

NAG Fortran Compiler, Release 5.3

May 16, 2012

NAG[®] Fortran Compiler

© 2011 The Numerical Algorithms Group Limited

All rights reserved. No part of this Manual may be reproduced, transcribed, stored in a retrieval system, translated into any language or computer language or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except for the purpose of using the NAG Fortran Compiler.

The copyright owner gives no warranties and makes no representations about the contents of this Manual and specifically disclaims any implied warranties of merchantability or fitness for any purpose.

The copyright owner reserves the right to revise this Manual and to make changes from time to time in its contents without notifying any person of such revisions or changes.

5th Edition – October 2011

NAG is a registered trademark of:

The Numerical Algorithms Group Limited
The Numerical Algorithms Group Inc
Nihon Numerical Algorithms Group KK

All other trademarks are acknowledged.

NAG Ltd

Wilkinson House
Jordan Hill Road
OXFORD
OX2 8DR
United Kingdom
Tel: +44 (0)1865 511245
Fax: +44 (0)1865 310139
Web: www.nag.co.uk

Nihon NAG KK

Hatchobori Frontier Building 2F
4-9-9
Hatchobori
Chuo-ku
Tokyo
104-0032
Japan
Tel: +81 (0)3 5542 6311
Fax: +81 (0)3 5542 6312
Web: www.nag-j.co.jp

NAG Inc

801 Warrenville Road, Suite 185
Lisle, IL 60532-4332
USA
Tel: +1 630 971 2337
Fax: +1 630 971 2706
Web: www.nag.com

NAG also has a number of distributors throughout the world. Please contact NAG for further details.

Introduction to the Compiler	Page 1
 Using the Compiler	
Usage	2
Debugging with dbx90	22
Preprocessing with fpp	28
 Extensions	
Non-standard Extensions	36
Obsolete Extensions	37
 Intrinsic Modules	
NAG-specific Modules	40
Modern Fortran API to Posix	47
High Performance Fortran Library Module	67
 Standard Fortran 95	
Fortran 95 Program Structure	70
Fortran 95 Expressions	73
Fortran 95 Statements	77
Fortran 95 Intrinsic Procedures	93
 Fortran 2003 Extensions	
Fortran 2003 Overview	97
Object-oriented Programming	98
ALLOCATABLE extensions	103
Other data-oriented enhancements	109
C interoperability	115
IEEE arithmetic support	118
Input/output Features	125
Miscellaneous Fortran 2003 Features	132
 Fortran 2008 Extensions	
Fortran 2008 Overview	142
 Appendices	
Mixing Fortran and C	145
ASCII Collating Sequence	151

Introduction

1 Introduction to the Compiler

The NAG Fortran Compiler is based on the NAGWare f90 Compiler which was the world's first Fortran 90 compiler. The design goals of the development were to produce a compiler with the following characteristics:

- compiles Fortran 90/95 to host-compatible C;
- good speed of compilation, reasonable efficiency of execution;
- good error checking, comprehensible error messages;
- full standard implementation, standard-conforming compiler (i.e. all constraints identified);
- modular construction;
- compiler written in C;
- maintainability, portability and re-usability.

The compiler is multi-pass; the passes have been kept distinct to improve maintainability and to allow re-use of the components.

Pass 1: Lexical and syntactic analysis, build symbol table and abstract syntax tree.

Pass 2: Semantic analysis, annotate parse tree and fill in symbol table; all major error and constraint detection takes place in this pass.

Pass 3: Code generation by parse tree transformation.

Pass 4: Code output, generate declarations and flatten transformed parse tree to C source code.

Pass 5: Compilation using the host operating system's C compiler.

Pass 6: Linking to executable code using the host system's linker, including linking in the Fortran run-time libraries.

1.1 Other Fortran-related Activities at NAG

NAG has released several Fortran-based numerical procedure libraries: the Fortran Library, SMP Library, Parallel Library and the Fortran 90 Library. A number of implementations of these libraries, using the NAG compiler and other compilers, are available.

NAG has supported the development and standardisation of Modern Fortran, participating both in the ISO Working Group and the technical development committee INCITS/J3. The head of the NAG compiler team, Malcolm Cohen, is the current editor of the Fortran standard.

From this it can be seen that NAG is committed to Fortran.

1.2 This Manual

This is the documentation for the NAG Fortran Compiler. This is not intended to be a language description or tutorial, but rather a guide to the use of the software and a quick reference for some of the features of the language.

The compiler is a full implementation of the ISO Fortran 95 programming language [IS 1539-1:1997(E)], with most of the features from the previous ISO Fortran 2003 standard [IS 1539-1:2004(E)] and some from the current ISO Fortran 2008 standard [ISO/IEC 1539-1:2010(E)].

Using the Compiler

2 Usage

nagfor [mode] [option]... *file*...

3 Description

nagfor is the interface to the NAG Fortran Compiler system. The compiler translates programs written in Fortran into executable programs, relocatable binary modules, assembler source files or C source files.

The *mode* determines the action performed, and can be one of

=compiler

Compile (and/or link) the files; this is the default mode if none is specified.

=callgraph

Produce a callgraph of the Fortran routines in the files (see the Producing a Call Graph section).

=depend

Produce a dependency analysis of the Fortran files (see the Dependency Analysis section).

=polish Pretty-print (polish) the Fortran files (see the Source File Polishing section).

Options that do not apply to the current mode of operation (e.g. polish options when the mode is for compilation) are ignored.

4 File Types

A file ending in `‘.f90’` or `‘.f95’` is taken to be a Fortran free-form source file, a file ending in `‘.f’`, `‘.for’` or `‘.ftn’` is taken to be a Fortran fixed-form source file; these assumptions can be overridden with the `–fixed` or `–free` option. A file ending in `‘.ff90’` or `‘.ff95’` is taken to be a free-form file requiring preprocessing by fpp, and a file ending in `‘.ff’` is taken to be a fixed-form file requiring preprocessing by fpp. On Unix, a file ending in `‘.F90’` or `‘.F95’` is taken to be a free-form file requiring preprocessing by fpp, and a file ending in `‘.F’` is taken to be a fixed-form files requiring preprocessing by fpp. (Note that on Mac OS X and Windows, the file system is not case-sensitive so uppercase and lowercase letters are equivalent in filenames including in the suffixes.)

If a filename without a suffix is provided nagfor will look for a file with the suffix `‘.f95’`, and if that does not exist, the suffix `‘.f90’`.

Non-intrinsic modules and `INCLUDE` files are expected to exist in the current working directory or in a directory named by the `–I` option.

5 Compiler Options

–132 Increase the length of each fixed source form input line from 72 characters to 132 characters. This has no effect on free source form input.

–abi=abi

(Linux x86-64 only) Specify the ABI to compile for, either **32** (the 32-bit ABI), **64t** (the AMD 64-bit ABI) or **64** (the AMD 64-bit ABI with object sizes limited to 2 GB). Programs compiled with `–abi=64t` or `–abi=64` will only run on 64-bit kernels. Programs compiled with `–abi=32` will run on any x86 Linux system. The default is `–abi=64t`. The `–abi=64` option should only be used to generate code compatible with NAGWare Release 5.1(346) or earlier, it will be withdrawn at a future release.

-abi=abi

(64-bit Intel Mac only) Specify the ABI to compile for, either **32** (the 32-bit ABI) or **64** (the 64-bit ABI). The default is *-abi=64*.

-abi=abi

(Windows only) Specify the ABI to compile for, either **32** (the 32-bit ABI) or **64** (the 64-bit ABI). The default is *-abi=64* on Windows x64; on 32-bit Windows the default is *-abi=32* and the *-abi=64* option is not available.

-align=alignment

(Mac OS X only) Specify the alignment of variables and components, which must be one of:

natural (natural alignment for best performance; this can alter the interpretation of **COMMON** block or **SEQUENCE** type layout in a non-standard-conforming manner), or

standard (use standard-conforming alignment; this is the default).

The whole program should be compiled with the same alignment option.

-Bbinding

Specify **static** or **dynamic** binding. This only has effect if specified during the link phase. The default is dynamic binding. On SPARC and SG/Irix, these options are positional and can be used to selectively bind some libraries statically and some dynamically. This option is not available on IBM z9 Open Edition.

-c

Compile only (produce .o file for each source file), do not link the .o files to produce an executable file. This option is equivalent to *-otype=obj*.

-C

Compile with all but the most expensive runtime checks; this omits the *-C=dangling* and *-C=undefined* options.

-C=check

Compile checking code according to the value of *check*, which must be one of:

all	(perform all checks except for <i>-C=undefined</i>),
array	(check array bounds),
bits	(check bit intrinsic arguments),
calls	(check procedure references),
dangling	(check for dangling pointers),
do	(check DO loops for zero step values),
none	(do no checking: this is the default),
present	(check OPTIONAL references),
pointer	(check POINTER references),
recursion	(check for invalid recursion) or
undefined	(check for undefined variables).

The *-C=undefined* option is subject to a number of limitations and is not binary-compatible with code compiled without that option; see the Undefined Variable Detection section for further details.

-colour

Colour the message output from the compiler using ANSI escape sequences and the default foreground colouring scheme which is: red for error messages (including fatal errors), blue for warning messages and green for information messages.

-colour=scheme

Colour the message output from the compiler according to the specified *scheme*. This is a comma-separated list of colour specifications, each consisting of a message category name (“error”, “warn” or “info”) followed by a colon and the foreground colour name, optionally following by a plus sign and the background colour name. The colouring for unspecified categories will be the default.

Colours are: black, red, green, yellow, blue, magenta, cyan and white.

E.g. *—colour=error:red+blue,warn:cyan,info:magenta+yellow* would be a rather garish colour scheme.

-compatible

Make external linkages compatible with other compilers where possible; on Windows this is Microsoft Fortran, on Mac OSX and Linux this is g77, g95 and gfortran, and on other systems this is the operating system vendor’s compiler. This affects the naming convention and procedure calling convention (for example, on Windows it causes use of the “STDCALL” calling convention that is commonly used for most DLLs).

-convert=*format*

Set the default conversion mode for unformatted files to *format*. This format may be overridden by an explicit **CONVERT=** specifier in the **OPEN** statement, or by the environment variable **FORT_CONVERT_n** (where *n* is the unit number). The value of *format* must be one of the following (not case-sensitive):

Format	Description
BIG_ENDIAN	synonym for BIG_IEEE
BIG_IEEE_DD	big-endian with IEEE floating-point, quad precision is double-double
BIG_IEEE	big-endian with IEEE floating-point, including quad precision
BIG_NATIVE	big-endian with native floating-point format
LITTLE_ENDIAN	synonym for LITTLE_IEEE
LITTLE_IEEE_DD	little-endian with IEEE floating-point, quad precision is double-double
LITTLE_IEEE	little-endian with IEEE floating-point, including quad precision
LITTLE_NATIVE	little-endian with native floating-point format
NATIVE	no conversion (the default)

-D*name*

Defines *name* to fpp as a preprocessor variable. This only affects files that are being preprocessed by fpp.

-dcfuns Enable recognition of non-standard double precision complex intrinsic functions. These act as specific versions of the standard generic intrinsics as follows:

Non-standard	Equivalent Standard Fortran Generic Intrinsic Function
CDABS(A)	ABS(A)
DCMPLX(X,Y)	CMPLX(X,Y,KIND=KIND(0d0))
DCONJG(Z)	CONJG(Z)
DIMAG(Z)	AIMAG(Z)
DREAL(Z)	REAL(Z) or DBLE(Z)

-double Double the size of default **INTEGER**, **LOGICAL**, **REAL** and **COMPLEX**. Entities specified with explicit kind numbers or byte lengths are unaffected. If quadruple precision **REAL** is available, the size of **DOUBLE PRECISION** is also doubled.**-dryrun** Show but do not execute commands constructed by the compiler driver.**-dusty** Allows the compilation and execution of “legacy” software by downgrading the category of common errors found in such software from “Error” to “Warning” (which may then be suppressed entirely with the *-w* option). This option disables *-C=calls*, and also enables Hollerith i/o (see the *-hollerith_io* option).**-encoding=*charset***

Specifies that the encoding system of the Fortran source files is *charset*, which must be one of **ISO_Latin_1**, **Shift_JIS** or **UTF_8**. If this option is not specified, the default encoding is UTF-8 for Fortran source files that begin with a UTF-8 Byte Order Mark, and ISO Latin-1 (if the language setting is English) or Shift-JIS (if the language setting is Japanese) for other Fortran source files.

-english Produce compiler messages in English (default).**-F** Preprocess only, do not compile. Each file that is preprocessed will produce an output file of the same name with the suffix replaced by **.f**, **.f90** or **.f95** according to the suffix of the input file. This option is equivalent to *-otype=Fortran*.**-f90_sign**

Use the Fortran 77/90 version of the **SIGN** intrinsic instead of the Fortran 95 one (they differ in the treatment of negative zero).

-f95 Specify that the base language is Fortran 95. This only affects extension message generation (Fortran 2003 and 2008 features will be reported as extensions).**-f2003** Specify that the base language is Fortran 2003. This only affects extension message generation (Fortran 2008 features will be reported as extensions).**-f2008** Specify that the base language is Fortran 2008. This is the default.**-fixed** Interpret all Fortran source files according to fixed-form rules.

-float-store

(Gnu C based systems only) Do not store floating-point variables in registers on machines with floating-point registers wider than 64 bits. This can avoid problems with excess precision.

-fpp Preprocess the source files using fpp even if the suffix would normally indicate an ordinary Fortran file.

-framework *f*

(Mac OS X only) Use framework *f* during linking.

-free Interpret all Fortran source files according to free-form rules.

-g Produce information for interactive debugging by the host system debugger.

-g90 Produce debugging information for dbx90, a Fortran 90 aware front-end to the host system debugger. This produces a debug information (.g90) file for each Fortran source file. This option must be specified for both compilation and linking.

-gc Enables automatic garbage collection of the executable program. This option must be specified for both compilation and linking, and is unavailable on IBM z9 OpenEdition and Windows x64. It is incompatible with the *-thread_safe* and *-mtrace* options. For more details see the Automatic Garbage Collection section.

-gline Compile code to produce a traceback when a runtime error message is generated. Only routines compiled with this option will appear in such a traceback. This option increases both executable file size and execution time. For example:

```
Runtime Error: Invalid input for real editing
Program terminated by I/O error on unit 5 (Input_Unit,Formatted,Sequential)
main.f90, line 28: Error occurred in READ_DATA
main.f90, line 57: Called by READ_COORDS
main.f90, line 40: Called by INITIAL
main.f90, line 13: Called by $main$
```

-help Display a one-line summary of the options available for the current mode (*=compiler*, *=callgraph*, *=depend* or *=polish*).

-hollerith_io

Enable Fortran-66 compatible input/output of character data stored in numeric variables using the A edit descriptor. This was superseded by the CHARACTER datatype in Fortran 77.

-hpf Accept the extensions to Fortran specified by the High Performance Fortran Forum in HPF 1.0. These consist of the EXTRINSIC keyword and a large number of compiler directives. The compiler directives are checked for correctness but have no effect on compilation.

-I *pathname*

Add *pathname* to the list of directories which are to be searched for module information (.mod) files and INCLUDE files. The current working directory is always searched first, then any directories named in *-I* options, then the compiler's library directory (see the *-Qpath* option).

-indirect *file*

Read the contents of *file* as additional arguments to the compiler driver. This option may also be given by "@*file*"; note in this case there is no space between the '@' and the file name.

