the current run is multiple processor.
- Defined, but some value other than 1, 2, 3, ..., the results are unpredictable.

The function is called by:

<code>e = etime(tarray)</code>			
<i>tarray</i>	real(2)	Output	<i>e</i> = -1.0: Error: <i>tarray</i> values are undefined. <i>e</i> ≠ -1.0: 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	real	Output	<i>e</i> = -1.0: Error <i>e</i> ≠ -1.0: 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% f77 -silent tetime.f
demo% a.out
elapsed:  0.190000, user:    6.00000E-02, sys:    0.130000
demo%
```

See also `times(2)`, `f77(1)`, and the *Fortran Programming Guide*.

exit: Terminate a Process and Set the Status

The subroutine is called by:

call exit(status)		
status	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)`.

fdate: Return Date and Time in an ASCII String

The subroutine or function is called by:

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

or:

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

`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	<code>n = 0</code> no error <code>n > 0</code> 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)`.

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)`.

`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.

`fseek`: Reposition a File on a Logical Unit

The function is called by:

INTEGER*4 <code>fseek</code> <code>n = fseek(lunit, offset, from)</code>			
<i>lunit</i>	INTEGER*4	Input	Open logical unit
<i>offset</i>	INTEGER*4 or 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 <code>-xarch=v9</code> . If a literal constant is supplied, it must be a 64-bit constant, for example: <code>100_8</code>		
<i>from</i>	INTEGER*4	Input	0=Beginning of file 1=Current position 2=End of file
Return value	INTEGER*4	Output	<code>n=0</code> : OK; <code>n>0</code> : 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

```

`ftell`: Return Current Position of File

The function is called by:

<pre> INTEGER*4 ftell n = ftell(lunit) </pre>			
<i>lunit</i>	INTEGER*4	Input	Open logical unit
Return value	INTEGER*4 or INTEGER*8	Output	$n \geq 0$: n =Offset in bytes from start of file $n < 0$: n =System error code
<p>An INTEGER*8 offset value is returned when compiling for a 64-bit environment, such as Solaris 7 or 8, with <code>-xarch=v9</code>. <code>ftell</code> and variables receiving this return value should be declared INTEGER*8.</p>			

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)`.

`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 on Solaris 2.6, 7, or 8 operating environments. (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.

`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
```

`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'
...
```

getarg, iargc: Get Command-Line Arguments

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

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

iargc: Get the Number of Command-Line Arguments

The function is called by:

<code>m = 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% f77 -silent yarg.f
demo% a.out *.f
1 first.f
2 yarg.f
```

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

`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.

`getc`: Get Next Character from `stdin`

The function is called by:

<code>INTEGER*4 getc</code> <code>status = getc(char)</code>			
<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
^D            terminated by a CONTROL-D.
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
-1 012       Next attempt to read returns CONTROL-D
demo%
```

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

fgetc: Get Next Character from Specified Logical Unit

The function is called by:

<pre>INTEGER*4 fgetc status = fgetc(lunit, char)</pre>			
<i>lunit</i>	INTEGER*4	Input	Logical unit
<i>char</i>	character	Output	Next character
Return value	INTEGER*4	Output	<i>status</i> =-1: End of File <i>status</i> >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)`.

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>`.

getenv: Get Value of Environment Variables

The subroutine is called by:

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

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

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

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

```
character*18 evalue
call getenv( 'SHELL', evalue )
write(*,*) "'", evalue, "'"
end
```

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

getfd: Get File Descriptor for External Unit Number

The function is called by:

<code>INTEGER*4 getfd</code> <code>fildes = getfd(unitn)</code>			
<i>unitn</i>	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)`.

getfilep: Get File Pointer for External Unit Number

The function is:

<i>irtn = c_read(getfilep(unitn), inbyte, 1)</i>			
<i>c_read</i>	C function	Input	User's own C function. See example.
<i>unitn</i>	INTEGER*4	Input	External unit number.
<i>getfilep</i>	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% f77 tgetfilepC.o tgetfilepF.f
tgetfileF.f:
MAIN:
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)`.

getlog: Get User's Login Name

The subroutine is called by:

call <code>getlog(name)</code>			
<i>name</i>	character* <i>n</i>	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).

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).

getuid, getgid: Get User or Group ID of Process

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

getuid: Get User ID of the Process

The function is called by:

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

getgid: Get Group ID of the Process

The function is called by:

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

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

<pre>INTEGER*4 getuid, getpid, gid, uid uid = getuid() gid = getgid() write(*,*) uid, gid end</pre>

See also: `getuid(2)`.

hostnm: Get Name of Current Host

The function is called by:

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

Example: `hostnm()`:

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

See also `gethostname(2)`.

idate: Return Current Date

idate has two versions:

- *Standard*—Put the current system date into an integer array: day, month, and year.
- *VMS*—Put the current system date into three integer variables: month, day, and year. This version is not “Year 2000 Safe”.

The `-1V77` compiler option accesses the VMS library and links the VMS versions of both `time()` and `idate()`; otherwise, the linker accesses the standard versions. (VMS versions of library routines are only available with `f77` through the `-1V77` library option, and not with `f95`).

The standard version 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% f77 -silent tidate.f
demo% a.out
      The date is:10 8 1998
demo%
```

The VMS idate() subroutine is called by:

call idate(<i>m</i> , <i>d</i> , <i>y</i>)		<i>VMS Version</i>	
<i>m</i>	INTEGER*4	Output	Month (1 - 12)
<i>d</i>	INTEGER*4	Output	Day (1 - 7)
<i>y</i>	INTEGER*4	Output	Year (1 - 99) <i>Not year 2000 safe!</i>

Using the VMS `idate()` routine will cause a warning message at link time and the first time the routine is called in execution.

Note – The VMS version of the `idate()` 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 `idate()` 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.

Example: `idate` (VMS version):