In an indirect file, arguments may be given on separate lines; on a single line, multiple arguments may be separated by blanks. A blank can be included in an option or file name by putting the whole option or file name in quotes (""); this is the only quoting mechanism. An indirect file may reference other indirect files.

-ieee=*mode*

Set the mode of IEEE arithmetic operation according to *mode*, which must be one of **full**, **nonstd** or **stop**.

full enables all IEEE arithmetic facilities including non-stop arithmetic.

nonstd Disables non-stop arithmetic, terminating execution on floating overflow, division by zero or invalid operand. If the hardware supports it, this also disables IEEE gradual underflow, producing zero instead of a denormalised number; this can improve performance on some systems.

stop enables all IEEE arithmetic facilities except for non-stop arithmetic; execution will be terminated on floating overflow, division by zero or invalid operand.

The `-ieee` option must be specified when compiling the main program unit, and its effect is global. The default mode is `-ieee=stop`. For more details see the IEEE 754 Arithmetic Support section. This option is not available on IBM z9 Open Edition with hexadecimal floating point.

-info Request output of information messages. The default is to suppress these messages.

-kind=option

Specify the kind numbering system to be used; *option* must be one of **byte** or **sequential**.

For `-kind=byte`, the kind numbers for **INTEGER**, **REAL** and **LOGICAL** will match the number of bytes of storage (e.g., default **REAL** is 4 and **DOUBLE PRECISION** is 8). Note that **COMPLEX** kind numbers are the same as its **REAL** components, and thus half of the total byte length in the entity.

For `-kind=sequential` (the default), the kind numbers for all datatypes are numbered sequentially from 1, increasing with precision (e.g., default **REAL** is 1 and **DOUBLE PRECISION** is 2).

This option does not affect the interpretation of byte-length specifiers (an extension to Fortran 77).

-lx Link with library `libx.a`. The linker will search for this library in the directories specified by `-Ldir` options followed by the normal system directories (see the `ld(1)` command).

-Ldir Add *dir* to the list of directories for library files (see the `ld(1)` command).

-M Produce module information files (`.mod` files) only. This option is equivalent to `-otype=mod`.

-max_parameter_size=N

Set the maximum size of a **PARAMETER** to *N* MB (megabytes). *N* must be in the range 1 to 1048576 (1MB to 1TB); the default is 50 MB.

-maxcontin=N

Increase the limit on the number of continuation lines from 255 to *N*. This option will not decrease the limit below the standard number.

-mdir dir

Write any module information (`.mod`) files to directory *dir* instead of the current working directory.

-mismatch

Downgrade consistency checking of procedure argument lists so that mismatches produce warning messages instead of error messages. This only affects calls to a routine which is not in the current file; calls to a routine in the file being compiled must still be correct. This option disables `-C=calls`.

-mismatch_all

Further downgrade consistency checking of procedure argument lists so that calls to routines in the same file which are incorrect will produce warnings instead of error messages. This option disables `-C=calls`.

-mtrace Trace memory allocation and deallocation. This option is a synonym for `-mtrace=on`.

-mtrace=trace_opt_list

Trace memory allocation and deallocation according to the value of *trace_opt_list*, which must be a comma separated list of one or more of:

address	(display addresses),
all	(all options except for off),
line	(display file/line info if known),
off	(disable tracing output),
on	(enable tracing output),
paranoia	(protect memory allocator data structures against the user program),
size	(display size in bytes) or
verbose	(all options except for off and paranoia).

This option should be specified during both compilation and linking, and is incompatible with the `-gc` option. For more details see the Memory Tracing section. The `-mtrace=paranoia` option is not available on IBM z9 Open Edition.

- nan** Initialise `REAL` and `COMPLEX` variables to IEEE Signalling NaN, causing a runtime crash if the values are used before being set. This affects local variables, module variables, and `INTENT(OUT)` dummy arguments only; it does not affect variables in `COMMON` or `EQUIVALENCE`. This option is not available on IBM z9 Open Edition with hexadecimal floating point.
 - nihongo** Produce compiler messages in Japanese (using Shift-JIS encoding). This option is not available on IBM z9 Open Edition.
 - no_underflow_warning** Suppress the warning message that normally appears if a floating-point underflow occurred during execution. This option is only effective if specified during the link phase.
 - nocheck_modtime** Do not check for `.mod` files being out of date.
 - nomod** Suppress module information (`.mod`) file production. Combining this with `-M` will produce no output (other than error and warning messages) at all, equivalent to `-otype=none`.
 - noqueue** If no licence for the compiler is immediately available, exit with an error instead of queueing for it.
 - o output** Name the output file *output* instead of the default. If an executable is being produced the default is `a.out`; otherwise it is *file.o* with the `-c` option, *file.c* with the `-S` option, and *file.f*, *file.f90* or *file.f95* with the `-F` option, where *file* is the base part of the source file (i.e. with the suffix removed).
 - O** Normal optimisation, equivalent to `-O2`.
 - ON** Set the optimisation level to *N*. The optimisation levels are:
 - O0** No optimisation. This is the default, and is recommended when debugging.
 - O1** Minimal quick optimisation.
 - O2** Normal optimisation.
 - O3** Further optimisation.
 - O4** Maximal optimisation.
 - Oassumed** This is a synonym for `-Oassumed=contig`.
 - Oassumed=shape** Optimises assumed-shape array dummy arguments according to the value of *shape*, which must be one of
 - always_contig** Optimised for contiguous actual arguments. If the actual argument is not contiguous a runtime error will occur (the compiler is not standard-conforming under this option).
 - contig** Optimised for contiguous actual arguments; if the actual argument is not contiguous (i.e. it is an array section) a contiguous local copy is made. This may speed up array section accessing if a sufficiently large number of array element or array operations is performed (i.e. if the cost of making the local copy is less than the overhead of discontinuous array accesses), but usually makes such accesses slower. Note that this option does not affect dummy arguments with the `TARGET` attribute; these are always accessed via the dope vector.
 - section** Optimised for low-moderate accesses to array section (discontiguous) actual arguments. This is the default.
- Note that `CHARACTER` arrays are not affected by these options.
- Oblock=N** Specify the dimension of the blocks used for evaluating the `MATMUL` intrinsic. The default value (only for `-O1` and above) is system and datatype dependent.
 - Onopropagate** Disable the optimisation of constant propagation. This is the default for `-O1` and lower.

-Opropagate

Enable the optimisation of constant propagation. This is the default for *-O2* and higher.

-Orounding

Specify that the program does not alter the default rounding mode. This enables the use of faster code for the **ANINT** intrinsic.

-Ounroll=*N*

Specify the depth to which simple loops and array operations should be unrolled. The default is no unrolling (i.e. a depth of 1) for *-O0* and *-O1*, and a depth of 2 for *-O* and higher optimisation levels. It can be advantageous to disable the Fortran compiler's loop unrolling if the C compiler normally does a very good job itself — this can be accomplished with *-Ounroll=1*.

-Ounsafe

Perform possibly unsafe optimisations that may depend on the numerical stability of the program. On IBM z9 Open Edition this option, in conjunction with *-O4*, passes **NOSTRICT** to the C compiler.

-openmp

Recognise OpenMP directives and link with the OpenMP support library. For more details see the OpenMP Support section.

-otype=*filetype*

Specify the type of output file required to *filetype*, which must be one of

c	(C source file),
exe	(executable file),
fortran	(Fortran source file),
mod	(module information file),
none	(no output file),
obj	(object file).

The *-c*, *-F* and *-M* options are equivalent to *-otype=obj*, *-otype=Fortran* and *-otype=mod* respectively.

-pg

Compile code to generate profiling information which is written at run-time to an implementation-dependent file (usually *gmon.out* or *mon.out*). An execution profile may then be generated using *gprof*. This option must be specified for compilation and linking and may be unavailable on some implementations.

-pic

Produce position-independent code (small model), for use in a shared library. If the shared library is too big for the small model, use *-PIC*. This option is not available on IBM z9 Open Edition.

-PIC

Produce position-independent code (large model), for use in a shared library. This option is not available on IBM z9 Open Edition.

-Qpath *pathname*

Change the compiler library pathname from its default location to *pathname*. (The default location on Unix is usually *'/usr/local/lib/NAG_Fortran'* or *'/opt/NAG_Fortran/lib'*) This option is unnecessary on Windows as the installed location is automatically detected.

-r8

Double the size of default **REAL** and **COMPLEX**, and on machines for which quadruple-precision floating-point arithmetic is available, double the size of **DOUBLE PRECISION** (and the non-standard **DOUBLE COMPLEX**). **REAL** or **COMPLEX** specified with explicit **KIND** numbers or byte lengths are unaffected — but since the **KIND** intrinsic returns the correct values, **COMPLEX(KIND(0d0))** on a machine with quad-precision floating-point will correctly select quad-precision **COMPLEX**.

This has no effect on **INTEGER** sizes, and so the compiler is not standard-conforming in this mode.

Note: This option has been superseded by the *-double* option which doubles the size of all numeric data types.

-s

Strip symbol table information from the executable file. This option is only effective if specified during the link phase.

-S

Produce assembler (actually C source code). The resulting *.c* file should be compiled with the NAG Fortran compiler, not with the C compiler directly. This option is equivalent to *-otype=c*.

-save

This is equivalent to inserting the **SAVE** statement in all subprograms which are not declared **RECURSIVE**, thus causing all non-automatic local variables in such subprograms to be statically allocated.

-strict95

Produce obsolescence warning messages for use of ‘**CHARACTER***’ syntax. This message is not produced by default since many programs contain this syntax.

-target=*machine*

Specify the machine for which code should be generated and optimised.

- For Linux (x86-32) and Windows, *machine* may be one of
i486, **i586**, **i686**, **pentium2**, **pentium3**, **pentium4**
the specified Intel processor,
k6, **k6-2**, **k6-3**, **k6-4**, **athlon**, **athlon-4**, **athlon-xp**, **athlon-mp**
the specified AMD processor,
pentium (equivalent to *i586*) or
pentiumpro (equivalent to *i686*).

The default is to compile for i686.

- For Sun/SPARC, *machine* may be one of

V7 SPARCstation 1 et al,
V8 SPARCstation 2 et al,
super SuperSPARC,
ultra UltraSPARC or
native the current machine.

The default is to compile for SPARC V7.

Note that programs compiled for later versions of the architecture may not run, or may run much more slowly, on an earlier machine. The *-target=**native* option is not available with gcc.

- For HP9000/700, *machine* may be one of
2.0 the specified revision of the PA-RISC architecture (default) or
native the current machine.

-tempdir *directory*

Set the directory used for the compiler’s temporary files to *directory*. The default is to use the directory named by the TMPDIR environment variable, or if that is not set, /tmp.

-thread_safe

Compile code for safe execution in a multi-threaded environment. This must be specified when compiling and also during the link phase. It is incompatible with *-gc*.

-time Report execution times for the various compilation phases.**-u** Specify that IMPLICIT NONE is in effect by default, unless overridden by explicit IMPLICIT statements.**-unsharedrts**

Bind with the unshared (static) version of the Fortran runtime system; this allows a dynamically linked executable to be run on systems where the NAG Fortran Compiler is not installed. This option is only effective if specified during the link phase.

-v Verbose. Print the name of each file as it is compiled.**-V** Print version information about the compiler.**-w** Suppress all warning messages. This option is a synonym for *-w=all*.**-w=*class***

Suppress the warning messages specified by *class*, which must be one of **all**, **alloctr**, **obs**, **ques**, **uda**, **uei**, **uep**, **uip**, **ulv**, **unreffed**, **unused**, **usf**, **usy**, **x77** or **x95**.

-w=all suppresses all warning messages;

-w=alloctr suppresses warning messages about the use of allocatable components, dummy arguments and functions;

-w=obs suppresses warning messages about the use of obsolescent features;

-w=ques suppresses warning messages about questionable usage;

- w=uda** suppresses warning messages about unused dummy arguments;
- w=uei** suppresses warning messages about unused explicit imports;
- w=uep** suppresses warning messages about unused external procedures;
- w=uip** suppresses warning messages about unused intrinsic procedures;
- w=ulv** suppresses warning messages about unused local variables;
- w=unreffed** suppresses warning messages about variables set but never referenced;
- w=unused** suppresses warning messages about unused entities — this is equivalent to ‘-w=uda -w=uei -w=uep -w=uip -w=ulv -w=usf -w=usy’;
- w=usf** suppresses warning messages about unused statement functions;
- w=usy** suppresses warning messages about unused symbols;
- w=x77** suppresses extension warnings for common extensions to Fortran 77 — these are TAB format, byte-length specifiers and Hollerith constants;
- w=x95** suppresses extension warnings for extensions to Fortran 95.

-W options

The **-W** option can be used to specify the path to use for a compilation component or to pass an option directly to such a component. The possible combinations are:

- W0=path** Specify the path used for the Fortran Compiler front-end. Note that this does not affect the library directory; the **-Qpath** option should be used to specify that.
- Wc=path** Specify the path to use for invoking the C compiler; this is used both for the final stage of compilation and for linking.
- Wc,option** Pass *option* directly to the host C compiler when compiling (producing the .o file). Multiple options may be specified in a single **-Wc,** option by separating them with commas.
- Wl=path** Specify the path to use for invoking the linker (producing the executable).
- Wl,option** Pass *option* directly to the host C compiler when linking (producing the executable). Multiple options may be specified in a single **-Wl,** option by separating them with commas. A comma may be included in an option by repeating it, e.g. **-Wl,-filelist=file1,file2,file3** becomes the linker option **-filelist=file1,file2,file3**.
- Wp=path** Specify the path to use for invoking the fpp preprocessor.
- Wp,option** Pass *option* directly to fpp when preprocessing.

-wmismatch=proc-name-list

Specify a list of external procedures for which to suppress argument data type and arrayness consistency checking. The procedure names should be separated by commas, e.g. **-wmismatch=p_one,p2**. Unlike the **-mismatch** option, this only affects data type and arrayness checking, and no warning messages are produced.

-xlicinfo

Report on the availability of licences for the compiler instead of compiling anything.

-xs

(Sun/SPARC option only) Store the symbol tables in the executable (otherwise debugging is only possible if the object files are kept).

6 Files

- file.a** Library of object files.
- file.c** C source file.
- file.f** Fortran source file in fixed format (obsolete).
- file.f90** Fortran source file in free format.
- file.f95** Fortran source file in free format.

<i>file.ff</i>	Preprocessor source file for fixed-form Fortran (obsolete).
<i>file.F</i>	(Unix) Preprocessor source file for fixed-form Fortran (obsolete).
<i>file.ff90</i>	Preprocessor source file for free-form Fortran.
<i>file.F90</i>	(Unix) Preprocessor source file for free-form Fortran.
<i>file.ff95</i>	Preprocessor source file for free-form Fortran.
<i>file.F95</i>	(Unix) Preprocessor source file for free-form Fortran.
<i>name.mod</i>	Compiled module information file; <i>name</i> is the name of the module in lower case.
<i>file.o</i>	Object file
<i>/opt/NAG_Fortran/lib</i>	Default NAG Fortran Compiler library directory on Sun Solaris (see <i>-Qpath</i>); referred to as <i>library</i> hereafter.
<i>/usr/local/lib/NAG_Fortran</i>	Default NAG Fortran Compiler library directory on other Unix-based operating systems.
<i>C:\Program Files\NAG\EFBuilder 5.3\nagfor\lib</i>	Default NAG Fortran Compiler library directory on 32-bit Windows.
<i>C:\Program Files (x86)\NAG\EFBuilder 5.3\nagfor\lib</i>	Default NAG Fortran Compiler library directory on 64-bit Windows.
<i>library/f90_iostat.f90</i>	Source code for the <i>f90_iostat</i> module.
<i>library/f90_kind.f90</i>	Source code for the <i>f90_kind</i> module.
<i>library/f90_stat.f90</i>	Source code for the <i>f90_stat</i> module.
<i>library/f90_util.f90</i>	A sample Fortran 90 program that displays implementation-specific information
<i>library/iso_fortran_env.f90</i>	Source code for the <i>iso_fortran_env</i> module.
<i>library/nagfmcheck.f90</i>	Source code for the <i>nagfmcheck</i> program, see the Memory Tracing section.