```
demo% cat titime.f
      INTEGER*4 m, d, y
      call idate ( m, d, y )
      write (*, "(' The date is: ',3i5)" ) m, d, y
      end
demo% f77 -silent tidateV.f -lv77
"titime.f", line 2: Warning: Subroutine "idate" is not safe after
year 2000; use "date_and_time" instead
demo% a.out
Computing time differences using the 2 digit year from subroutine
idate is not safe after year 2000.
The date is:      7   10   98
```

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

<code>ieeeer = ieee_flags(action, mode, in, out)</code>		
<code>ieeeer = ieee_handler(action, exception, hdl)</code>		
<code>ieeeer = sigfpe(code, hdl)</code>		
<i>action</i>	character	Input
<i>code</i>	<code>sigfpe_code_type</code>	Input
<i>mode</i>	character	Input
<i>in</i>	character	Input
<i>exception</i>	character	Input
<i>hdl</i>	<code>sigfpe_handler_type</code>	Input
<i>out</i>	character	Output
Return value	<code>INTEGER*4</code>	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, in, out</i> , unused; returns 0		
<i>action</i> = 'clear' clear <i>mode, in</i> <i>out</i> is unused; returns 0	<i>mode</i> = 'direction'		
	<i>mode</i> = 'precision' (on x86 platforms only)		
	<i>mode</i> = 'exception'	<i>in</i> = 'inexact' 'division' 'underflow' 'overflow' 'invalid' 'all' 'common'	<i>or</i> <i>or</i> <i>or</i> <i>or</i> <i>or</i> <i>or</i>
<i>action</i> = 'set' set floating-point <i>mode,in</i> <i>out</i> is unused; returns 0	<i>mode</i> = 'direction'	<i>in</i> = 'nearest' 'tozero' 'positive' 'negative'	<i>or</i> <i>or</i> <i>or</i> <i>or</i>
	<i>mode</i> = 'precision' (on x86 only)	<i>in</i> = 'extended' 'double' 'single'	<i>or</i> <i>or</i> <i>or</i>
	<i>mode</i> = 'exception'	<i>in</i> = 'inexact' 'division' 'underflow' 'overflow' 'invalid' 'all' 'common'	<i>or</i> <i>or</i> <i>or</i> <i>or</i> <i>or</i> <i>or</i>

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

<i>action</i> = 'get'	<i>mode</i> =	<i>out</i> = 'nearest'	<i>or</i>
test <i>mode</i> settings	'direction'	'tozero'	<i>or</i>
<i>in</i> , <i>out</i> may be blank or one of the settings to test		'positive'	<i>or</i>
returns the current setting depending on <i>mode</i> , or 'not available'	<i>mode</i> =	<i>out</i> = 'extended'	<i>or</i>
The function returns 0 or the current exception flags if <i>mode</i> = 'exception'	'precision' (on x86 only)	'double'	<i>or</i>
		'single'	
	<i>mode</i> =	<i>out</i> = 'inexact'	<i>or</i>
	'exception'	'division'	<i>or</i>
		'underflow'	<i>or</i>
		'overflow'	<i>or</i>
		'invalid'	<i>or</i>
		'all'	<i>or</i>
		'common'	

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

<i>action</i> = 'clear'	<i>in</i> = 'inexact'	<i>or</i>
clear user exception handling 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 handling of <i>in</i> ; <i>out</i> is address of handler routine, or SIGFPE_DEFAULT, or SIGFPE_ABORT, or SIGFPE_IGNORE defined in <code>f77/f77_floating_point.h</code>	'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 (f77):

```
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)`, `f77_floatingpoint(3F)`, `ieee_flags(3M)`, and `ieee_handler(3M)`.

f77_floatingpoint.h: Fortran IEEE Definitions

The header file `f77_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 77 source program as follows:

```
#include "f77_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.

Fortran 95 programs should include the file `floatingpoint.h` instead.

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(3M)` and `f77_ieee_environment(3F)`.

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

These functions search through a character string:

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

`index` has the following forms:

index: First Occurrence of a Substring in a String

The index is an intrinsic function called by:

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

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).

rindex: Last Occurrence of a Substring in a String

The function is called by:

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

lnblnk: Last Nonblank in a String

The function is called by:

<code>n = lnblnk(a1)</code>			
<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% f77 -silent 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.

imax: Return Maximum Positive Integer

The function is called by:

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

Example: imax:

```
INTEGER*4 imax, m
m = imax()
write(*,*) m
end
demo% f77 -silent tinmax.f
demo% a.out
      2147483647
demo%
```

See also `libm_single(3F)` and `libm_double(3F)`. See also the intrinsic function `ephuge()` described in the *FORTRAN 77 Language Reference Manual*.

ioinit: Initialize I/O Properties

The IOINIT routine (FORTRAN 77 only) establishes properties of file I/O for files opened after the call to IOINIT. The file I/O properties that IOINIT controls are as follows:

- Carriage control: Recognize carriage control on any logical unit.
- Blanks/zeros: Treat blanks in input data fields as blanks or zeroes.
- File position: Open files at beginning or at end-of-file.
- Prefix: Find and open files named *prefixNN*, $0 \leq NN \leq 19$.

IOINIT does the following:

- Initializes global parameters specifying `f77` file I/O properties

- Opens logical units 0 through 19 with the specified file I/O properties—attaches externally defined files to logical units at runtime

Persistence of File I/O Properties

The file I/O properties apply as long as the connection exists. If you close the unit, the properties no longer apply. The exception is the preassigned units 5 and 6, to which carriage control and blanks/zeroes apply at any time.

Internal Flags

IOINIT uses labeled common to communicate with the runtime I/O system. It stores internal flags in the equivalent of the following labeled common block:

```
INTEGER*2 IEOF, ICTL, IBZR  
COMMON /__IOIFLG/ IEOF, ICTL, IBZR ! Not in user name space
```

In earlier releases (prior to 3.0.1) the labeled common block was named IOIFLG. The name changed subsequently to `__IOIFLG` to prevent conflicts with any user-defined common blocks.

Source Code

Some user needs are not satisfied with a generic version of IOINIT. You can find the Fortran 77 source code at:

```
<install>/SUNWspro/<release>/src/ioint.f
```

where `<install>` is usually `/opt` for a standard installation of the Sun WorkShop software packages, and `<release>` path changes with every release of the compilers.

Usage: `ioinit`

The `ioinit` subroutine is called by:

<code>call ioinit (<i>cctl</i>, <i>bzro</i>, <i>apnd</i>, <i>prefix</i>, <i>vrbose</i>)</code>			
<i>cctl</i>	logical	Input	True: Recognize carriage control, all formatted output (except unit 0)
<i>bzro</i>	logical	Input	True: Treat trailing and imbedded blanks as zeroes.
<i>apnd</i>	logical	Input	True: Open files at EoF. Append.
<i>prefix</i>	character*n	Input	Nonblank: For unit <i>NN</i> , seek and open file <i>prefixNN</i>
<i>vrbose</i>	logical	Input	True: Report <code>ioinit</code> activity as it happens

See also `getarg(3F)` and `getenv(3F)`.

Restrictions

Note the following restrictions:

- *prefix* can be no longer than 30 characters.
- A path name associated with an environment name can be no longer than 255 characters.

Description of Arguments

These are the arguments for `ioinit`.

cctl (Carriage Control)

By default, carriage control is not recognized on any logical unit. If *cctl* is `.TRUE.`, then carriage control is recognized on formatted output to all logical units, except unit 0, the diagnostic channel. Otherwise, the default is restored.

bzro (Blanks)

By default, trailing and embedded blanks in input data fields are ignored. If *bzro* is `.TRUE.`, then such blanks are treated as zeros. Otherwise, the default is restored.

apnd (Append)

By default, all files opened for sequential access are positioned at their beginning. It is sometimes necessary or convenient to open at the end-of-file, so that a write will append to the existing data. If *apnd* is `.TRUE.`, then files opened subsequently on any logical unit are positioned at their end upon opening. A value of `.FALSE.` restores the default behavior.

prefix (Automatic File Connection)

If the argument *prefix* is a nonblank string, then names of the form *prefixNN* are sought in the program environment. The value associated with each such name found is used to open the logical unit *NN* for formatted sequential access.

This search and connection is provided only for *NN* between 0 and 19, inclusive. For *NN* > 19, nothing is done; see “Source Code” on page 64.

vrbose (IOINIT Activity)

If the argument *vrbose* is `.TRUE.`, then IOINIT reports on its own activity.

Example: The program `myprogram` has the following `ioinit` call:

```
call ioinit( .true., .false., .false., 'FORT', .false.)
```

You can assign file name in at least two ways.

In `sh`:

```
demo$ FORT01=mydata  
demo$ FORT12=myresults  
demo$ export FORT01 FORT12  
demo$ myprogram
```

In csh:

```
demo% setenv FORT01 mydata
demo% setenv FORT12 myresults
demo% myprogram
```

With either shell, the `ioinit` call in the above example gives these results:

- Open logical unit 1 to the file, `mydata`.
- Open logical unit 12 to the file, `myresults`.
- Both files are positioned at their beginning.
- Any formatted output has column 1 removed and interpreted as carriage control.
- Embedded and trailing blanks are to be ignored on input.

Example: `ioinit()`—list and compile:

```
demo% cat tioint.f
character*3 s
call ioint( .true., .false., .false., 'FORT', .false.)
do i = 1, 2
    read( 1, '(a3,i4)') s, n
    write( 12, 10 ) s, n
end do
10    format(a3,i4)
end
demo% cat tioint.data
abc 123
PDQ 789
demo% f77 -silent tioint.f
demo%
```

You can set environment variables as follows, using either `sh` or `csh`:

`ioinit()`—`sh`:

```
demo$ FORT01=tioint.data
demo$ FORT12=tioint.au
demo$ export FORT01 FORT12
demo$
```

ioinit()—csh:

```
demo% setenv FORT01 tioinit.data
demo% setenv FORT12 tioinit.au
```

ioinit()—Run and test:

```
demo% a.out
demo% cat tioinit.au
abc 123
PDQ 789
```

itime: Current Time

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

call itime(<i>iarray</i>)			
<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% f77 -silent titime.f
demo% a.out
      The time is: 15 42 35
```

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

kill: Send a Signal to a Process

The function is called by:

<i>status</i> = kill(<i>pid</i> , <i>signum</i>)			
<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)`.

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

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% f77 -silent 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%
```

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 tsymlink.f
      character*34 name1/'tlink.db.data.1'/, name2/'data1'/
      INTEGER*4 status, symlink
      status = symlink( name1, name2 )
      if ( status .ne. 0 ) stop 'symlink: error'
      end
demo% f77 -silent tsymlink.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>`.

loc: Return the Address of an Object

This intrinsic function is called by:

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

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.

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

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.

`longjmp`, `isetjmp`: Return to Location Set by `isetjmp`

`isetjmp` sets a location for `longjmp`; `longjmp` returns to that location.

`isetjmp`: Set the Location for `longjmp`

This intrinsic function is called by:

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

longjmp: Return to the Location Set by iset jmp

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 <code>iset jmp</code> . In 64-bit environments it must be declared INTEGER*8
<i>ival</i>	INTEGER*4	Output	<i>ival</i> = 0 if <code>iset jmp</code> is called explicitly <i>ival</i> ≠ 0 if <code>iset jmp</code> is called through <code>long jmp</code>

Description

The `iset jmp` and `long jmp` routines are used to deal with errors and interrupts encountered in a low-level routine of a program. They are `f77` intrinsics.

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

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

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

The integer expression *ival* returned from `iset jmp` is zero if `long jmp` is not called, and nonzero if `long jmp` 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
```

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)`.

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 are 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`.

Fortran 77 programs can 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.

Allocate Memory: `malloc`, `malloc64`

The `malloc()` function is called by:

<code>k = malloc(n)</code>			
<code>n</code>	<code>INTEGER</code>	Input	Number of bytes of memory
Return value	<code>INTEGER</code> (Cray <code>POINTER</code>)	Output	<code>k>0</code> : <code>k</code> = address of the start of the block of memory allocated <code>k=0</code> : 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.			

Note – This function is intrinsic in Fortran 95 and external in Fortran 77. Fortran 77 programs compiled to run in 64-bit environments must declare the `malloc()` function and the variables receiving its output as `INTEGER*8`. Portability issues can be solved by using `malloc64()` instead of `malloc()` in Fortran 77 programs that must run in both 32-bit or 64-bit environments.

The function `malloc64(3F)` is provided to make programs portable between 32-bit and 64-bit environments:

<code>k = malloc64(n)</code>			
<code>n</code>	<code>INTEGER*8</code>	Input	Number of bytes of memory
Return value	<code>INTEGER*8</code> (<i>Cray</i> <code>POINTER</code>)	Output	<code>k>0</code> : <code>k</code> =address of <i>the</i> start of the block of memory allocated <code>k=0</code> : 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.

Reallocate Memory: `realloc` (*Fortran 95 Only*)

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

<code>k = realloc(ptr, n)</code>			
<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	<i>k</i> >0: <i>k</i> =address of <i>the</i> start of the new block of memory allocated <i>k</i> =0: Error
	An INTEGER*8 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 f77 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.

free: Deallocate Memory Allocated by Malloc

The subroutine is called by:

call free (<i>ptr</i>)		
<i>ptr</i>	Cray POINTER	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
```

mvbits: Move a Bit Field

The subroutine is called by:

call mvbits(<i>src</i> , <i>ini1</i> , <i>nbits</i> , <i>des</i> , <i>ini2</i>)			
<i>src</i>	INTEGER*4	Input	Source
<i>ini1</i>	INTEGER*4	Input	Initial bit position in the source
<i>nbits</i>	INTEGER*4	Input	Number of bits to move
<i>des</i>	INTEGER*4	Output	Destination
<i>ini2</i>	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
&      / 7, 0, 3, 1, 3 /
      call mvbits ( src, ini1, nbits, des, ini2 )
      write (*,"(5o3)") src, ini1, nbits, des, ini2
      end
demo% f77 -silent 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:
 - $ini1 + nbits \geq 32$
 - $ini2 + nbits \leq 32$

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.

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" )
```

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
```

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 `perror` is determined by the calling program.
- Runtime I/O error codes for `f77` and `f95` are listed in the *Fortran User's Guide*.

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.

putc: Write to Logical Unit 6

The function is called by:

INTEGER*4 putc <i>status</i> = putc(<i>char</i>)			
<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% f77 -silent tputc.f
demo% a.out
OK by putc
demo%
```

fputc: Write to Specified Logical Unit

The function is called by:

INTEGER*4 fputc <i>status</i> = fputc(<i>lunit</i> , <i>char</i>)			
<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% f77 -silent tfputc.f
demo% a.out
demo% cat tfputc.data
OK by fputc
demo%
```

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

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

The subroutine is called by:

call qsort(<i>array</i> , <i>len</i> , <i>isize</i> , <i>compar</i>) call qsort64(<i>array</i> , <i>len8</i> , <i>isize8</i> , <i>compar</i>)			
<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: <i>compar</i> (<i>arg1</i> , <i>arg2</i>)

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/, 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% f77 -silent tqsort.f
demo% a.out
 0 1 2 3 4 5 6 7 8 9
```

ran: Generate a Random Number Between 0 and 1

Repeated calls to `ran` generate a sequence of random numbers with a uniform distribution.