7 Compilation Messages

The messages produced by the NAG Fortran Compiler itself during compilation are intended to be self-explanatory. The linker, or more rarely the host C compiler, may produce occasional messages.

Messages produced by the compiler are classified by severity level; these levels are:

Info	informational message, noting an aspect of the source code in which the user may be interested.
Warning	the source code appears likely to be in error.
Questionable	some questionable usage has been found in the source code which may indicate a programming error. This has the same severity as “warning”.
Extension	some non-standard-conforming source code has been detected but has successfully been compiled as an extension to the language. This has the same severity as “warning”.
Obsolescent	some archaic source code has been detected which although standard-conforming was classified as obsolescent by the Fortran 95 standard. This has the same severity as “warning”.

Deleted feature used

a feature that was present in Fortran 90 but deleted from the Fortran 95 standard was used. This has the same severity as “warning”.

Error

the source code does not conform to the Fortran standard or does not make sense. Compilation continues after recovery.

Fatal

a serious error in the user’s program from which the compiler cannot recover, the compilation is immediately terminated.

Panic

an internal inconsistency is found by one of the compiler’s self-checks; this is a bug in the compiler itself and NAG should be notified.

8 Compiler Limits

Item	Limit
Maximum <code>INCLUDE</code> file nesting	20
Maximum number of <code>INCLUDE</code> file references per compilation	2047
Maximum <code>DO</code> loop nesting level	199
Maximum <code>CASE</code> construct nesting level	30
Maximum <code>DATA</code> -implied- <code>DO</code> loop nesting	99
Maximum array-constructor-implied- <code>DO</code> loop nesting	99
Maximum number of dummy arguments	32767
Maximum number of arguments to <code>MIN</code> and <code>MAX</code>	100
Maximum character length	2147483647
Maximum array size (32-bit systems)	2147483647 bytes
Maximum array size (64-bit systems)	64GB
Maximum unit number	2147483647
Maximum I/O record length	2147483647

9 Input/Output Information

Item	Value
Standard error (stderr) unit number	0
Standard input (stdin) unit number	5
Standard output (stdout) unit number	6
Default maximum record length for formatted output	1024 characters
Default maximum record length for unformatted output	2147483647 bytes

The default directory used for files opened with `STATUS='SCRATCH'` is `'/tmp'` on Unix and the Windows temporary directory on Windows. This default may be overridden with the `TMPDIR` environment variable.

10 OpenMP Support

The most commonly-used features of OpenMP 3.0 are supported. The following table describes the level of support for each OpenMP directive in the initial 5.3 release.

Executable directive	Level of support
PARALLEL	Supported except for the COPYIN clause.
DO	Supported except for COLLAPSE, LASTPRIVATE and ORDERED.
SECTIONS	Supported except for LASTPRIVATE.
SINGLE	Fully supported.
MASTER	Fully supported.
WORKSHARE	Not supported.
PARALLEL DO	Supported except as noted for PARALLEL and DO.
PARALLEL SECTIONS	Fully supported.
PARALLEL WORKSHARE	Not supported.
TASK	Not supported.
MASTER	Fully supported.
CRITICAL	Fully supported.
BARRIER	Fully supported.
TASKWAIT	Not supported.
ATOMIC	Fully supported.
FLUSH	Fully supported.
ORDERED	Not supported.
Data directive/clauses	Level of support
THREADPRIVATE	Not supported.
DEFAULT	Fully supported.
SHARED	Fully supported.
PRIVATE	Fully supported.
FIRSTPRIVATE	Supported in PARALLEL directives.
LASTPRIVATE	Not supported.
REDUCTION	Fully supported.
COPYIN	Not supported.
COPYPRIVATE	Not supported.

All the procedures in section 3.2 of the OpenMP standard are supported; these are `omp_set_num_threads`, `omp_get_num_threads`, `omp_get_max_threads`, `omp_get_thread_num`, `omp_get_num_procs`, `omp_in_parallel`, `omp_set_dynamic`, `omp_get_dynamic`, `omp_set_nested`, `omp_get_nested`, `omp_set_schedule`, `omp_get_schedule`, `omp_get_thread_limit`, `omp_set_max_active_levels`, `omp_get_max_active_levels`, `get_level`, `omp_get_ancestor_thread_num`, `omp_get_team_size` and `omp_get_active_level`.

The lock procedures in section 3.3 of the OpenMP standard are not supported in release 5.3.

The timing routines in section 3.4 of the OpenMP standard are supported; these are `omp_get_wtime` and `omp_get_wtick`.

When using the IEEE arithmetic support modules, the IEEE modes (rounding, halting and underflow) are propagated into spawned OpenMP threads at the beginning of a `PARALLEL` construct, and any IEEE flag that are set by an OpenMP thread is passed back to the parent thread at the end of the `PARALLEL` construct.

The following table lists the OpenMP environment variables with their default values and, if applicable, their limits.

Environment Variable	Default	Limits
OMP_NUM_THREADS	1	1-32768
OMP_DYNAMIC	False	true or false
OMP_NESTED	False	true or false
OMP_STACKSIZE	0	<1GB (32-bit) or 16GB (64-bit)
OMP_WAIT_POLICY	None	active or passive
OMP_MAX_ACTIVE_LEVELS	1	1-64
OMP_THREAD_LIMIT	32768	1-32768

Note that although the NAG runtime supports up to 32768 threads, operating system limits may prevent usage of so many.

OpenMP is not compatible with the `-C=undefined` option.

11 Automatic File Preconnection

All logical unit numbers are automatically preconnected to specific files. These files need not exist and will only be opened or created if they are accessed with `READ` or `WRITE` without an explicit `OPEN`. By default the specific filename for unit n is `fort.n`; however if the environment variable `FORTnn` exists its value is used as the filename. Note that there are two digits in this variable name, e.g. the variable controlling unit 1 is `FORT01` whereas the default filename is `'fort.1'` (unless the prefix has been changed, see the description of module `F90_PRECONN_IO`).

A file preconnected in this manner is opened with `ACCESS='SEQUENTIAL'`. If the initial `READ` or `WRITE` is an unformatted i/o statement, it is opened with `FORM='UNFORMATTED'` otherwise it is opened with `FORM='FORMATTED'`. By default a formatted connection is opened with `BLANK='NULL'` and `POSITION='REWIND'` (see module `F90_PRECONN_IO`).

Automatic preconnection applies only to the initial use of a logical unit; once `CLOSED` the unit will not be reconnected automatically but must be explicitly `OPENed`.

Note that this facility means that it is possible for a `READ` or `WRITE` statement with an `IOSTAT=` clause to receive an i/o error code associated with the implicit `OPEN`.

12 IEEE 754 Arithmetic Support

If no floating-point option is specified, any floating divide-by-zero, overflow or invalid operand exception will cause the execution of the program to be terminated (with an informative message and usually a core dump). Occurrence of floating underflow may be reported on normal termination of the program. On hardware supporting IEEE 754 standard arithmetic gradual underflow with denormalised numbers will be enabled. Note that this mode of operation is the only one available on hardware which does not support IEEE 754.

If the `-ieee=full` option is specified, non-stop arithmetic is enabled; thus `REAL` variables may take on the values +Infinity, -Infinity and NaN (Not-a-Number). If any of the floating exceptions listed above are detected by the hardware during execution, this fact will be reported on normal termination. The `-ieee=full` option must be specified when compiling the main program and has global effect; that is, it affects the entire executable program.

If the `-ieee=nonstd` option is specified, floating-point exceptions are handled in the default manner (i.e. execution is terminated). However, gradual underflow is **not** enabled, so results which would have produced a denormalised number produce zero instead. This option can only be used on hardware for which this mode of operation is faster. Like `-ieee=full`, the `-ieee=nonstd` option must be specified when compiling the main program and has global effect.

13 Random Number Algorithm

The random number generator supplied as the intrinsic subroutine `RANDOM_NUMBER` is the “Mersenne Twister”.

14 Automatic Garbage Collection

The `-gc` option enables use of the runtime garbage collector. It is necessary to use this option during the link phase for it to have effect; specifying it additionally during the compilation phase can result in improved performance.

The supplied Technical Information note (TECHINFO) lists whether garbage collection is available for your system. If it is available, there will be a file `'gc.o'` in the compiler's library directory.

The collector used is based on version 5.3 of the publicly available general purpose garbage collecting storage allocator of Hans-J Boehm, Alan J Demers and Xerox Corporation, described in “Garbage Collection in an Uncooperative Environment” (H Boehm and M Weiser, Software Practice and Experience, September 1988, pp 807-820).

The copyright notice attached to their latest version is as follows:

Copyright 1988, 1989 Hans-J. Boehm, Alan J. Demers
Copyright (c) 1991-1995 by Xerox Corporation. All rights reserved.

Copyright 1996-1999 by Silicon Graphics. All rights reserved.
Copyright 1999 by Hewlett-Packard Company. All rights reserved.

THIS MATERIAL IS PROVIDED AS IS, WITH ABSOLUTELY NO WARRANTY EXPRESSED
OR IMPLIED. ANY USE IS AT YOUR OWN RISK.

Permission is hereby granted to use or copy this program
for any purpose, provided the above notices are retained on all copies.
Permission to modify the code and to distribute modified code is granted,
provided the above notices are retained, and a notice that the code was
modified is included with the above copyright notice.

Note that the “NO WARRANTY” disclaimer refers to the original copyright holders Boehm, Demers, Xerox Corporation, Silicon Graphics and Hewlett-Packard Company. The modified collector distributed in binary form with the NAG Fortran Compiler is subject to the same warranty and conditions as the rest of the NAG Fortran compilation system.

The module F90_GC is provided; it contains functions and variables that can control the behaviour of the garbage collector.

15 Memory Tracing

Tracing of memory allocation and deallocation is provided by the `-mtrace` option. Control is provided over whether the address, size, and line number of each allocation is displayed, or the tracing output can be suppressed entirely. A “paranoia” mode is provided where the memory allocator protects its data structures against inadvertent modification by the user program.

Runtime environment variables may be used to override the tracing options a program was built with, and to specify where to write the tracing output. These are only operative if the program was built with some tracing option; `-mtrace=off` will build a program with the tracing-capable memory allocator.

If `-mtrace=off` is not specified, use of any `-mtrace` option will implicitly do a `-mtrace=on`.

Basic tracing produces a message to the memory tracing file (normally standard error) for each allocation and deallocation, including those for automatic variables, i/o buffers and compiler-generated temporaries. Each allocation is numbered sequentially; the first three items are the i/o buffers for units 0, 5 and 6 (standard error, standard input and standard output).

All `-mtrace=` suboptions may be overridden at run time by the `NAGFORTRAN_MTRACE_OPTIONS` environment variable, which should be set to the required `trace_opt_list` (e.g. ‘`on,size`’). The memory tracing file may be specified at run time by the `NAGFORTRAN_MTRACE_FILE` environment variable.

The `-mtrace` option must be specified when linking, and is incompatible with `-gc`. Additionally, line number information is only available for those files compiled with `-mtrace=line`.

The `nagfmcheck` program can be used to check the output from the `-mtrace` option. It is designed to be used as a filter. Any lines that do not look like memory tracing output are ignored. It reports to standard output any errors it detects such as deallocating something twice, deallocating something that was never allocated, or deallocating something with a size different from that with which it was allocated. It also reports any apparent memory leaks, though this is less useful if the program terminated prematurely.

16 Undefined Variable Detection

Use of undefined variables can be detected with the `-C=undefined` option. Program units compiled with this option are incompatible with program units compiled without this option (i.e. the whole program must be compiled the same way). For this reason, `-C=undefined` is not part of `-C` or `-C=all`.

Currently, there are a number of other limitations on the use of `-C=undefined`.

1. It is incompatible with pointers in an initialised `COMMON`.
2. All intrinsic modules are available, but the `ISO_C_BINDING` module can only be used with all-Fortran programs as the option makes changes to the ABI.
3. Internal `READ` from a `CHARACTER` array requires the entire specified array subobject to be “defined”, even those elements corresponding to records not actually read.
4. Internal `WRITE` to a `CHARACTER` array is considered to define the entire specified array subobject, even those elements corresponding to records not actually written.
5. Certain intrinsic functions require the entirety of their arguments to be defined, even if some portions are not actually required for the value of the function. For example, the `PAD` argument to `RESHAPE` when no padding is actually required, and elements of the `ARRAY` argument to `PACK` that correspond to false elements of the `MASK`.
6. It is incompatible with the use of OpenMP directives.

17 Data Types

The table below lists the data types provided by the NAG Fortran Compiler together with their kind numbers. There are two possibilities for the KIND numbers: the default mode of operation (which may be specified explicitly by the `-kind=sequential` option) and the “byte” numbering scheme (specified by the `-kind=byte` option).

Type Name	KIND Number (kind=sequential)	KIND Number (kind=byte)	Description
REAL	1	4	Single precision floating-point
REAL	2	8	Double precision floating-point
REAL	3	16	Quadruple precision floating-point
COMPLEX	1	4	Single precision complex
COMPLEX	2	8	Double precision complex
COMPLEX	3	16	Quadruple precision complex
LOGICAL	1	1	Single byte logical
LOGICAL	2	2	Double byte logical
LOGICAL	3	4	Default logical
LOGICAL	4	8	Eight byte logical
INTEGER	1	1	8-bit integer
INTEGER	2	2	16-bit integer
INTEGER	3	4	32-bit (default) integer
INTEGER	4	8	64-bit integer
CHARACTER	1	1	ASCII or ISO 8859-1 character
CHARACTER	2	2	JIS X 0213 character
CHARACTER	3	3	Unicode (UCS-2) character
CHARACTER	4	4	ISO 10646 (UCS-4) character

Note that on all machines except Sun Solaris with the SunPro C compiler, quadruple precision is actually “double double” precision; this provides nearly twice the precision of Double precision but with a reduced exponent range.

The `F90_KIND` module contains named parameters useful for specifying which kind you want regardless of whether the numbering system is “sequential” or “byte”.

Additional intrinsic modules built into the NAG Fortran Compiler are described in the “`nag_modules`” document.

18 Modules

To use a module it must be previously compiled, or must be defined in the file prior to its use. When separately compiling a module the `-c` option should be specified.

Compiling a module creates a ‘.mod’ file and a ‘.o’ file. The ‘.mod’ file is used by the compiler at compile time to provide information about module contents, the ‘.o’ file (if generated) contains the code of any module procedures and must be specified when creating an executable file.

Note that the name of the ‘.mod’ file will be the name of the module, the ‘.o’ file will be named after the original source file.