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

See `lcrans(3m)`.

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% f77 -silent 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:

```
SEED = 6909 * SEED + 1 (MOD 2**32)
```

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

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>			
<i>k</i>	INTEGER*4	Input	<i>k</i> =0: Get next random number in the sequence <i>k</i> =1: Restart sequence, return first number <i>k</i> >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% f77 -silent trand.f
demo% a.out
      2078917053 143302914 1027100827 1953210302 755253631
demo%
```

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% f77 -silent 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>`.

secnds: Get System Time in Seconds, Minus Argument

<code>t = secnds(t0)</code>			
<code>t0</code>	REAL	Input	Constant, variable, or array element
Return Value	REAL	Output	Number of seconds since midnight, minus <code>t0</code>

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% f77 -silent 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.

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.

signal: Change the Action for a Signal

The function is called by:

INTEGER*4 signal or INTEGER*8 signal <i>n</i> = signal(<i>signum</i> , <i>proc</i> , <i>flag</i>)			
<i>signum</i>	INTEGER*4	Input	Signal number; see <i>signal</i> (3)
<i>proc</i>	Routine name	Input	Name of user signal handling routine; must be in an external statement
<i>flag</i>	INTEGER*4	Input	<i>flag</i> < 0: Use <i>proc</i> as the signal handler <i>flag</i> ≥ 0: Ignore <i>proc</i> ; pass <i>flag</i> as the action: <i>flag</i> = 0: Use the default action <i>flag</i> = 1: Ignore this signal
Return value	INTEGER*4 INTEGER*8	Output	<i>n</i> =-1: System error <i>n</i> >0: Definition of previous action <i>n</i> >1: <i>n</i> =Address of routine that would have been called <i>n</i> <-1: If <i>signum</i> is a valid signal number, then: <i>n</i> =address of routine that would have been called. If <i>signum</i> is a <i>not</i> a valid signal number, then: <i>n</i> is an error number. On 64-bit environments, <i>signal</i> and the variables receiving its output must be declared INTEGER*8

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.

£77 arranges to trap certain signals when a process is started. The only way to restore the default £77 action is to save the returned value from the first call to *signal*.

`f77_floatingpoint.h` defines *proc* values `SIGFPE_DEFAULT`, `SIGFPE_IGNORE`, and `SIGFPE_ABORT`. See page 59. (Use `floatingpoint.h` with `f95`).

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*.

sleep: Suspend Execution for an Interval

The subroutine is called by:

call sleep(<i>itime</i>)			
<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)`.

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.

stat: Get Status for File, by File Name

The function is called by:

INTEGER*4 stat <i>ierr</i> = stat (<i>name</i> , <i>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
```

`fstat`: Get Status for File, by Logical Unit

The function

<code>INTEGER*4 fstat</code> <code>ierr = fstat (lunit, statb)</code>			
<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
```

lstat: Get Status for File, by File Name

The function is called by:

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

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>`.

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`.

system: Execute a System Command

The function is called by:

<code>INTEGER*4 system</code> <code>status = system(string)</code>			
<code>string</code>	<code>character*n</code>	Input	String containing command to do
Return value	<code>INTEGER*4</code>	Output	Exit status of the shell executed. See <code>wait(2)</code> 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` and `f77` developed with different assumptions:

- If `cc` calls `system`, the shell is always the Bourne shell.
- If `f77` calls `system`, then which shell is called depends on the environment variable `SHELL`.

The `system` function 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.

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

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

time, ctime, ltime, gmtime: Get System Time

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.

time: Get System Time

For `time()`, there are two versions, a standard version and a VMS version. If you use the `f77` command-line option `-lV77`, then you get the VMS version for `time()` and for `idate()`; otherwise, you get the standard versions. (The VMS versions of certain library routines is only available with `f77` through the `-lV77` library option, and not with `f95`.)

The standard function is called by:

INTEGER*4 <i>time</i> or INTEGER*8 <i>n</i> = <i>time</i> ()		<i>Standard Version</i>	
Return value	INTEGER*4	Output	Time, in seconds, since 0:0:0, GMT, 1/1/70
	INTEGER*8	Output	In 64-bit environments, <i>time</i> 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% f77 -silent ttime.f
demo% a.out
      Seconds since 0 1/1/70 GMT =    913240205
demo%
```

The VMS version of *time* is a subroutine that gets the current system time as a character string.

The VMS subroutine is called by:

call <i>time</i> (<i>t</i>)		<i>VMS Version</i>	
<i>t</i>	character*8	Output	Time, in the form <i>hh:mm:ss</i> <i>hh</i> , <i>mm</i> , and <i>ss</i> are each two digits: <i>hh</i> is the hour; <i>mm</i> is the minute; <i>ss</i> is the second

Example: `time(t)`, VMS version, `ctime`—convert the system time to ASCII:

```
character t*8
call time( t )
write(*, "(' The current time is ', A8)") t
end
demo% f77 -silent ttimeV.f -lv77
demo% a.out
The current time is 08:14:13
demo%
```

`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 <code>string</code> 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% f77 -silent tctime.f
demo% a.out
ctime: Wed Dec 9 13:50:05 1998
demo%
```

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 ltime(<i>stime</i> , <i>tarray</i>)			
<i>stime</i>	INTEGER*4	Input	System time from time() (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% f77 -silent tltime.f
demo% a.out
ltime: 25 49 10 12 7 91 1 223 1
demo%
```

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 gmtime(<i>stime</i> , <i>tarray</i>)			
<i>stime</i>	INTEGER*4	Input	System time from time() (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% f77 -silent 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)`.

`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.

topen, tclose, tread,..., tstate: Tape I/O

(FORTRAN 77 Only) These routines provide an alternative way to manipulate magnetic tape:

topen	Associate a device name with a tape logical unit.
tclose	Write EOF, close tape device channel, and remove association with <i>tlu</i> .
tread	Read next physical record from tape into buffer.
twrite	Write the next physical record from buffer to tape.
trewin	Rewind the tape to the beginning of the first data file.
tskipf	Skip forward over files and/or records, and reset EOF status.
tstate	Determine the logical state of the tape I/O channel.

On any one unit, do not mix these functions with standard Fortran I/O.

You must first use `topen()` to open a tape logical unit, *tlu*, for the specified device. Then you do all other operations on the specified *tlu*. *tlu* has no relationship at all to any normal Fortran logical unit.

Before you use one of these functions, its name must be in an `INTEGER*4` type statement.

topen: Associate a Device with a Tape Logical Unit

The function is called by:

<code>INTEGER*4 topen</code> <code>n = topen(tlu, devnam, islabeled)</code>			
<i>tlu</i>	<code>INTEGER*4</code>	Input	Tape logical unit, in the range 0 to 7.

INTEGER*4 <code>topen</code> <code>n = topen(tlu, devnam, islabeled)</code>			
<code>devnam</code>	CHARACTER	Input	Device name; for example: <code>'/dev/rst0'</code>
<code>islabeled</code>	LOGICAL	Input	True=the tape is labeled A label is the first file on the tape.
Return value	INTEGER*4	Output	<code>n=0</code> : OK <code>n<0</code> : Error

This function does *not* move the tape. See `perror(3F)` for details.