When a previously compiled module is USED the NAG Fortran Compiler attempts to find its source file and, if that is successful, checks the modification times producing a warning message if the ‘.mod’ file is out of date.

19 Runtime Environment Variables

The following variables control the runtime environment for programs compiled with the NAG Fortran Compiler.

NAGFORTRAN_MTRACE_FILE

Programs compiled using any *-mtrace=* option will write the memory trace to this file. The default is standard error.

NAGFORTRAN_MTRACE_OPTIONS

Changes the memory tracing options for programs compiled using any *-mtrace=* option.

NAGFORTRAN_RUNTIME_ERROR_FILE

Runtime error messages will be written to this file. The default is standard error.

NAGFORTRAN_RUNTIME_LANGUAGE

Controls the language used for runtime error messages. This may be ‘English’ or ‘Japanese’ (not case-sensitive); the default is English.

TMPDIR Controls the directory used for scratch files (the default is system-dependent).

20 Debugging

On Windows debugging is built-in to the Fortran Builder. For operating systems other than Windows a Modern Fortran-aware debugger might be available as dbx90; see TECHINFO.txt for details.

In general, host system debuggers, such as dbx or gdb, may be used successfully on Fortran code as the names of the original source files, plus line numbers, are passed through to the intermediate C files. In using such debuggers it should be noted that most local variables have an underscore appended to their names. It may be useful to look at the intermediate C code when debugging; this is produced by the *-S* option.

21 Producing a Call Graph

The call graph generator takes a set of Fortran source files and produces a call graph with optional index and called-by tables. C files and fpp-processed files are not handled.

The call graph generator understands the following compiler options with the same meaning: *-132*, *-compatible*, *-dcfuns*, *-double*, *-dryrun*, *-dusty*, *-encoding*, *-english*, *-f2003*, *-f2008*, *-f95*, *-fixed*, *-free*, *-help*, *-hollerith_io*, *-hpf*, *-I*, *-indirect*, *-info*, *-kind*, *-maxcontin*, *-mismatch*, *-mismatch_all*, *-nihongo*, *-nocheck_modtime*, *-nomod*, *-o*, *-Qpath*, *-r8*, *-strictf95*, *-u*, *-v*, *-V*, *-w* and *-xlicinfo*.

The “@filename” syntax may also be used, with the same effect as the “*-indirect filename*” option.

The call graph is written to the file specified by the *-o* option, or to standard output if no *-o* option is specified.

The following additional options control the output produced.

-calledby

Produce a “called-by” table showing, for each routine, the routines that call it directly or indirectly. This is produced at the end of the output.

-indent=*N*

Indent by *N* for each level in the graph, up to the maximum. The default is *-indent=4*.

-indent_max=*N*

The maximum indentation is *N*. The default is *-indent_max=70*.

-index Produce an alphabetic index listing, for each routine, the line of the call graph where the routine first appears. This follows the call graph itself and precedes the called-by table (when the *-calledby* option is used).

-show_entry

Show ENTRY point names in the call graph; without this option, calls to an ENTRY point are shown as calls to the containing subprogram.

-show_generic

If a call is via a generic identifier, show the generic identifier in the call graph.

-show_host

Show the host scope names for calls to internal and module procedures.

-show_pclass

Show the class of each procedure (e.g. 'module', 'internal', ...).

-show_rename

If a called procedure was renamed on a USE statement, show the renaming.

22 Dependency Analysis

The dependency analyser takes a set of Fortran source files and produces dependency information in the form specified. C files and fpp-processed files are not handled.

The dependency analyser understands the following compiler options with the same meaning: *-132*, *-dryrun*, *-english*, *-fixed*, *-free*, *-help*, *-hpf*, *-I*, *-indirect*, *-nihongo*, *-o*, *-Qpath*, *-tempdir*, *-v* and *-V*. The “@*filename*” syntax may also be used with the same effect as the “*-indirect filename*” option.

The output form is controlled by the *-otype=type* option, where **type** is one of:

blist (the filenames as an ordered build list),
dfile (the dependencies in Makefile format, written to separate *file.d* files),
info (the dependencies as English descriptions) or
make (the dependencies in Makefile format).

The default is *-otype=info*. If *-otype=dfile* is specified, no *-o* option is permitted; otherwise, the result is written to the file specified by the *-o* option or to standard output if no *-o* option is specified.

23 Source File Polishing

The polisher takes a set of Fortran source files, which may be in fixed or free form, and produces a free form “polished” version of each file. C files and fpp-processed files are not handled.

The polisher understands the following compiler options with the same meaning: *-132*, *-encoding*, *-english*, *-fixed*, *-free*, *-help*, *-I*, *-indirect*, *-info*, *-nihongo*, *-noqueue*, *-o*, *-Qpath*, *-tempdir*, *-u*, *-v*, *-V*, *-w* and *-xlicinfo*.

The polished output is written to the file specified by the *-o* option, or to the same filename with the extension replaced by '.f90_pol' if no *-o* option is specified. The output file cannot have the same name as the input file.

The following additional options control the operation of the polisher:

-alter_comments

Enable options to alter comments; without this option, any options that would otherwise alter the comments are ignored.

-blank_cmt_to_blank_line

Turn comment lines that have no text (other than the comment-initiating character) into plain blank lines; this is the default if the *-alter_comments* option is set.

-blank_line_after_decls

Ensure that there is a blank line after the declarations and before the first executable statement; this is the default.

-bom=X

Specify whether to write a Unicode Byte-Order Mark at the beginning of the output file; *X* must be one of **Asis** (same as the input file), **Insert** (insert a byte-order mark) or **Remove** (remove any byte-order mark). This option only has effect if the input file is known to be in UTF-8 encoding, either because it begins with a byte-order mark or the *-encoding=UTF8* option was used. The default is *-bom=Asis*.

-break_long_comment_word

If a comment line will be split into two lines, the comment may be broken in the middle of a long word.

-delete_all_comments

Delete all comments (if the *-alter_comments* option is set).

-delete_blank_lines

Delete blank lines and comment lines that have no text (other than the comment-initiating character), if the *-alter_comments* option is set.

-delete_unused_labels

Delete labels that are never referenced; this is the default.

-format_start=N

If renumbering **FORMAT** statements in a separate sequence, the first **FORMAT** statement will be *N*; the default is *-format_start=90000*.

-format_step=N

If renumbering **FORMAT** statements in a separate sequence, the step from one label to the next will be *N*; the default is *-format_step=10*. Note that this may be negative (but not zero).

-idcase=X

Set the case to use for identifiers; *X* must be one of **C** (for Capitalised), **L** (for lowercase) or **U** (for UPPERCASE); the default is *-idcase=L*.

-indent=N

Indent statements within a construct by *N* spaces from the current indentation level; the default is *-indent=2*.

-indent_comment_marker

When indenting comments, the comment-initiating character should be indented to the indentation level; this is the default.

-indent_comments

Indent comments; this is the default if the *-alter_comments* option is set. The result is also affected by the *-indent_comment_marker* option.

-indent_continuation=N

Indent continuation lines by an additional *N* spaces; the default is *-indent_continuation=2*.

-indent_max=N

Set the maximum indentation level to *N* spaces; the default is *-indent_max=60*. The value must be at least 10 less than the output line length (*-width=*).

-inline_comment_indent=N

Set the indentation level for inline comments to column *N*; the default is *-inline_comment_index=35*.

-keep_blank_lines

Do not delete blank lines or comment lines with no text; this is the opposite of *-delete_blank_lines* and is the default.

-keep_comments

Do not delete non-blank comment lines; this is the opposite of *-delete_comments* and is the default.

-kind_keyword=*X*

Specifies how to handle the **KIND=** specifier in declarations; *X* must be one of **Asis** (take no action but preserve the input status), **Insert** (insert **KIND=** if not present), or **Remove** (remove **KIND=** if present); the default is *-kind_keyword=Asis*.

-kwcasing=*X*

Set the case to use for language keywords; *X* must be one of **C** (for Capitalised), **L** (for lowercase) or **U** (for UPPERCASE); the default is *-kwcasing=C*.

-label_after_indent

Indent labels; this is the opposite to *-label_before_indent*.

-label_before_indent

Output the statement label, if any, before indenting the statement; this is the default.

-leave_formats_in_place

Leave **FORMAT** statements in the same position as they are in the input file; this is the opposite of *-move_formats_to_end*, and is the default.

-margin=*N*

Set the left margin (initial indent) to *N*. The value must be at least 10 less than the output line length (*-width=*).

-name_scopes=*X*

Specify whether to add optional keywords and scope names to the **END** or **END TYPE** statement for a scope; *X* must be one of **Asis** (leave as is), **Insert** (insert keywords and/or names) or **Remove** (remove optional keywords and names). This option also applies to the **END INTERFACE** statement. The default is *-name_scopes=Insert*.

-noalter_comments

Do not alter comments in any way; this is the default.

-noblank_cmt_to_blank_line

Do not turn blank comments to blank lines.

-noblank_line_after_decls

Do not insert a blank line between the last declaration and the first executable statement.

-nobreak_long_comment_word

If a comment line will be split into two lines, do not break the comment in the middle of a long word; this is the default.

-noindent_comment_marker

Place the comment-initiating character for a comment line in column 1.

-noindent_comments

Indent the text of a comment line.

-norenumber

Do not renumber statement labels.

-noseparate_format_numbering

When renumbering statement labels, use a single sequence for both **FORMAT** and non-**FORMAT** statements; this is the default.

-noterminate_do_with_enddo

Do not change **DO** loop terminating statements.

-relational=*X*

Specifies the form to use for relational operators, *X* must be either **F77-** (use **.EQ.**, **.LE.**, etc.) or **F90+** (use **==**, **<=**, etc.); the default is *-relational=F90+*.

-renumber

Renumber statement labels; this is the default.

-renumber_start=*N*

When renumbering statement labels, the first label will be *N*; the default is *-renumber_start=100*.

-renumber_step=*N*

When renumbering statement labels, the step from one label to the next will be *N*; the default value is *-renumber_step=10*.

–`separate_format_numbering`

When renumbering statement labels, renumber `FORMAT` statements in a separate sequence from non-`FORMAT` statements.

–`terminate_do_with_enddo`

Change the terminating statements of all `DO` loops so that each loop ends with an `ENDDO` statement; this is the default.

–`width=N`

Set the maximum length of the text on each output line to N ; the default is `–width=78`. Note that in the case of continuation lines, an additional two characters (`'&'`) will be produced after the last text on a line and this may take the line length over the limit. The width must be at least 10 more than the left margin (`–margin=`) and the maximum indent (`–indent_max=`). The maximum width setting is 1024, however values higher than 130 will produce output that does not conform to the Fortran standard.

Debugging with dbx90

24 dbx90 command line

dbx90 [option]... *executable-file*

25 Description of dbx90

dbx90 is a Fortran-oriented debugger for use with the NAG Fortran Compiler on Unix-like systems (e.g. Linux, Solaris). Its syntax is quite similar to that of **dbx**, which it invokes as a sub-process to carry out the actual machine-dependent debugging commands. (On gcc-based implementations, **gdb** is used.)

The program to be debugged should be compiled and linked with the *-g90* option. This creates a debug information (*.g90*) file for each Fortran source file.

If the environment variable **DBX90_DBXPATH** is defined, **dbx90** will use it to locate the native debugger instead of the built-in path.

26 dbx90 options

-compatible

This permits **dbx90** to debug code compiled by the NAG Fortran Compiler using the *-compatible* option.

-I pathname

Add *pathname* to the list of directories which are to be searched for module information (*.mod*) files and debug information (*.g90*) files. The current working directory is always searched first, then any directories named in *-I* options, then the compiler's library directory (usually *'/usr/local/lib/NAG.Fortran'* or *'/opt/NAG.Fortran/lib'*)

-Qpath path

Set the compiler's library directory to *path*. This is only used for finding the module information (*.mod*) files for intrinsic modules.

27 dbx90 Commands

alias List command aliases.

alias name text

Create a new alias *name* with the replacement value *text*. The replacement text is not limited to a single word but may contain spaces.

assign var = expr

Assign the value of *expr* to variable *var*. The variable can be scalar, an array (including an array section), a scalar component, or an element of an array component, but cannot be of derived type. The value must be a scalar expression (see "Expressions" for more details).

cont Continue execution from where it was stopped (either by a breakpoint or an interrupt).

delete n Delete breakpoint number *n*.

delete all Delete all breakpoints.

display List the expressions to be displayed after each breakpoint is reached.

display expr

Add *expr* to the list of expressions to display after each breakpoint. *Expr* may also be an array section.

dump Display all local variables and their values.

- down** [*n*] Move the focus down (i.e. to the procedure called by the current procedure). If *n* is present, move down *n* levels.
- help** [*topic*]
Display a brief help message on the specified *topic*, or on using dbx90 generally.
- history** List the history command buffer. Old commands can be executed with:
- !!** repeat last command,
 - ! *n*** repeat command *n* in the history buffer, and
 - ! -*n*** repeat the *n*th last command.
- history *n*** Set the size of the history command buffer to *n* commands. The default size of the history command buffer is 20.
- if *expr*** This is actually a suffix to the breakpoint ('stop') commands not an independent command. It prevents the triggering of the breakpoint until the expression *expr* (a scalar expression) is **.TRUE.** (or non-zero).
- list** [*line1* [,*line2*]]
Display the next 10 lines of the program, display line *line1* of the current file or display lines *line1* to *line2* of the current file.
- next** [*n*] Execute the next *n* lines (default 1) of the current (or parent) procedure. Any procedure reference in these lines will be executed in its entirety unless a breakpoint is contained therein.
- print *expr*** [,*expr*]...
Display the value of *expr*, which can be a scalar expression, array section, derived type component, or a variable of any data type. Several expressions may be given separated by commas.
- quit** Exit from dbx90, immediately terminating any program being debugged.
- raw *dbx-command***
Pass *dbx-command* directly to "dbx". This is not recommended for normal use.
- rerun** [*command-line*]
Begin a new execution of the program, passing *command-line* to it (if present) or the *command-line* from the previous run or rerun command if not.
- run** [*command-line*]
Begin a new execution of the program, passing *command-line* to it (if present) or blank command line if not.
- scope** [*name*]
Display the current procedure name or set the focus to the specified procedure *name*.
- status** List the breakpoints which are currently set.
- step** [*n*] Execute the next *n* lines (default 1) of the program, counting lines in referenced procedures (i.e. step into procedure references).
- stop *name***
Set a breakpoint which triggers when variable *name* is accessed. Note that *name* cannot be 'at' or 'in'. This command is not available on Solaris or HP-UX.
- stop at *lineno***
Set a breakpoint at line *lineno* of the file containing the current procedure.
- stop in *name***
Set a breakpoint at the beginning of procedure *name*. Note that this breakpoint occurs at the beginning of procedure initialisation, not at the first executable statement. If there is no procedure called 'MAIN', the main program can be specified using that name.
- undisplay *expr***
Remove *expr* from the "display" list.
- up** [*n*] Move the focus up (i.e. to the caller of the current procedure). If *n* is present, move up *n* levels.

whatis *name*

Describe how *name* would be explicitly declared.

where Display the stack of active procedures with their dummy argument names.

which *name*

Display the fully qualified form of *name* which can be used for access from another scope.

28 dbx90 Expressions

28.1 Scalar expressions

Scalar expressions in dbx90 are composed of literal constants, scalar variable references including array elements, intrinsic operations and parentheses.

Literal constants can be of any intrinsic type, e.g.

```
INTEGER    42
REAL       1.2
           1.3e2
COMPLEX    (5.2,6.3)
CHARACTER  "string"
LOGICAL    .TRUE.
           .T.
```

Subscript expressions must be scalar and of type INTEGER.

All intrinsic operations are supported except for exponentiation and concatenation, that is:

`+, -, *, /, ==, /=, <, <=, >, >=, .AND., .OR., .NOT., .EQV., .NEQV., .EQ., .NE., .LT., .LE., .GT., .GE.`

(Operator names are not case-sensitive).

Note: array operations and operations involving variables of complex, character or derived type are not supported.

28.2 Array sections

Assignment, printing and displaying of array sections follows the Fortran syntax, e.g.

```
ARRAY(:)
ARRAY(1:5)
ARRAY(1:10:2)
```

If the stride is supplied it must be a positive scalar expression – negative strides are not supported. All subscript expressions must be scalar – vector subscripts are not supported.

28.3 Derived type component specification

Individual components of a derived type scalar or array may be printed using normal Fortran syntax.

For example,

```
print var%a
```

will print the “a” component of derived type “var”.

Components of all data types are supported.

Components which are of derived type will be displayed recursively until either:

- a. there are no further nested derived types, or
- b. a derived type array component is reached.

Array components of intrinsic data types will be truncated to '`<array>`', and derived type array components will be truncated to '`<derived type array>`'.

Allocatable components of derived types are supported.

Derived type assignment is not supported; however, scalar non-derived-type components may be assigned values.

29 dbx90 Command aliases

The following set of command aliases are defined:

```
a  assign
b  stop at
bp stop in
c  cont
h  history
l  list
n  next
p  print
q  quit
r  rerun
s  step
```

New aliases may be created using the `alias` command, e.g.

```
alias xp1 print x+1
```

30 dbx90 limitations

Breakpoints set at the beginning of a routine occur before procedure initialisation; at this point attempting to print or display an assumed-shape dummy argument, a variable with a complicated `EQUIVALENCE` or an automatic variable will produce dbx crashes, dbx90 crashes or nonsensical output. Execution must be stepped to the first executable statement (e.g. using the `next` command or by setting a second breakpoint) before any of these will work satisfactorily.

Strides in array sections must be positive.

31 Example of dbx90

This is an example of the use of dbx90 in debugging some Fortran code which contains both `COMMON` blocks and modules.

The file to be debugged is called '`fh4.f90`' and contains:

```
MODULE fh4
  REAL r
END MODULE fh4

PROGRAM fh4_prog
  USE fh4
  COMMON/fh4com/i
  i = 2
  CALL sub
```

```
      PRINT *,i,r
END PROGRAM fh4_prog

SUBROUTINE sub
  USE fh4
  COMMON/fh4com/i
  r = 0.5*i
  i = i*3
END SUBROUTINE sub
```

It is first compiled with the `-g90` option and then run under dbx90:

```
% nagfor -g90 -o fh4 fh4.f90
% dbx90 fh4
NAG dbx90 Version 5.2(22)
Copyright 1995-2008 The Numerical Algorithms Group Ltd., Oxford, U.K.
GNU gdb Red Hat Linux (6.5-15.fc6rh)
Copyright (C) 2006 Free Software Foundation, Inc.
GDB is free software, covered by the GNU General Public License, and you are
welcome to change it and/or distribute copies of it under certain conditions.
Type "show copying" to see the conditions.
There is absolutely no warranty for GDB. Type "show warranty" for details.
This GDB was configured as "i386-redhat-linux-gnu"...Using host libthread_db lib
rary "/lib/libthread_db.so.1".

(dbx90)
```

Setting a breakpoint in routine SUB and running the program.

```
(dbx90) stop in sub
[1] stop in SUB in file "fh4.f90"
(dbx90) run
stopped in SUB at line 16 in file "fh4.f90"
    16      r = 0.5*i
(dbx90)
```

Printing the value of a variable, which may be local, in a COMMON block, or in a USED module.

```
(dbx90) print i
I = 2
(dbx90) next
17      i = i*3
(dbx90) print r
R = 1
(dbx90) next
18      END SUBROUTINE sub
(dbx90) print i
I = 6
```

Variables can also be assigned values.

```
(dbx90) assign i = 7
I = 7
(dbx90) cont
7      1.0000000
```

```
Program exited normally.
(dbx90) quit
%
```

32 Troubleshooting dbx90

The diagnostic messages produced by dbx90 itself are intended to be self-explanatory.

If you receive the error message '**Cannot exec dbx**' when starting dbx90 then you must set the environment variable DBX90.DBXPATH to the pathname of dbx (or gdb, or xdb).

Preprocessing with fpp

33 Overview of fpp

The **fpp** preprocessor is automatically invoked by the compiler driver when the *-fpp* option is used or the source file has an option that implies preprocessing (e.g. *‘.ff90’*), but may also be invoked from the compiler library directory (usually */usr/local/lib/NAG_Fortran*).

34 fpp command line

fpp [option]... [*input-file* [*output-file*]]

35 Description of fpp

fpp is the preprocessor used by the NAG Fortran compiler. It optionally accepts two filenames as arguments: *input-file* and *output-file* are, respectively, the input file read and the output file written by the preprocessor. By default standard input and output are used.

36 fpp options

-c_com={yes|no}

By default, C style comments are recognized. Turn this off by specifying *-c_com=no*.

-Dname

Define the preprocessor variable *name* to be 1 (one). This is the same as if a *-Dname=1* option appeared on the fpp command line, or as if a

#define name 1

line appeared in the input file.

-Dname=def

Define *name* as if by a **#define** directive. This is the same as if a

#define name def

line appeared at the beginning of the input file. The *-D* option has lower precedence than the *-U* option. Thus, if the same name is used in both a *-U* option and a *-D* option, the name will be undefined regardless of the order of the options.

-e Accept extended source lines, up to 132 characters long.

-fixed Specifies fixed format input source.

-free Specifies free format input source.

-Ipathname

Add *pathname* to the list of directories which are to be searched for **#include** files whose names do not begin with *‘/’*. If the **#include** file name is enclosed in double-quotes (*"..."*), it is searched for first in the directory of the file with the **#include** line; if the file name was enclosed in angle brackets (*<...>*) this directory is not searched. Then, the file is searched for in directories named in *-I* options, and finally in directories from the standard list.

-M Generate a list of makefile dependencies and **write** them to the standard output. This list indicates that the object file which would be generated from the input file depends on the input file as well as the include files referenced.

-macro={yes|no_com|no}

By default, macros are expanded **everywhere**. Turn off macro expansion in comments by specifying *-macro=no_com* and turn off macro expansion all together by specifying *-macro=no*

- P** Do not put line numbering directives to the output file. Line numbering directives appear as *#line-number file-name*
- Uname** Remove any initial definition of *name*, where *name* is an fpp variable that is predefined by a particular preprocessor. Here is a partial list of variables that might be predefined, depending upon the architecture of the system:
- Operating System:** `unix`, `__unix` and `__SVR4`;
Hardware: `sun`, `__sun`, `sparc` and `__sparc`.
- undef** Remove initial definitions for all predefined symbols.
- w** Suppress warning messages.
- w0** Suppress warning messages.
- Xu** Convert upper-case letters to lower-case, except within character-string constants. The default is not to convert upper-case letters to lower-case.
- Xw** For fixed source form only, treat blanks as insignificant. The default for fpp is that blanks are significant in both source forms.
- Ydirectory** Use the specified *directory* instead of the standard list of directories when searching for `#include` files.

37 Using fpp

37.1 Source files

fpp operates on both fixed and free form source files. Files with the (non-case-sensitive) extension `.f`, `.ff`, `.for` or `.ftn` are assumed to be fixed form source files. All other files (e.g. those with the extension `.ff90`) are assumed to be free form source files. These assumptions can be overridden by the `-fixed` and `-free` options. Tab format lines are recognised in fixed form.

A source file may contain **fpp** tokens. An fpp token is similar to a Fortran token, and is one of:

- an fpp directive name;
- a symbolic name or Fortran keyword;
- a literal constant;
- a Fortran comment;
- an fpp comment;
- a special character which may be a blank character, a control character, or a graphic character that is not part of one of the previously listed tokens.

37.2 Output

Output consists of a modified copy of the input plus line numbering directives (unless the `-P` option is used). A line numbering directive has the form

#line-number file-name

and these are inserted to indicate the original source line number and filename of the output line that follows.

37.7 Conditional selection of source text

There are three forms of conditional selection of source text:

1.

```
#if condition_1
    block_1
#elif condition_2
    block_2
#else
    block_n
#endif
```
2.

```
#ifdef name
    block_1
#elif condition
    block_2
#else
    block_n
#endif
```
3.

```
#ifndef name
    block_1
#elif condition
    block_2
#else
    block_n
#endif
```

The “`#else`” and “`#elif`” parts are optional. There may be more than one “`#elif`” part. Each condition is an expression consisting of fpp constants, macros and macro functions. Condition expressions are similar to cpp expressions, and may contain any cpp operations and operands with the exception of C long, octal and hexadecimal constants. Additionally, fpp will accept and evaluate the Fortran logical operations `.NOT.`, `.AND.`, `.OR.`, `.EQV.`, `.NEQV.`, the relational operators `.GT.`, `.LT.`, `.LE.`, `.GE.`, and the logical constants `.TRUE.` and `.FALSE.`.

38 Preprocessing details

38.1 Scope of macro or variable definitions

The scope of a definition begins from the place of its definition and encloses all the source lines (and source lines from `#included` files) from that definition line to the end of the current file.

However, it does not affect:

- files included by Fortran `INCLUDE` lines;
- fpp and Fortran comments;
- `IMPLICIT` single letter specifications;
- `FORMAT` statements;
- numeric and character constants.

The scope of the macro effect can be limited by means of the `#undef` directive.

38.2 End of macro definition

A macro definition can be of any length but is only one logical line. These may be split across multiple physical lines by ending each line but the last with the macro continuation character ‘\’ (backslash). The backslash and newline are not part of the replacement text. The macro definition is ended by a newline that is not preceded by a backslash.

For example:

```
#define long_macro_name(x,\
y) x*y
```

38.3 Function-like macro definition

The number of macro call arguments must be the same as the number of arguments in the corresponding macro definition. An error is produced if a macro is used with the wrong number of arguments.

38.4 Cancelling macro definitions

```
#undef name
```

After this directive, ‘*name*’ will not be interpreted by fpp as a macro or variable name. This directive has no effect if ‘*name*’ is not a macro name.

38.5 Conditional source code selection

```
#if condition
```

Condition is a constant expression, as specified below. Subsequent lines up to the first matching **#elif**, **#else** or **#endif** directive appear in the output only if the *condition* is true.

The lines following a **#elif** directive appear in the output only if

- the condition in the matching **#if** directive was false,
- the conditions in all previous matching **#elif** directives were false, and
- the condition in this **#elif** directive is true.

If the condition is true, all subsequent matching **#elif** and **#else** directives are ignored up to the matching **#endif**.

The lines following a **#else** directive appear in the output only if all previous conditions in the construct were false.

The macro function ‘defined’ can be used in a constant expression; it is true if and only if its argument is a defined macro name.

The following operations are allowed.

- C language operations: <, >, ==, !=, >=, <=, +, -, /, *, %, <<, >>, &, ~, |, && and ||. These are interpreted in accordance with C language semantics, for compatibility with cpp.
- Fortran language operations: .AND., .OR., .NEQV., .XOR., .EQV., .NOT., .GT., .LT., .LE., .GE., .NE., .EQ. and **.
- Fortran logical constants: .TRUE. and .FALSE..

Only these items, integer literal constants, and names can be used within a constant expression. Names that are not macro names are treated as if they were ‘0’. The C operation ‘!=’ (not equal) can be used in **#if** or **#elif** directives, but cannot be used in a **#define** directive, where the character ‘!’ is interpreted as the start of a Fortran comment.

#ifdef *name*

Subsequent lines up to the matching **#else**, **#elif**, or **#endif** appear in the output only if the name has been defined, either by a **#define** directive or by the **-D** option, and in the absence of an intervening **#undef** directive. No additional tokens are permitted on the directive line after name.

#ifndef *name*

Subsequent lines up to the matching **#else**, **#elif**, or **#endif** appear in the output only if name has not been defined, or if its definition has been removed with an **#undef** directive. No additional tokens are permitted on the directive line after name.

#elif *constant-expression*

Any number of **#elif** directives may appear between an **#if**, **#ifdef**, or **#ifndef** directive and a matching **#else** or **#endif** directive.

#else

This inverts the sense of the conditional directive otherwise in effect. If the preceding conditional would indicate that lines are to be included, then lines between the **#else** and the matching **#endif** are ignored. If the preceding conditional indicates that lines would be ignored, subsequent lines are included in the output.

#endif

End a section of lines begun by one of the conditional directives **#if**, **#ifdef**, or **#ifndef**. Each such directive must have a matching **#endif**.

38.6 Including external files

Files are searched as follows:

for **#include "filename"**:

- in the directory, in which the processed file has been found;
- in the directories specified by the **-I** option;
- in the default directory.

for **#include <filename>**:

- in the directories specified by the **-I** option;
- in the default directory.

Fpp directives (lines beginning with the **#** character) can be placed anywhere in the source code, in particular immediately before a Fortran continuation line. The only exception is the prohibition of fpp directives within a macro call divided on several lines by means of continuation symbols.

38.7 Comments

Fpp permits comments of two kinds:

1. Fortran language comments. A source line containing one of the characters 'C', 'c', '*', 'd' or 'D' in the first column is considered to be a comment line. Within such lines macro expansions are not performed. The '!' character is interpreted as the beginning of a comment extending to the end of the line. The only exception is the case when this symbol occurs within a constant expression in a **#if** or **#elif** directive.
2. Fpp comments enclosed in the **/*** and ***/** character sequences. These are excluded from the output. Fpp comments can be nested so that for each opening sequence **/*** there must be a corresponding closing sequence ***/**. Fpp comments are suitable for excluding the compilation of large portions of source instead of commenting every line with Fortran comment symbols. Using **"#if 0 ... #endif"** achieves the same effect without being ridiculous.

38.8 Macro functions

The macro function

`defined(name)` or `defined name`

expands to `.TRUE.` if *name* is defined as a macro, and to `.FALSE.` otherwise.

38.9 Macro expression

If, during expansion of a macro, the column width of a line exceeds column 72 (for fixed form) or column 132 (for free form), fpp inserts appropriate continuation lines.

In fixed form there are limitations on macro expansion in the label part of the line (columns 1-5):

- a macro call (together with possible arguments) should not extend past column 5;
- a macro call whose name begins with one of the Fortran comment characters is treated as a comment;
- a macro expansion may produce text extending beyond the column 5 position. In such a case a warning will be issued.

In the fixed form when the `-Xw` option has been specified an ambiguity may appear if a macro call occurs in a statement position and a macro name begins or coincides with a Fortran keyword. For example, in the following text:

```
#define call p(x)    call f(x)
                call p(0)
```

fpp can not determine with certainty how to interpret the ‘call p’ token sequence. It could be considered as a macro name. The current implementation does the following:

- the longer identifier is chosen (callp in this case);
- from this identifier the longest macro name or keyword is extracted;
- if a macro name has been extracted a macro expansion is performed. If the name begins with some keyword **fpp** outputs an appropriate warning;
- the rest of the identifier is considered as a whole identifier.