Example: `topen()`—open a 1/4-inch tape file:

```

CHARACTER devnam*9 / '/dev/rst0' /
INTEGER*4 n / 0 /, tlu / 1 /, topen
LOGICAL islabeled / .false. /
n = topen( tlu, devnam, islabeled )
IF ( n .LT. 0 ) STOP "topen: cannot open"
WRITE(*,('"topen ok:", 2I3, 1X, A10')) n, tlu, devnam
END

```

The output is:

```
topen ok: 0 1 /dev/rst0
```

`tclose`: Write EOF, Close Tape Channel, Disconnect `tlu`

The function is called by:

INTEGER*4 <code>tclose</code> <code>n = tclose (tlu)</code>			
<code>tlu</code>	INTEGER*4	Input	Tape logical unit, in range 0 to 7
<code>n</code>	INTEGER*4	Return value	<code>n=0</code> : OK <code>n<0</code> : Error

Caution – `tclose()` places an EOF marker immediately after the current location of the unit pointer, and then closes the unit. So if you `trewin()` a unit before you `tclose()` it, its contents are discarded.

Example: `tclose()`—close an opened 1/4-inch tape file:

```
CHARACTER devnam*9 / '/dev/rst0' /
INTEGER*4 n / 0 /, tlu / 1 /, tclose, topen
LOGICAL islabeled / .false. /
n = topen( tlu, devnam, islabeled )
n = tclose( tlu )
IF ( n .LT. 0 ) STOP "tclose: cannot close"
WRITE(*, '( "tclose ok:", 2I3, 1X, A10)') n, tlu, devnam
END
```

The output is:

```
tclose ok: 0 1 /dev/rst0
```

`twrite`: Write Next Physical Record to Tape

The function is called by:

<pre>INTEGER*4 twrite n = twrite(tlu, buffer)</pre>			
<i>tlu</i>	INTEGER*4	Input	Tape logical unit, in range 0 to 7
<i>buffer</i>	character	Input	Must be sized at a multiple of 512
<i>n</i>	INTEGER*4	Return value	<i>n</i> >0: OK, and <i>n</i> = the number of bytes written <i>n</i> =0: End of Tape <i>n</i> <0: Error

The physical record length is the size of `buffer`.

Example: `twrite()`—write a 2-record file:

```
CHARACTER devnam*9 / '/dev/rst0' /, recl*512 / "abcd" /,  
&      rec2*512 / "wxyz" /  
INTEGER*4 n / 0 /, tlu / 1 /, tclose, topen, twrite  
LOGICAL islabeled / .false. /  
n = topen( tlu, devnam, islabeled )  
IF ( n .LT. 0 ) STOP "topen: cannot open"  
n = twrite( tlu, recl )  
IF ( n .LT. 0 ) STOP "twrite: cannot write 1"  
n = twrite( tlu, rec2 )  
IF ( n .LT. 0 ) STOP "twrite: cannot write 2"  
WRITE(*, '( "twrite ok:", 2I4, 1X, A10)') n, tlu, devnam  
END
```

The output is:

```
twrite ok: 512 1 /dev/rst0
```

tread: Read Next Physical Record from Tape

The function is called by:

INTEGER*4 tread <i>n</i> = tread(<i>tlu</i> , <i>buffer</i>)			
<i>tlu</i>	INTEGER*4	Input	Tape logical unit, in range 0 to 7.
<i>buffer</i>	character	Input	Must be sized at a multiple of 512, and must be large enough to hold the largest physical record to be read.
<i>n</i>	INTEGER*4	Return value	<i>n</i> >0: OK, and <i>n</i> is the number of bytes read. <i>n</i> <0: Error <i>n</i> =0: EOF

If the tape is at EOF or EOT, then `tread` does a return; it does not read tapes.

Example: `tread()`—read the first record of the file written above:

```
CHARACTER devnam*9 / '/dev/rst0' /, onerec*512 / " " /
INTEGER*4 n / 0 /, tlu / 1 /, topen, tread
LOGICAL islabeled / .false. /
n = topen( tlu, devnam, islabeled )
IF ( n .LT. 0 ) STOP "topen: cannot open"
n = tread( tlu, onerec )
IF ( n .LT. 0 ) STOP "tread: cannot read"
WRITE(*, '( "tread ok:", 2I4, 1X, A10)') n, tlu, devnam
WRITE(*, '( A4)') onerec
END
```

The output is:

```
tread ok: 512 1 /dev/rst0
abcd
```

trewin: Rewind Tape to Beginning of First Data File

The function is called by:

INTEGER*4 trewin <i>n</i> = trewin (<i>tlu</i>)			
<i>tlu</i>	INTEGER*4	Input	Tape logical unit, in range 0 to 7
<i>n</i>	INTEGER*4	Return value	<i>n</i> =0: OK <i>n</i> <0: Error

If the tape is labeled, then the label is skipped over after rewinding.

Example 1: `trewin()`—typical fragment:

```
CHARACTER devnam*9 / '/dev/rst0' /
INTEGER*4 n / 0 /, tlu / 1 /, tclose, topen, tread, trewin
...
n = trewin( tlu )
IF ( n .LT. 0 ) STOP "trewin: cannot rewind"
WRITE(*, '( "trewin ok:", 2I4, 1X, A10)') n, tlu, devnam
...
END
```

Example 2: `trewin()`—in a two-record file, try to read three records, rewind, read one record:

```
CHARACTER devnam*9 / '/dev/rst0' /, onerec*512 / " " /
INTEGER*4 n / 0 /, r, tlu / 1 /, topen, tread, trewin
LOGICAL islabeled / .false. /
n = topen( tlu, devnam, islabeled )
IF ( n .LT. 0 ) STOP "topen: cannot open"
DO r = 1, 3
    n = tread( tlu, onerec )
    WRITE(*, '(1X, I2, 1X, A4)') r, onerec
END DO
n = trewin( tlu )
IF ( n .LT. 0 ) STOP "trewin: cannot rewind"
WRITE(*, '( "trewin ok:" 2I4, 1X, A10)') n, tlu, devnam
n = tread( tlu, onerec )
IF ( n .LT. 0 ) STOP "tread: cannot read after rewind"
WRITE(*, '(A4)') onerec
END
```

The output is:

```
1 abcd
2 wxyz
3 wxyz
trewin ok: 0 1 /dev/rst0
abcd
```

tskipf: Skip Files and Records; Reset EOF Status

The function is called by:

<code>INTEGER*4 tskipf</code> <code>n = tskipf(tlu, nf, nr)</code>			
<code>tlu</code>	<code>INTEGER*4</code>	Input	Tape logical unit, in range 0 to 7
<code>nf</code>	<code>INTEGER*4</code>	Input	Number of end-of-file marks to skip over first
<code>nr</code>	<code>INTEGER*4</code>	Input	Number of physical records to skip over after skipping files
<code>n</code>	<code>INTEGER*4</code>	Return value	<code>n=0</code> : OK <code>n<0</code> : Error

This function does *not* skip backward.