In the above example the macro expansion would be performed and the following warning would be output:

warning: possibly incorrect substitution of macro callp

It should be noted that this situation appears only when preprocessing fixed form source code and when the blank character is not being interpreted as a token delimiter. It should be said also that if a macro name coincides with a keyword beginning part, as in the following case:

```
#define INT    INTEGER*8
                INTEGER k
```

then in accordance with the described algorithm, the `INTEGER` keyword will be found earlier than the `INT` macro name. Thus, there will be no warning when preprocessing such a macro definition.

39 fpp diagnostics

There are three kinds of diagnostic messages:

- warnings. preprocessing of source code is continued and the return value remains to be 0.

- errors. Fpp continues preprocessing but sets the return code to a nonzero value, namely the number of errors.
- fatal error. Fpp cancels preprocessing and returns a nonzero return value.

The messages produced by fpp are intended to be self-explanatory. The line number and filename where the error occurred are printed along with the diagnostic.

Extensions

40 Non-standard Extensions

The following extensions to the standard Fortran language are accepted by the NAG Fortran Compiler.

40.1 BOZ literal constants outside DATA statements

The NAG Fortran Compiler accepts hexadecimal, octal and binary literal constants in any context where an integer literal would be accepted. Such constants have a data type of default `INTEGER`, unless there are more digits in the `BOZ` constant than will fit in a default `INTEGER`; e.g., `Z'12345678'` would be of type default `INTEGER`, and `Z'123456789'` would be a 64-bit integer.

40.2 High Performance Fortran

All of the High Performance Fortran 1.0 language can be used by specifying the `-hpf` option. When this option is not used, any HPF directives are ignored, the `EXTRINSIC` prefix is not allowed, the `ILEN`, `NUMBER_OF_PROCESSORS` and `PROCESSORS_SHAPE` intrinsic functions are not recognised, and the `HPF_LIBRARY` module is not available.

40.2.1 HPF Directives

HPF directives are comments beginning with `!HPF$`; for full details see the *HPF Language Specification*; this appeared as a special issue of *Scientific Programming*, published by John Wiley & Sons (1993).

All HPF directives are checked for syntactic and static semantic correctness. However, note that as this is a serial (single processor) implementation, use of these directives has no effect on the code produced.

40.2.2 HPF intrinsic functions

Three new intrinsic functions are available with the `-hpf` option; these are:

ILEN This function is elemental with one argument of type `INTEGER`, returning a value of the same type. It returns one less than the number of bits needed to store the value of its argument in two's-complement arithmetic. For positive numbers, this is one more than the highest bit set (according to the model in section 13.3 of the Fortran 2003 standard, the least significant bit is bit zero); for negative numbers, it is one more than the highest bit not set.

NUMBER_OF_PROCESSORS

This is a default integer scalar function with no arguments and returns (in NAG Fortran) the value 1.

PROCESSORS_SHAPE

This function has no arguments and returns an array of type default `INTEGER`; in NAG Fortran the value returned is a zero-sized array.

40.2.3 EXTRINSIC procedures

Extrinsic procedure declarations are intended to facilitate calling non-HPF routines from an HPF program. The extrinsic clause is another prefix similar to `RECURSIVE` and `PURE`, but which can only occur within an interface block. Its syntax is:

extrinsic-prefix **is** `EXTRINSIC` (*extrinsic-kind-keyword*)
extrinsic-kind-keyword **is** `HPF` **or** `HPF_LOCAL`

Declaring a procedure to be of extrinsic kind `HPF` is the same as leaving the extrinsic clause out altogether. In NAG Fortran, declaring a procedure to be of extrinsic kind `HPF_LOCAL` has no effect.

40.3 Longer Names

Names may be up to 199 characters long.

40.4 Dollar Sign in Names

The dollar sign '\$' may appear in names as if it were a letter. On some systems this name is directly used in the link name, but on others, it is translated to `DOLLAR`.

40.5 Input/output endian/format conversion

Runtime conversion of unformatted input/output files is enabled by the `CONVERT=` specifier. This is an `OPEN` statement specifier, and takes a scalar default character argument; the accepted values are:

`BIG_ENDIAN`

synonym for `BIG_IEEE`;

`BIG_IEEE_DD`

big-endian with IEEE single and double precision floating point but with double-double for quad precision;

`BIG_IEEE` big-endian with IEEE floating point include 128-bit quad precision;

`BIG_NATIVE`

big-endian with native floating-point formats;

`LITTLE_ENDIAN`

synonym for `LITTLE_IEEE`;

`LITTLE_IEEE_DD`

little-endian with IEEE single and double precision floating point but with double-double for quad precision;

`LITTLE_IEEE`

little-endian with IEEE floating point include 128-bit quad precision;

`LITTLE_NATIVE`

little-endian with native floating-point formats;

`NATIVE` native endianness and data format.

It is also possible to specify this with an environment variable at runtime; when unit n is opened, if the environment variable `FORT_CONVERT n` exists, its value is interpreted as a `CONVERT=` specifier value. If the environment variable exists it takes precedence over any `CONVERT=` specifier in the `OPEN` statement. If the environment variable does not exist, and there is no `CONVERT=` specifier in the `OPEN` statement, the conversion mode is taken from the `-convert=` option at compile time; the default is `NATIVE` (i.e. no conversion).

The `CONVERT=` specifier is also available in the `INQUIRE` statement, and sets its argument to the current conversion mode or to 'UNKNOWN' if the file is not connected for unformatted input/output.

40.6 Elemental BIND(C) procedures

Interoperable (`BIND(C)`) procedures are permitted to be declared as `ELEMENTAL`. Such a procedure must satisfy the normal Fortran requirements for elemental procedures, in particular they must have scalar dummy arguments and must be pure (free from side-effects).

41 Obsolete Extensions

The following extensions were common in the Fortran 77 era, and are still in frequent use (though they have been superseded and are thus unnecessary). Warning messages (marked 'Extension:') are produced for each occurrence of any extension; these particular ones may be suppressed with the `-w=x77` option.

41.1 Byte Sizes

Byte sizes for `REAL`, `INTEGER`, `LOGICAL` and `COMPLEX` are accepted and mapped to Fortran 90/95 `KINDS`.

Byte Size Specification	Standard Fortran	Using <code>F90_KIND</code>
<code>REAL*4</code>	Real	<code>Real(real32)</code>
<code>REAL*8</code>	Double Precision	<code>Real(real64)</code>
<code>REAL*16</code>	<code>Real(Selected_Real_Kind(30))</code>	<code>Real(real128)</code>
<code>COMPLEX*8</code>	Complex	<code>Complex(real32)</code>
<code>COMPLEX*16</code>	<code>Complex(Kind(0d0))</code>	<code>Complex(real64)</code>
<code>COMPLEX*32</code>	<code>Complex(Selected_Real_Kind(30))</code>	<code>Complex(real128)</code>
<code>INTEGER*1</code>	<code>Integer(Selected_Int_Kind(2))</code>	<code>Integer(int8)</code>
<code>INTEGER*2</code>	<code>Integer(Selected_Int_Kind(4))</code>	<code>Integer(int16)</code>
<code>INTEGER*4</code>	Integer	<code>Integer(int32)</code>
<code>INTEGER*8</code>	<code>Integer(Selected_Int_Kind(18))</code>	<code>Integer(int64)</code>
<code>LOGICAL*1</code>	Logical(1)	<code>Logical(byte)</code>
<code>LOGICAL*2</code>	Logical(2)	<code>Logical(twobyte)</code>
<code>LOGICAL*4</code>	Logical	<code>Logical(word)</code>
<code>LOGICAL*8</code>		<code>Logical(logical64)</code>

41.2 TAB Format

The occurrence of a `TAB` character in fixed-form source is treated as follows:

- an initial `TAB` followed by a digit is expanded to 5 spaces (putting the digit in the continuation column),
- an initial `TAB` followed by any other character is expanded to 6 spaces (putting the character after the `TAB` in column 7, making the line the initial line of a statement), and
- other `TAB` characters are treated as single blanks except in character context where they remain `TABs` (but they are still treated as taking one column for the purposes of line length).

41.3 Hollerith Constants

Hollerith constants may be used as actual arguments or to initialise objects in `DATA`. In either case they must be associated with an object of intrinsic numeric data type, not with a `CHARACTER` or derived type object.

Hollerith i/o (i.e., use of the `A` edit descriptor with non-`CHARACTER` data) is only enabled if the using subprogram was compiled with the `-hollerith_io` or `-dusty` option.

41.4 Increased Line Length in Fixed Source Form

The `-132` option increases the effective length of each fixed source form input line from 72 characters to 132 characters.

Note that when this option is used the NAG Fortran Compiler no longer conforms to the Fortran language standard. The meaning of a program will change if it contains a character constant which is continued across a line boundary. A standard-conforming program containing an `H`-edit descriptor which is continued across a line boundary will very likely be rejected.

For new Fortran programs we recommend the use of free source form instead of this option. Free source form provides superior detection of typographical errors and is also part of the Fortran standard and thus fully portable to all standard-conforming compilers.

41.5 Increased Maximum Number of Continuation Lines

The `-maxcontin=N` option increases the limit on the maximum number of continuation lines to N . Note that this option can affect both fixed source form and free source form, but never decreases the continuation line limit below the standard.

The Fortran 90 and 95 standards specified that the maximum number of continuation lines in fixed source form was 19, and that the maximum number of continuation lines in free source form was 39. The Fortran 2003 standard increased this to 255 lines regardless of source form.

41.6 Legacy Application Support

The `-dusty` option downgrades the category of common errors found in “legacy” software from “Error” to “Warning”, allowing such programs to be compiled and run. (The messages may be suppressed altogether by additionally specifying the `-w` option.)

This option also effectively provides the extensions of allowing `COMMON` blocks to be initialised outside of `BLOCK DATA`, and of accepting VAX format octal and hexadecimal constants (these have the forms `'... 'O` and `'... 'X` respectively).

41.7 Mismatched Argument Lists

The `-mismatch` option downgrades checking of procedure argument lists so that inconsistencies between calls to routines which are not in the current file being processed produce warning messages rather than error messages.

The `-mismatch_all` option further downgrades argument list checking so that incorrect calls to routines present in the current file being processed produce warning messages instead of error messages.

41.8 Double Precision Complex Extensions

Double precision complex entities may be declared with the `DOUBLE COMPLEX` keywords instead of the standard Fortran `'Complex(Kind(0d0))'` specification.

If the `-defuns` option has been used, additional intrinsic functions are available (see the documentation of the option for full details). These functions have all been redundant since Fortran 90.

Intrinsic Modules

42 Intrinsic Module Overview

A number of intrinsic modules are provided that are available for use in programs. An intrinsic module is one that is pre-compiled or built in to the compiler system; several of these are part of Fortran 2003, others are specific to NAG.

The standard intrinsic modules from Fortran 2003 that are available are:

<code>ieee_arithmetic</code>	Special facilities for IEEE floating-point arithmetic
<code>ieee_exceptions</code>	IEEE arithmetic exception handling
<code>ieee_features</code>	IEEE feature control
<code>iso_c_binding</code>	Facilities for interface to C functions
<code>iso_fortran_env</code>	Fortran-specific environmental facilities

The non-standard intrinsic modules supplied by NAG are:

<code>f90_gc</code>	Garbage collector control
<code>f90_iostat</code>	Input/output error codes (source form provided)
<code>f90_kind</code>	Kind number parameters (source form provided)
<code>f90_preconn_io</code>	File preconnection control
<code>f90_stat</code>	STAT= error codes (source form provided)
<code>f90_unix_dir</code>	Unix system functions — directories and files
<code>f90_unix_dirent</code>	Unix system functions — directory reading
<code>f90_unix_env</code>	Unix system functions — environment
<code>f90_unix_errno</code>	Unix system functions — error codes
<code>f90_unix_io</code>	Unix system functions — input/output (incomplete)
<code>f90_unix_proc</code>	Unix system functions — processes
<code>hpf_library</code>	High Performance Fortran Library module

Note that although the above `f90_unix_*` modules contain Unix-specific functions, in many cases these are also usable on Windows and Mac OS X; a function that is not available can still be called but will return the `ENOSYS` error code.

43 `f90_gc`

A module is provided for controlling the garbage collector more precisely, called “`F90_GC`”. It contains the following procedures described below.

In the description of each procedure, an argument whose `KIND` is denoted by ‘*’ can accept any kind of that type. Other `KIND` indications use the named parameters from the `F90_KIND` module; these named parameters are not, however, exported from `F90_GC`.

```
SUBROUTINE ENABLE_INCREMENTAL_GC()
```

Enables incremental garbage collection; once enabled, it cannot be disabled. Full collections will still be performed with a frequency determined by the “Full GC frequency” setting (see `GET_FULL_GC_FREQUENCY` and `SET_FULL_GC_FREQUENCY`).

```
LOGICAL(word) FUNCTION EXPAND_HEAP(N)  
INTEGER(*), INTENT(IN) :: N
```

This function attempts to expand the heap by N blocks (of 4K bytes), returning `.TRUE.` if and only if it is successful. This is unaffected by the heap expansion setting (see `GET_HEAP_EXPANSION` and `SET_HEAP_EXPANSION`).

Note that on some systems this may return a false indication of success as the operating system delays the actual allocation of memory until an attempt is made to use it — at which point the program may be aborted. Therefore this function should *not* be used to attempt to allocate all the virtual memory on a system.

```
SUBROUTINE GCOLLECT()
```

Manually initiates a garbage collection; this is a full collection even if incremental collection has been enabled (see `ENABLE_INCREMENTAL_GC`). This is not affected by the “GC allowed” setting (see `GET_GC_ALLOWED` and `SET_GC_ALLOWED`).

```
SUBROUTINE GET_BYTES_SINCE_GC(NBYTES)
INTEGER(int32 or int64),INTENT(OUT) :: NBYTES
```

Returns the number of bytes allocated since the last collection.

```
SUBROUTINE GET_FREE_BYTES(NBYTES)
INTEGER(int32 or int64),INTENT(OUT) :: NBYTES
```

Returns a conservative estimate of the number of free bytes in the heap.

```
SUBROUTINE GET_FULL_GC_FREQUENCY(FREQUENCY)
INTEGER(int16 or int32 or int64),INTENT(OUT) :: FREQUENCY
```

Returns the number of partial collections that are done between each full collection when incremental collection is enabled (see `ENABLE_INCREMENTAL_GC`). This has no effect on manual collection (see `GC_COLLECT`) or when incremental collection has not been enabled. The default value in Release 5.2 is 4.

```
SUBROUTINE GET_GC_ALLOWED(ALLOWED)
LOGICAL(*),INTENT(OUT) :: ALLOWED
```

Returns the “GC allowed” setting; the default setting is `.TRUE.`

```
SUBROUTINE GET_GC_VERBOSITY(VERBOSITY)
INTEGER(*),INTENT(OUT) :: VERBOSITY
```

Returns the “GC verbosity” level; this has a range of 0-100, and the default level is zero. Increasing the verbosity level increases the number of informational messages producing during garbage collection. In Release 5.2, the only significant verbosity levels are 0 and 10.

```
SUBROUTINE GET_HEAP_EXPANSION(ALLOWED)
LOGICAL(*),INTENT(OUT) :: ALLOWED
```

Returns the automatic heap expansion setting; the default setting is `.TRUE.`. Even when automatic heap expansion is disabled, the heap will still be expanded if that is necessary to satisfy an allocation request.

```
SUBROUTINE GET_HEAP_SIZE(NBYTES)
INTEGER(int32 or int64),INTENT(OUT) :: NBYTES
```

Returns the current heap size in bytes, or -1 if the heap size cannot be represented in the integer variable (i.e. if the current heap size is more than 2G bytes and the integer variable is only a 32-bit one).

```
SUBROUTINE GET_MAX_GC_RETRIES(N_RETRIES)
INTEGER(*),INTENT(OUT) :: N_RETRIES
```

Returns the maximum number of garbage collections attempted before reporting failure to the caller; if the allocation was not initiated by an `ALLOCATE` statement with a `STAT=` clause, this will result in program termination. The default value in Release 5.2 is zero.

```
SUBROUTINE GET_MAX_HEAP_SIZE(NBYTES)
INTEGER(int32 or int64),INTENT(OUT) :: NBYTES
```

Returns the maximum heap size setting, or zero if no maximum heap size has been set. The default setting is zero.

```
INTEGER(int32) FUNCTION NCOLLECTIONS()
```

Returns the number of garbage collections that have been performed.

```
SUBROUTINE SET_FULL_GC_FREQUENCY(FREQUENCY)
INTEGER(int16 or int32 or int64), INTENT(IN) :: FREQUENCY
```

Sets the number of partial garbage collections to be done between each full collection; this has no effect unless incremental collection has been enabled (see `ENABLE_INCREMENTAL_GC`). A full collection may still be done if there is a substantial increase in the number of in-use blocks.

```
SUBROUTINE SET_GC_ALLOWED(ALLOWED)
LOGICAL(*), INTENT(IN) :: ALLOWED
```

Controls the “GC allowed” setting; when this setting is `.FALSE.`, automatic garbage collection is inhibited. This does not affect manual garbage collection (see `GCOLLECT`).

```
SUBROUTINE SET_GC_VERBOSITY(VERBOSITY)
INTEGER(*), INTENT(IN) :: VERBOSITY
```

Sets the verbosity level; the default value is zero, and the range is limited to 0-100. In Release 5.2, the only significant levels are 0 and 10.

```
SUBROUTINE SET_HEAP_EXPANSION(ALLOWED)
LOGICAL(*), INTENT(IN) :: ALLOWED
```

Controls whether automatic heap expansion is allowed; if this is `.FALSE.`, automatic heap expansion will only occur if necessary to satisfy an allocation request (normally the heap is expanded when an heuristic determines that it would be advantageous). This does not affect manual heap expansion (see `EXPAND_HEAP`).

```
SUBROUTINE SET_MAX_GC_RETRIES(N_RETRIES)
INTEGER(*), INTENT(IN) :: N_RETRIES
```

Sets the maximum number of garbage collections attempted before reporting out of memory after heap expansion fails (i.e. is refused by the operating system). The default setting in Release 5.2 is zero.

```
SUBROUTINE SET_MAX_HEAP_SIZE(N)
INTEGER(int32 or int64), INTENT(IN) :: N
```

Sets the maximum size of the heap to N bytes. This will prevent the heap from automatic expansion beyond the specified limit, and prevent it from automatic expansion entirely if it is already beyond the limit.

44 f90_iostat

This module contains definitions of integer parameters for all the `IOSTAT` values that can be returned as a result of use of an input/output or data transfer statement.

For example:

```
USE f90_iostat
INTEGER ios
OPEN (10, FILE='a.b', IOSTAT=ios, STATUS='NEW')
IF (ios==IOERR_NEW_FILE_EXISTS) PRINT *, "File a.b existed already"
```

45 f90_kind

This module contains definitions of integer parameters that can be used as kind numbers. Users wishing to write portable software making use of non-default kinds should **USE** this module and use the parameters instead of numeric values. For example, users should use **LOGICAL(KIND=BYTE)** instead of **LOGICAL(KIND=1)**. The available **KIND** parameters are shown below; their exact meanings (i.e. the values they represent) are implementation dependent.

INTEGER,PARAMETER :: SINGLE

For **REAL** and **COMPLEX**, selects the default real or default complex kind; this is equivalent to leaving the **KIND** selector off entirely.

INTEGER,PARAMETER :: DOUBLE

Selects the double precision real kind; this is equivalent to declaring **REAL** entities using the **DOUBLE PRECISION** type specifier, to declaring **COMPLEX** entities using **COMPLEX(KIND(0d0))**, and to using the exponent letter **D** on literal constants.

INTEGER,PARAMETER :: QUAD

REAL/COMPLEX kind selector for real and complex types with approximately twice the precision of **DOUBLE**. This might not be available on some systems; on a system without this type, the value of this parameter will be negative.

INTEGER,PARAMETER :: REAL32

REAL/COMPLEX kind selector for real and complex types that are represented using 32-bit floating-point numbers.

INTEGER,PARAMETER :: REAL64

REAL/COMPLEX kind selector for real and complex types that are represented using 64-bit floating-point numbers.

INTEGER,PARAMETER :: REAL64x2

REAL/COMPLEX kind selector for real and complex types that are represented using “double-double” floating-point numbers. A double-double floating-point number consists of two 64-bit values, one of which is at least **DIGITS(1._REAL64)** smaller than the other; this has almost twice the precision of **REAL64** (except when near zero), but a smaller exponent range.

This type is not available on all systems; on a system without this type, the value of this parameter is -1 .

INTEGER,PARAMETER :: REAL128

REAL/COMPLEX kind selector for real and complex types that are represented using 128-bit floating-point numbers. This will select a “true 128-bit” floating-point type if one is available, and if not it will select a “double-double” floating-point type if that is available; if no 128-bit floating-point type is available the value of this parameter is -1 .

INTEGER,PARAMETER :: INT8

INTEGER kind selector for integer types with at least 8 bits of precision.

INTEGER,PARAMETER :: INT16

INTEGER kind selector for integer types with at least 16 bits of precision.

`INTEGER,PARAMETER :: INT32`

INTEGER kind selector for integer types with at least 32 bits of precision.

`INTEGER,PARAMETER :: INT64`

INTEGER kind selector for integer types with at least 64 bits of precision.

`INTEGER,PARAMETER :: BYTE`

LOGICAL kind selector for logical types occupying only one byte of memory.

`INTEGER,PARAMETER :: TWOBYTE`

LOGICAL kind selector for logical types occupying the same space as `INTEGER(INT16)` entities.

`INTEGER,PARAMETER :: WORD`

LOGICAL kind selector for a 32-bit logical type.

`INTEGER,PARAMETER :: LOGICAL64`

LOGICAL kind selector for a 64-bit logical type.

`INTEGER,PARAMETER :: ASCII`

CHARACTER kind selector for the ASCII character set.

`INTEGER,PARAMETER :: JIS`

CHARACTER kind selector for the JIS X 0213:2004 character set.

`INTEGER,PARAMETER :: UCS2`

CHARACTER kind selector for the UCS-2 (Unicode) character set.

`INTEGER,PARAMETER :: UCS4`

CHARACTER kind selector for the UCS-4 (ISO 10646) character set.

46 f90_preconn_io

This module enables alteration of the default values used by automatically preconnected files.

46.1 Procedures

Note that in the descriptions below, `LOGICAL(*)` means any kind of `LOGICAL`. Also, note that although the `VERBOSE` argument is an optional argument, the actual argument is not permitted to be a non-present optional dummy argument.

```
SUBROUTINE GET_PCIO_OPTIONS(BLANK,POSITION,PREFIX,VERBOSE)
CHARACTER(*),OPTIONAL,INTENT(OUT) :: BLANK,POSITION,PREFIX
LOGICAL(*),OPTIONAL,INTENT(OUT) :: VERBOSE
```

Returns the current settings for file preconnection; the meanings for these settings are described under the `IOINIT` routine. Both the `BLANK` and `POSITION` values are returned in upper case, even if lower case was used in a call to `IOINIT` to set them.

```
SUBROUTINE IOINIT(BLANK,POSITION,PREFIX,VERBOSE)
CHARACTER(*),OPTIONAL,INTENT(IN) :: BLANK,POSITION,PREFIX
LOGICAL(*),OPTIONAL,INTENT(IN) :: VERBOSE
```

This procedure controls automatic file preconnection. Files can be automatically preconnected when a logical unit is initially referenced without an `OPEN` statement; preconnection does not occur again after file `CLOSE`.

The name of the file to be preconnected is determined first by searching for an environment variable of the form `FORTnn` where *nn* is the two-digit logical unit number; e.g. `FORT03`. If this environment variable is found its value is used as the filename; otherwise the name “`fort.n`” is used; e.g. ‘`fort.3`’. The file is opened either with `FORM='FORMATTED'` or `FORM='UNFORMATTED'` depending on whether the initial reference is with a formatted or unformatted i/o statement.

The `BLANK` and `POSITION` arguments control the `BLANK=` and `POSITION=` keywords in the implied `OPEN` statement when the file is preconnected; initially these are set to `BLANK='NULL'` and `POSITION='REWIND'`. Note that these have no meaning (and no effect) for unformatted files.

The `PREFIX` argument changes the *prefix* used to find the environment variable containing the preconnected file name. Initially this is ‘`FORT`’. Only the first 30 characters of `PREFIX` are used.

The `VERBOSE` argument controls activity reporting; if it is `.TRUE.`, subsequent preconnection activity will produce informative messages on the standard error unit.

46.2 Example

```
USE F90_PRECONN_IO
CALL IOINIT(BLANK='ZERO',PREFIX='MYFILE')
READ (99,10) I,J,K,L,M
10 FORMAT(I3,I1,I3,I1,I3)
```

When the `READ` statement is executed, the environment variable `MYFILE99` is interrogated to discover which file is to be automatically preconnected to unit 99. The unit will be connected with `BLANK='ZERO'`. If no environment variable is found the file ‘`fort.99`’ will be opened.

47 f90_stat

This module contains definitions of integer parameters for all the `STAT=` values that can be returned as a result of use of an `ALLOCATE` or `DEALLOCATE` statement.

47.1 Parameters

```
INTEGER,PARAMETER :: STAT_ALREADY_ALLOCATED
```

An allocatable variable in an `ALLOCATE` statement is already currently allocated.

```
INTEGER,PARAMETER :: STAT_MEMORY_LIMIT_EXCEEDED
```

An allocation in an `ALLOCATE` statement requested more memory than the limit in this version of the NAG Fortran compiler.

```
INTEGER,PARAMETER :: STAT_NO_MEMORY
```


Insufficient free memory available to satisfy the requested allocation.

```
INTEGER,PARAMETER :: STAT_NOT_ALLOCATED
```

An allocatable variable in a DEALLOCATE statement is not currently allocated.

```
INTEGER,PARAMETER :: STAT_NOT_ASSOCIATED
```

A pointer in a DEALLOCATE statement is disassociated.

```
INTEGER,PARAMETER :: STAT_PART_OF_A_LARGER_OBJECT
```

A pointer in a DEALLOCATE statement refers to part of a larger object.

```
INTEGER,PARAMETER :: STAT_POINTER_UNDEFINED
```

A pointer in a DEALLOCATE statement is undefined. (This value is never returned to the user program, which is always immediately terminated if the use of an undefined pointer is detected.)

```
INTEGER,PARAMETER :: STAT_WRONG_COLOUR
```

A pointer in a DEALLOCATE statement is associated with a target that was not created by pointer allocation.

47.2 Example

```
USE f90_stat
REAL,ALLOCATABLE :: big(:,:,:)
INTEGER :: status
ALLOCATE(big(100,1024,1024),STAT=status)
IF (status==STAT_NO_MEMORY) PRINT *, 'Out of memory'
```

48 f90_unix_*

These modules are described in the next part, *Modern Fortran API to Posix*.

49 ieee_*, iso_c_binding, iso_fortran_env

These modules are described in the *Fortran 2003 Extensions* part, under *10.6 IEEE arithmetic support*, *10.5.1 The ISO_C_BINDING module* and *10.8.5 The ISO_FORTTRAN_ENV module* respectively.

50 hpf_library

This module is described in *High Performance Fortran Library Module*.

Modern Fortran API to Posix

51 Posix Module Overview

The following modules are provided by NAG as a partial interface to the operating system facilities defined by ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) – Part 1: System Application Program Interface (API) [C Language].

<code>f90_unix_dir</code>	Directories and Files
<code>f90_unix_dirent</code>	Directory Reading
<code>f90_unix_env</code>	Environment
<code>f90_unix_errno</code>	Error Codes
<code>f90_unix_file</code>	File Characteristics
<code>f90_unix_io</code>	Input/Output (incomplete)
<code>f90_unix_proc</code>	Processes

52 `f90_unix_dir`

This module contains part of a Fortran API to functions detailed in ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) [C Language].

The procedures in this module are from sections 5.2: Working Directory, 5.3.3 Set File Creation Mask, 5.3.4 Link to a File, 5.4 Special File Creation and 5.5 File Removal.

Error handling is described in `F90_UNIX_ERRNO`. Note that for procedures with an optional `ERRNO` argument, if an error occurs and `ERRNO` is not present, the program will be terminated.

All the procedures in this module are both generic and specific.

52.1 Parameters

```
INTEGER,PARAMETER :: MODE_KIND
```

The integer kind used to represent file permissions (see ISO/IEC 9945-1). Parameters for specific permissions are contained in `F90_UNIX_FILE`.

52.2 Procedures

```
SUBROUTINE CHDIR(PATH,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Sets the current working directory to `PATH`. Note that any trailing blanks in `PATH` may be significant. If `ERRNO` is present it receives the error status.

Possible error conditions include `EACCES`, `ENAMETOOLONG`, `ENOTDIR` and `ENOENT` (see `F90_UNIX_ERRNO`).

```
SUBROUTINE GETCWD(PATH,LENPATH,ERRNO)
  CHARACTER(*),OPTIONAL,INTENT(OUT) :: PATH
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: LENPATH
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Accesses the current working directory information. If `PATH` is present, it receives the name of the current working directory, blank-padded or truncated as appropriate if the length of the current working directory name differs from

that of PATH. If LENPATH is present, it receives the length of the current working directory name. If ERRNO is present it receives the error status.

If neither PATH nor LENPATH is present, error EINVAL is raised. If the path to current working directory cannot be searched, error EACCES is raised. If PATH is present and LENPATH is not present, and PATH is shorter than the current working directory name, error ERANGE is raised. (See F90_UNIX_ERRNO).

```
SUBROUTINE LINK(EXISTING,NEW,ERRNO)
  CHARACTER(*),INTENT(IN) :: EXISTING,NEW
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Creates a new link (with name given by NEW) for an existing file (named by EXISTING).

Possible errors include EACCES, EEXIST, EMLINK, ENAMETOOLONG, ENOENT, ENOSPC, ENOTDIR, EPERM, EROFS, EXDEV (see F90_UNIX_ERRNO).

```
SUBROUTINE MKDIR(PATH,MODE,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  INTEGER(mode_kind),INTENT(IN) :: MODE
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Creates a new directory with name given by PATH and mode MODE (see F90_UNIX_FILE for mode values). Note that any trailing blanks in PATH may be significant.