First, the function skips forward over `nf` end-of-file marks. Then, it skips forward over `nr` physical records. If the current file is at EOF, this counts as one file to skip. This function also resets the EOF status.

Example: `tskipf()`—typical fragment: skip four files and then skip one record:

```
INTEGER*4 nfiles / 4 /, nrecords / 1 /, tskipf, tlu / 1 /
...
n = tskipf( tlu, nfiles, nrecords )
IF ( n .LT. 0 ) STOP "tskipf: cannot skip"
...
```

Compare with `tstate()`.

tstate: Get Logical State of Tape I/O Channel

The function is called by:

<code>INTEGER*4 tstate</code> <code>n = tstate(tlu, fileno, recno, errf, eoff, eotf, tcsr)</code>			
<i>tlu</i>	INTEGER*4	Input	Tape logical unit, in range 0 to 7
<i>fileno</i>	INTEGER*4	Output	Current file number
<i>recno</i>	INTEGER*4	Output	Current record number
<i>errf</i>	LOGICAL	Output	True=an error occurred
<i>eoff</i>	LOGICAL	Output	True=the current file is at EOF
<i>eotf</i>	LOGICAL	Output	True=tape has reached logical end-of-tape
<i>tcsr</i>	INTEGER*4	Output	True=hardware errors on the device. It contains the tape drive control status register. If the error is software, then <i>tcsr</i> is returned as zero. The values returned in this status register vary grossly with the brand and size of tape drive.

For details, see `st(4s)`.

While *eoff* is true, you cannot read from that *tlu*. You can set this EOF status flag to false by using `tskipf()` to skip one file and zero records:

```
n = tskipf( tlu, 1, 0).
```

Then you can read any valid record that follows.

End-of-tape (EOT) is indicated by an empty file, often referred to as a double EOF mark. You cannot read past EOT, but you can write past it.

Example: Write three files of two records each:

```
CHARACTER devnam*10 / '/dev/nrst0' /,
&          f0rec1*512 / "eins" /, f0rec2*512 / "zwei" /,
&          flrec1*512 / "ichi" /, flrec2*512 / "ni__" /,
&          f2rec1*512 / "un__" /, f2rec2*512 / "deux" /
INTEGER*4 n / 0 /, tlu / 1 /, tclose, topen, trewin, twrite
LOGICAL islabeled / .false. /
n = topen( tlu, devnam, islabeled )
n = trewin( tlu )
n = twrite( tlu, f0rec1 )
n = twrite( tlu, f0rec2 )
n = tclose( tlu )
n = topen( tlu, devnam, islabeled )
n = twrite( tlu, flrec1 )
n = twrite( tlu, flrec2 )
n = tclose( tlu )
n = topen( tlu, devnam, islabeled )
n = twrite( tlu, f2rec1 )
n = twrite( tlu, f2rec2 )
n = tclose( tlu )
END
```

The next example uses `tstate()` to trap EOF and get at all files.

Example: Use `tstate()` in a loop that reads all records of the 3 files written in the previous example:

```
CHARACTER devnam*10 / '/dev/nrst0' /, onerec*512 / " " /
INTEGER*4 f, n / 0 /, tlu / 1 /, tcsr, topen, tread,
&   trewin, tskipf, tstate
LOGICAL errf, eoff, eotf, islabeled / .false. /
n = topen( tlu, devnam, islabeled )
n = tstate( tlu, fn, rn, errf, eoff, eotf, tcsr )
WRITE(*,1) 'open:', fn, rn, errf, eoff, eotf, tcsr
1  FORMAT(1X, A10, 2I2, 1X, 1L, 1X, 1L,1X, 1L, 1X, I2 )
2  FORMAT(1X, A10,1X,A4,1X,2I2,1X,1L,1X,1L,1X,1L,1X,I2)
n = trewin( tlu )
n = tstate( tlu, fn, rn, errf, eoff, eotf, tcsr )
WRITE(*,1) 'rewind:', fn, rn, errf, eoff, eotf, tcsr
DO f = 1, 3
  eoff = .false.
  DO WHILE ( .NOT. eoff )
    n = tread( tlu, onerec )
    n = tstate( tlu, fn, rn, errf, eoff, eotf, tcsr )
    IF (.NOT. eoff) WRITE(*,2) 'read:', onerec,
&   fn, rn, errf, eoff, eotf, tcsr
  END DO
  n = tskipf( tlu, 1, 0 )
  n = tstate( tlu, fn, rn, errf, eoff, eotf, tcsr )
  WRITE(*,1) 'tskip: ', fn, rn, errf, eoff, eotf, tcsr
END DO
END
```

The output is:

```
open: 0 0 F F F 0
rewind: 0 0 F F F 0
read: eins 0 1 F F F 0
read: zwei 0 2 F F F 0
tskip: 1 0 F F F 0
read: ichi 1 1 F F F 0
read: ni__ 1 2 F F F 0
tskip: 2 0 F F F 0
read: un__ 2 1 F F F 0
read: deux 2 2 F F F 0
tskip: 3 0 F F F 0
```

A summary of EOF and EOT follows:

- If you are at either EOF or EOT, then:

- Any `tread()` just returns; it does not read the tape.
- A successful `tskipf(tlu,1,0)` resets the EOF status to false, and returns; it does not advance the tape pointer.
- A successful `twrite()` resets the EOF and EOT status flags to false.
- A successful `tclose()` resets all those flags to false.
- `tclose()` truncates; it places an EOF marker immediately after the current location of the unit pointer, and then closes the unit. So, if you use `trewin()` to rewind a unit before you use `tclose()` to close it, its contents are discarded. This behavior of `tclose()` is inherited from the Berkeley code.

See also: `ioctl(2)`, `mtio(4s)`, `perror(3F)`, `read(2)`, `st(4s)`, and `write(2)`.

ttynam, isatty: Get Name of a Terminal Port

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

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:

<pre>CHARACTER ttynam*24 name = ttynam(lunit)</pre>			
<i>lunit</i>	INTEGER*4	Input	Logical unit
Return value	character*n	Output	<p>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.</p> <p>If empty string (all blanks) returned: <i>lunit</i> is not associated with a terminal device in the directory, /dev</p>

isatty: Is this Unit a Terminal?

The function

<code>terminal = isatty(lunit)</code>			
<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

is called by:

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/ttypl "
```

unlink: Remove a File

The function is called by:

INTEGER*4 unlink <code>n = unlink (patnam)</code>			
<i>patnam</i>	character*n	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% f77 -silent 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>`.

wait: Wait for a Process to Terminate

The function is:

INTEGER*4 wait n = wait(status)			
<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).

Index

SYMBOLS

(e**x)-1, 13, 16

Δ, blank character, 2

A

abort, 20

access

time, 94

access, 20

action for signal, change, signal, 92

address

loc, 71

alarm, 21

and, 22

append on open

ioinit, 63

arc

cosh, 13, 16

cosine, 13

sine, 13

sinh, 13

tangent, 13

tanh, 16

arc tangent, 13

arguments

command line, getarg, 41

B

bessel, 15, 16

bic, 23

bis, 23

bit

functions, 23

move bits, mvbits, 79

bit, 23

bitwise

and, 22

complement, 22

exclusive or, 22

inclusive or, 22

blocks allocated, 94

blocksize, 94

C

carriage control

initialize, ioinit, 63

ceiling, 13

change

action for signal, signal, 92

default directory, chdir, 26

mode of a file, chmod, 27

character

get a character `getc`, `fgetc`, 42

put a character, `putc`, `fputc`, 82

chdir, 26

chmod, 27

- clear
 - bit, 23
- command-line argument, `getarg`, 41
- compilers, accessing, 4
- complement, 22
- conversion by long, short, 72
- copy
 - process via `fork`, 36
- core file, 20
- `ctime`, convert system time to character, 98, 100
- `ctime64`, 102
- cube root, 13
- current working directory, `getcwd`, 45

D

- `d_acos(x)`, 16
- `d_acosd(x)`, 16
- `d_acosh(x)`, 16
- `d_acosp(x)`, 16
- `d_acospi(x)`, 16
- `d_addran()`, 17
- `d_addrans()`, 17
- `d_asin(x)`, 16
- `d_asind(x)`, 16
- `d_asinh(x)`, 16
- `d_asinp(x)`, 16
- `d_asinpi(x)`, 16
- `d_atan(x)`, 16
- `d_atan2(x)`, 16
- `d_atan2d(x)`, 16
- `d_atan2pi(x)`, 16
- `d_atand(x)`, 16
- `d_atanh(x)`, 16
- `d_atanp(x)`, 16
- `d_atanpi(x)`, 16
- `d_cbrt(x)`, 16
- `d_ceil(x)`, 16
- `d_erf(x)`, 16
- `d_erfc(x)`, 16
- `d_expml(x)`, 16
- `d_floor(x)`, 16
- `d_hypot(x)`, 16

- `d_infinity()`, 16
- `d_j0(x)`, 16
- `d_j1(x)`, 16
- `d_jn(n,x)`, 16
- `d_lcran()`, 17
- `d_lcrans()`, 17
- `d_lgamma(x)`, 17
- `d_log1p(x)`, 17
- `d_log2(x)`, 17
- `d_logb(x)`, 17
- `d_max_normal()`, 17
- `d_max_subnormal()`, 17
- `d_min_normal()`, 17
- `d_min_subnormal()`, 17
- `d_nextafter(x,y)`, 17
- `d_quiet_nan(n)`, 17
- `d_remainder(x,y)`, 17
- `d_rint(x)`, 17
- `d_scalbn(x,n)`, 17
- `d_shufrens()`, 17
- `d_signaling_nan(n)`, 17
- `d_significand(x)`, 17
- `d_sin(x)`, 17
- `d_sincos(x,s,c)`, 18
- `d_sincosd(x,s,c)`, 18
- `d_sincosp(x,s,c)`, 18
- `d_sincospi(x,s,c)`, 18
- `d_sind(x)`, 17
- `d_sinh(x)`, 17
- `d_sinp(x)`, 17
- `d_sinpi(x)`, 17
- `d_tan(x)`, 18
- `d_tand(x)`, 18
- `d_tanh(x)`, 18
- `d_tanp(x)`, 18
- `d_tanpi(x)`, 18
- `d_y0(x)`, `bessel`, 18
- `d_y1(x)`, `bessel`, 18
- `d_yn(n,x)`, 18
- data types, 9
- date
 - and time, as characters, `fdate`, 34
 - as integer, `idate`, 53

- current date, `date`, 27
- `date_and_time`, 28
- `date_and_time`, 28
- deallocate memory by `free`, 79
- default
 - directory change, `chdir`, 26
- delay execution, `alarm`, 21
- descriptor, `get file`, `getfd`, 47
- device name, type, size, 94
- directory
 - default change, `chdir`, 26
 - get current working directory, `getcwd`, 45
- documentation index, 5
- documentation, accessing, 5
- double-precision
 - functions, 15
- `drand`, 88
- `dtime`, 31

E

- elapsed time, 31
- embedded
 - blanks, `initialize`, `ioinit`, 63
- environment variables, `getenv`, 46
- EOF reset status for `tapeio`, 109
- error
 - function, 13
 - messages, `perror`, `gerror`, `ierrno`, 80
- errors and interrupts, `longjmp`, 74
- `etime`, 31
- exclusive or, 22
- execute an OS command, `system`, 91, 97
- execution time, 31
- existence of file, `access`, 20
- `exit`, 33

F

- `f77_floatingpoint` IEEE definitions, 59
- `f77_ieee_environment`, 55
- `fdate`, 34
- `fgetc`, 43

file

- connection, `automatic`, `ioinit`, 64
- descriptor, `get`, `getfd`, 47
- get file pointer, `getfilep`, 48
- mode, `access`, 20
- permissions, `access`, 20
- remove, `unlink`, 114
- rename, 89
- status, `stat`, 94
- status, `stat64`, 97
- find substring, `index`, 61
- floating-point
 - IEEE definitions, 59
- floor, 13
- flush, 35
- fork, 36
- `fputc`, 82
- `free`, 79
- free format, 2
- `fseek`, 37
- `fseeko64`, 39
- `fstat`, 94
- `fstat64`, 97
- `ftell`, 37
- `ftello64`, 39
- functions
 - quadruple-precision, `libm_quadruple`, 18
 - single-precision, `libm_single`, 13

G

- `gerror`, 80
- `get`
 - character `getc`, `fgetc`, 42
 - current working directory, `getcwd`, 45
 - environment variables, `getenv`, 46
 - file descriptor, `getfd`, 47
 - file pointer, `getfilep`, 48
 - group id, `getgid`, 51
 - login name, `getlog`, 50
 - process id, `getpid`, 50
 - user id, `getuid`, 51
- `getarg`, 41
- `getc`, 42
- `getcwd`, 45

getenv, 46
getfd, 47
getfilep, 48
getgid, 51
getlog, 50
getpid, 50
getuid, 51
gmtime, 98
gmtime(), GMT, 102
gmtime64, 102
Greenwich Mean Time, gmtime, 98
group, 94
group ID, get, getgid, 51

H

hard links, 94
host name, get, hostnm, 52
hostnm, 52
hyperbolic cos, 13
hyperbolic tan, 15, 18
hypotenuse, 13

I

iargc, 41
id, process, get, getpid, 50
id_finite(x), 17
id_fp_class(x), 17
id_rint(x), 17
id_isinf(x), 17
id_isnan(x), 17
id_isnormal(x), 17
id_issubnormal(x), 17
id_iszero(x), 17
id_logb(x), 17
id_signbit(x), 17
IEEE, 59
 environment, 55
ieee_flags, 55
ieee_handler>, 55
ierrno, 80

IMPLICIT, 9
inclusive or, 22
index, 60
initialize
 I/O, ioinit, 63
inmax, 63
inode, 94
integer
 conversion by long, short, 72
interrupts and errors, longjmp, 74
intrinsic math functions, 12
ioinit, 63
iq_finite(x), 19
iq_fp_class(x), 19
iq_isinf(x), 19
iq_isnan(x), 19
iq_isnormal(x), 19
iq_issubnormal(x), 19
iq_iszero(x), 19
iq_logb(x), 19
iq_signbit(x), 19
ir_finite(x), 14
ir_fp_class(x), 14
ir_rint(x), 14
ir_isinf(x), 14
ir_isnan(x), 14
ir_isnormal(x), 14
ir_issubnormal(x), 14
ir_iszero(x), 14
ir_logb(x), 14
ir_signbit(x), 14
irand, 88
isatty, 113
isetjmp, 73

J

jump, longjmp, setjmp, 74

K

kill, send signal, 69

L

- left shift, `lshift`, 22
- `libm_double`, 15
- `libm_quadruple`, 18
- `libm_single`, 12
- `link`, 69
- link to an existing file, `link`, 69
- `lnblnk`, 62
- local time zone, `lmtime()`, 101
- location of
 - a variable `loc`, 71
- log gamma, 14
- login name, `get` `getlog`, 49
- `long`, 72
- `longjmp`, 73
- `lshift`, 22
- `lstat`, 94
- `lstat64`, 97
- `ltime`, 98
- `ltime()`, local time zone, 101
- `ltime64`, 102

M

- `malloc`, 76
- man pages, accessing, 3
- `MANPATH` environment variable, setting, 5
- math functions, intrinsics, 12
- maximum
 - positive integer, `inmax`, 63
- memory
 - deallocate by `free`, 79
- mode
 - IEEE, 55
 - of file, access, 20
- modifying
 - time, 94
- `mvbits`, move bits, 80

N

- name
 - login, `get`, `getlog`, 49

- terminal port, `ttynam`, 113
- `not`, 22

O

- `or`, 22
- OS command, `execute`, `system`, 91, 97

P

- `PATH` environment variable, setting, 3
- permissions
 - access function, 20
- `perror`, 80
- `pid`, process id, `getpid`, 50
- pointer
 - get file pointer, `getfilep`, 48
- position file by `fseek`, `ftell`, 37
- position file by `fseeko64`, `ftello64`, 39
- process
 - create copy with `fork` function, 36
 - id, `get`, `getpid`, 50
 - send signal to, `kill`, 69
 - wait for termination, `wait`, 115
- protection, 94
- put a character, `putc`, `fputc`, 82
- `putc`, 82

Q

- `q_copysign(x)`, 19
- `q_fabs(x)`, 19
- `q_fmod(x)`, 19
- `q_infinity()`, 19
- `q_max_normal()`, 19
- `q_max_subnormal()`, 19
- `q_min_normal()`, 19
- `q_min_subnormal()`, 19
- `q_nextafter(x,y)`, 19
- `q_quiet_nan(n)`, 19
- `q_remainder(x,y)`, 19
- `q_scalbn(x,n)`, 19
- `q_signaling_nan(n)`, 19

qsort, qsort64, 85
quadruple-precision functions,
 libm_quadruple, 18
quick sort, qsort, 85

R

r_acos(x), 13
r_acosd(x), 13
r_acosh(x), 13
r_acosp(x), 13
r_acospi(x), 13
r_addran(), 14
r_addrans(), 14
r_asin(x), 13
r_asind(x), 13
r_asinh(x), 13
r_asinp(x), 13
r_asinpi(x), 13
r_atan(x), 13
r_atan2(x), 13
r_atan2d(x), 13
r_atan2pi(x), 13
r_atand(x), 13
r_atanh(x), 13
r_atanp(x), 13
r_atanpi(x), 13
r_cbrt(x), 13
r_ceil(x), 13
r_erf(x), 13
r_erfc(x), 13
r_expml(x), 13
r_floor(x), 13
r_hypot(x), 13
r_infinity(), 13
r_j0(x), 13
r_j1(x), 13
r_jn(n,x), 13
r_lcran(), 14
r_lcrans(), 14
r_lgamma(x), 14
r_loglp(x), 14
r_log2(x), 14

r_logb(x), 14
r_max_normal(), 14
r_max_subnormal(), 14
r_min_normal(), 14
r_min_subnormal(), 14
r_nextafter(x,y), 14
r_quiet_nan(n), 14
r_remainder(x,y), 14
r_rint(x), 14
r_scalbn(x,n), 14
r_shuftrans(), 14
r_signaling_nan(n), 14
r_significand(x), 14
r_sin(x), 14
r_sincos(x,s,c), 15
r_sincosd(x,s,c), 15
r_sincosp(x,s,c), 15
r_sincospi(x,s,c), 15
r_sind(x), 14
r_sinh(x), 14
r_sinp(x), 14
r_sinpi(x), 14
r_tan(x), 15
r_tand(x), 15
r_tanh(x), 15
r_tanp(x), 15
r_tanpi(x), 15
r_y0(x), *bessel*, 15
r_y1(x), *bessel*, 15
r_yn(n,x), *bessel*, 15
rand, 88
random
 number, 14
 values, rand, 88
read
 character *getc*, *fgetc*, 42
 remove a file, *unlink*, 114
 reposition file by *fseek*, *ftell*, 37
 reposition file by *fseeko64*, *ftello64*, 39
 reset EOF status for *tapeio*, 109
 right shift, *rshift*, 22
rindex, 61
rshift, 22

S

`secnds`, system time, 90
`send signal to process`, `kill`, 69
`setbit`, 23
`setjmp`, *See* `isetjmp`
shell prompts, 3
`short`, 72
`signal`, 92
`signal a process`, `kill`, 69
signals, IEEE, 55
`sine`, 14
single-precision functions, `libm_single`, 13
64-bit environments, 10
`skip`
 tape I/O files and records, 109
`sleep`, 93
Solaris operating environment versions
 supported, 3
`sort quick`, `qsort`, 85
`stat`, 94
`stat64`, 97
`status`
 file, `stat`, 94
 file, `stat64`, 97
 IEEE, 55
 termination, `exit`, 33
`substring`
 find, `index`, 61
`suspend execution for an interval`, `sleep`, 93
symbolic
 link to an existing file, `symlink`, 69
`symlink`, 69
`system`, 91, 97
system time
 `secnds`, 90
 `time`, 98
`system.inc` include file, 10

T

tab format, 2
`tangent`, 15
tape I/O, 103
 close files, 104

 open files, 103
 read from files, 106
 reset EOF status, 109
 rewind files, 107
 skip files and records, 109
 write to files, 105
`tarray()` values for various time routines, 102
`tclose`, 103
terminal
 port name, `ttynam`, 113
terminate
 wait for process to terminate, `wait`, 115
 with status, `exit`, 33
 write memory to core file, 20
`time`, 31
 in numerical form, 53
 `secnds`, 90
`time(t)`
 standard version, 99
 VMS version, 100
`time`, get system time, 98
`topen`, 103
trailing blanks, `initialize`, `ioinit`, 63
`tread`, 103
`trewin`, 103
`tskipf`, 103
`tstate`, 103
`ttynam`, 113
`twrite`, 103
typographic conventions, 2

U

`unlink`, 114
`user`, 94
user ID, `get`, `getuid`, 51

W

`wait`, 115
write a character `putc`, `fputc`, 82

X

xor, 22

Y

y0(x), y1(x), yn),bessel, 15

y0(x), y1(x), yn(x),bessel, 18