Possible errors include EACCES, EEXIST, EMLINK, ENAMETOOLONG, ENOENT, ENOSPC, ENOTDIR and EROFS (see F90_UNIX_ERRNO).

```
SUBROUTINE MKFIFO(PATH,MODE,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  INTEGER(mode_kind),INTENT(IN) :: MODE
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Creates a new FIFO special file with name given by PATH and mode MODE. Note that any trailing blanks in PATH may be significant.

Possible errors include EACCES, EEXIST, ENAMETOOLONG, ENOENT, ENOSPC, ENOTDIR and EROFS (see F90_UNIX_ERRNO).

```
SUBROUTINE RENAME(OLD,NEW,ERRNO)
  CHARACTER(*),INTENT(IN) :: OLD
  CHARACTER(*),INTENT(IN) :: NEW
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Changes the name of the file OLD to NEW. Any existing file NEW is first removed. Note that any trailing blanks in OLD or NEW may be significant.

Possible errors include EACCES, EBUSY, EEXIST, ENOTEMPTY, EINVAL, EISDIR, ENAMETOOLONG, EMLINK, ENOENT, ENOSPC, ENOTDIR, EROFS and EXDEV (see F90_UNIX_ERRNO).

```
SUBROUTINE RMDIR(PATH,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Removes the directory PATH. Note that any trailing blanks in PATH may be significant.

Possible errors include EACCES, EBUSY, EEXIST, ENOTEMPTY, ENAMETOOLONG, ENOENT, ENOTDIR and EROFS (see F90_UNIX_ERRNO).

```
SUBROUTINE UMASK(CMASK,PMASK)
  INTEGER(mode_kind),INTENT(IN) :: CMASK
  INTEGER(mode_kind),OPTIONAL,INTENT(OUT) :: PMASK
```

Sets the file mode creation mask of the calling process to `CMASK`. If `PMASK` is present it receives the previous value of the mask.

```
SUBROUTINE UNLINK(PATH,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Deletes the file `PATH`. Note that any trailing blanks in `PATH` may be significant.

Possible errors include `EACCES`, `EBUSY`, `ENAMETOOLONG`, `ENOENT`, `ENOTDIR`, `EPERM` and `EROFS` (see `F90_UNIX_ERRNO`).

53 f90_unix_dirent

This module contains part of a Fortran API to functions detailed in ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) [C Language].

The functions in this module are from Section 5.1.2: Directory Operations

Error handling is described in `F90_UNIX_ERRNO`. Note that for procedures with an optional `ERRNO` argument, if an error occurs and `ERRNO` is not present, the program will be terminated.

All the procedures in this module are specific and not generic.

53.1 Procedures

In the description of each procedure, an argument whose `KIND` is denoted by ‘*’ can accept any kind of that type. Other `KIND` indications use the named parameters from the `F90_KIND` or `F90_UNIX_ERRNO` modules; these named parameters are not, however, exported from `F90_UNIX_DIRENT`.

```
SUBROUTINE CLOSEDIR(DIRUNIT,ERRNO)
  INTEGER(*),INTENT(IN) :: DIRUNIT
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Close a directory stream that was opened by `OPENDIR`.

If `DIRUNIT` does not refer to an open directory stream, error `EBADF` (see `F90_UNIX_ERRNO`) is raised.

```
SUBROUTINE OPENDIR(DIRNAME,DIRUNIT,ERRNO)
  CHARACTER(*),INTENT(IN) :: DIRNAME
  INTEGER(*),INTENT(OUT) :: DIRUNIT
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Opens a directory stream, returning a handle to it in `DIRUNIT`.

Possible errors include `EACCES`, `ENAMETOOLONG`, `ENOENT`, `ENOTDIR`, `EMFILE` and `ENFILE` (see `F90_UNIX_ERRNO`).

```
SUBROUTINE REaddir(DIRUNIT,NAME,LENNAME,ERRNO)
  INTEGER(*),INTENT(IN) :: DIRUNIT
  CHARACTER(*),INTENT(OUT) :: NAME
  INTEGER(int32 or int64),INTENT(OUT) :: LENNAME
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Reads the first/next directory entry. The name of the file is placed into `NAME`, blank-padded or truncated as appropriate if the length of the file name differs from `LEN(NAME)`. The length of the file name is placed in `LENNAME`.

If there are no more directory entries, `NAME` is unchanged and `LENNAME` is negative.

If `DIRUNIT` is not a directory stream handle produced by `OPENDIR`, or has been closed by `CLOSEDIR`, error `EBADF` (see `F90_UNIX_ERRNO`) is raised.

```
SUBROUTINE REWINDDIR(DIRUNIT,ERRNO)
  INTEGER(*),INTENT(IN) :: DIRUNIT
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Rewinds the directory stream so that the next call to `REaddir` on that stream will return the name of the first file in the directory.

54 f90_unix_env

This module contains part of a Fortran API to functions detailed in ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) [C Language].

The functions in this module are from Section 4: Process Environment, plus `gethostname` from 4.3BSD.

Error handling is described in `F90_UNIX_ERRNO`. Note that for procedures with an optional `ERRNO` argument, if an error occurs and `ERRNO` is not present, the program will be terminated.

All the procedures in this module are generic; some are also specific (this may change in a future release).

54.1 Parameters

```
INTEGER(int32),PARAMETER :: clock_tick_kind
```

The integer kind used for clock ticks (see `TIMES`).

```
INTEGER(int32),PARAMETER :: id_kind
```

Integer kind for representing IDs; that is, users (`uids`), groups (`gids`), process groups (`pgids`) and processes (`pids`).

```
INTEGER(int32),PARAMETER :: long_kind
```

The integer kind corresponding to the C type 'long'. This is not actually used in this module, but is suitable for use when declaring a variable to receive a value from `SYSCONF`.

```
INTEGER(int32),PARAMETER :: sc_stdin_unit, sc_stdout_unit, sc_stderr_unit,
  sc_arg_max, sc_child_max, sc_clk_tck, sc_job_control, sc_open_max,
  sc_ngroups_max, sc_saved_ids, sc_stream_max, sc_tzname_max, sc_version
```

Values used as arguments to `SYSCONF`. The following table describes the returned information from `SYSCONF`, this is not the value of the `SC_*` constant.

`SC_STDIN_UNIT`

The logical unit number for standard input (`READ *,...`); this is the same value as `INPUT_UNIT` in the standard intrinsic module `ISO_FORTRAN_ENV`.

`SC_STDOUT_UNIT`

The logical unit number for standard output (`PRINT, WRITE(*,...)`); this is the same value as `OUTPUT_UNIT` in the standard intrinsic module `ISO_FORTRAN_ENV`.

SC_STDERR_UNIT

The logical unit number on which errors are reported; this is the same value as **ERROR_UNIT** in the standard intrinsic module **ISO_FORTRAN_ENV**.

SC_ARG_MAX

Maximum length of arguments for the **EXEC** functions, in bytes, including environment data.

SC_CHILD_MAX

Maximum number of simultaneous processes for a single user.

SC_CLK_TCK

Number of clock ticks per second. (This is the same value returned by the **CLK_TCK** function).

SC_JOB_CONTROL

Value available only if job control is supported by the operating system.

SC_NGROUPS_MAX

Maximum number of simultaneous supplementary group IDs per process.

SC_OPEN_MAX

Maximum number of files open simultaneously by a single process.

SC_SAVED_IDS

Value available only if each process has a saved set-uid and set-gid.

SC_STREAM_MAX

Maximum number of logical units that can be simultaneously open. Not always available.

SC_TZNAME_MAX

Maximum number of characters for the name of a time zone.

SC_VERSION

Posix version number. This will be 199009 if the underlying operating system's C interface conforms to ISO/IEC 9945-1:1990.

INTEGER(int32),PARAMETER :: time_kind

The integer kind used for holding date/time values (see **TIME**).

54.2 Types

TYPE tms

INTEGER(clock_tick_kind) utime, stime, cutime, cstime

END TYPE

Derived type holding CPU usage time in clock ticks. **UTIME** and **STIME** contain CPU time information for a process, **CUTIME** and **CSTIME** contain CPU time information for its terminated child processes. In each case this is divided into user time (**UTIME**, **CUTIME**) and system time (**STIME**, **CSTIME**).

TYPE utsname

CHARACTER(...) sysname, nodename, release, version, machine

END TYPE

Derived type holding data returned by **UNAME**. Note that the character length of each component is fixed, but may be different on different systems. The values in these components are blank-padded (if short) or truncated (if long). For further information see ISO/IEC 9945-1:1990.

54.3 Procedures

```
PURE INTEGER(KIND=clock_tick_kind) FUNCTION clk_tck()
```

Returns the number of clock ticks in one second of CPU time (see `TIMES`).

```
PURE SUBROUTINE ctermid(s,lens)
  CHARACTER(*),OPTIONAL,INTENT(OUT) :: s
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: lens
```

If present, `S` is set to the filename of the controlling terminal. If present, `LENS` is set to the length of the filename of the controlling terminal. If `S` is longer than the filename of the controlling terminal it is padded with blanks; if `S` is shorter it is truncated — it is the user's responsibility to check the value of `LENS` to detect such truncation.

If the filename of the controlling terminal cannot be determined for any reason `LENS` (if present) will be set to zero and `S` (if present) will be blank.

```
SUBROUTINE getarg(k,arg,lenarg,errno)
  INTEGER(*),INTENT(IN) :: k
  CHARACTER(*),OPTIONAL,INTENT(OUT) :: arg
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: lenarg
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Accesses command-line argument number `K`, where argument zero is the program name. If `ARG` is present, it receives the argument text (blank-padded or truncated as appropriate if the length of the argument differs from that of `ARG`). If `LENARG` is present, it receives the length of the argument. If `ERRNO` is present, it receives the error status.

Note that if `K` is less than zero or greater than the number of arguments (returned by `IARGC`) error `EINVAL` (see `F90_UNIX_ERRNO`) is raised.

This procedure is very similar to the standard intrinsic procedure `GET_COMMAND_ARGUMENT`.

```
PURE INTEGER(id_kind) FUNCTION getegid()
```

Returns the effective group number of the calling process.

```
SUBROUTINE getenv(name,value,lenvalue,errno)
  CHARACTER(*),INTENT(IN) :: name
  CHARACTER(*),OPTIONAL,INTENT(OUT) :: value
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: lenvalue
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Accesses the environment variable named by `NAME`. If `VALUE` is present, it receives the text value of the variable (blank-padded or truncated as appropriate if the length of the value differs from that of `VALUE`). If `LENVALUE` is present, it receives the length of the value. If `ERRNO` is present, it receives the error status.

If there is no such variable, error `EINVAL` (see `F90_UNIX_ERRNO`) is raised. Other possible errors include `ENOMEM`.

```
PURE INTEGER(id_kind) FUNCTION geteuid()
```

Returns the effective user number of the calling process.

```
PURE INTEGER(id_kind) FUNCTION getgid()
```

Returns the group number of the calling process.

```
SUBROUTINE getgroups(grouplist,ngroups,errno)
  INTEGER(id_kind),OPTIONAL,INTENT(OUT) :: grouplist(:)
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: ngroups
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Retrieves supplementary group number information for the calling process. If `GROUPLIST` is present, it is filled with the supplementary group numbers. If `NGROUPS` is present, it receives the number of supplementary group numbers. If `ERRNO` is present, it receives the error status.

If `GROUPLIST` is present but too small to contain the complete list of supplementary group numbers, error `EINVAL` (see `F90_UNIX_ERRNO`) is raised. The maximum number of supplementary group numbers can be found using `SYSCONF` (enquiry `SC_NGROUPS_MAX`); alternatively, `'CALL GETGROUPS(NGROUPS=N)'` will reliably return the actual number in use.

```
PURE SUBROUTINE gethostname(name,lenname)
  CHARACTER(*),OPTIONAL,INTENT(OUT) :: name
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: lenname
```

This provides the functionality of 4.3BSD's *gethostname*. If `NAME` is present it receives the text of the standard host name for the current processor, blank-padded or truncated if appropriate. If `LENNAME` is present it receives the length of the host name. If no host name is available `LENNAME` will be zero.

```
PURE SUBROUTINE getlogin(s,lens)
  CHARACTER(*),OPTIONAL,INTENT(OUT) :: s
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: lens
```

Accesses the user name (login name) associated with the calling process. If `S` is present, it receives the text of the name (blank-padded or truncated as appropriate if the length of the login name differs from that of `S`). If `LENS` is present, it receives the length of the login name.

```
PURE INTEGER(id_kind) FUNCTION getpgrp()
```

Returns the process group number of the calling process.

```
PURE INTEGER(id_kind) FUNCTION getpid()
```

Returns the process number of the calling process.

```
PURE INTEGER(id_kind) FUNCTION getppid()
```

Returns the process number of the parent of the calling process.

```
PURE INTEGER(id_kind) FUNCTION getuid()
```

Returns the user number of the calling process.

```
PURE INTEGER(int32) FUNCTION iargc()
```

Returns the number of command-line arguments; this is the same value as the intrinsic function `COMMAND_ARGUMENT_COUNT`, except that it returns -1 if even the program name is unavailable (the intrinsic function erroneously returns the same value, 0, whether the program name is available or not).

```
SUBROUTINE isatty(lunit,answer,errno)
  INTEGER(*),INTENT(IN) :: lunit
  LOGICAL(*),INTENT(OUT) :: answer
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```


ANSWER receives the value `.TRUE.` if and only if the logical unit identified by `LUNIT` is connected to a terminal.

If `LUNIT` is not a valid unit number or is not connected to any file, error `EBADF` (see `F90_UNIX_ERRNO`) is raised.

```
SUBROUTINE setgid(gid,errno)
  INTEGER(*),INTENT(IN) :: gid
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Sets the group number of the calling process to `GID`. For full details refer to section 4.2.2 of ISO/IEC 9945-1:1990.

If `GID` is not a valid group number, error `EINVAL` (see `F90_UNIX_ERRNO`) is raised. If the process is not allowed to set the group number to `GID`, error `EPERM` is raised.

```
SUBROUTINE setpgid(pid,pgid,errno)
  INTEGER(*),INTENT(IN) :: pid, pgid
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Sets the process group number of process `PID` (or, if `PID` is zero, the calling process) to `PGID`. For full details refer to section 4.3.3 of ISO/IEC 9945-1:1990.

Possible errors include `EACCES`, `EINVAL`, `ENOSYS`, `EPERM`, `ESRCH` (see `F90_UNIX_ERRNO`).

```
SUBROUTINE setsid(sid,errno)
  INTEGER(id_kind),OPTIONAL,INTENT(OUT) :: sid
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Creates a session and sets the process group number of the calling process. For full details refer to section 4.3.2 of ISO/IEC 9945-1:1990. If `SID` is present it receives the new session id (equal to the process id); if an error occurs it receives `-1`.

Possible errors include `EPERM` (see `F90_UNIX_ERRNO`).

```
SUBROUTINE setuid(uid,errno)
  INTEGER(*),INTENT(IN) :: uid
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Sets the user number of the calling process to `UID`. For full details refer to section 4.2.2 of ISO/IEC 9945-1:1990.

If `UID` is not a valid group number, error `EINVAL` (see `F90_UNIX_ERRNO`) is raised. If the process is not allowed to set the user number to `UID`, error `EPERM` is raised.

```
SUBROUTINE sysconf(name,val,errno)
  INTEGER(*),INTENT(IN) :: name
  INTEGER(*),INTENT(OUT) :: val
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Returns the value of a system configuration variable. The variables are named by integer `PARAMETER`s defined in this module, and are described in the **Parameters** section.

If `NAME` is not a valid configuration variable name, error `EINVAL` (see `F90_UNIX_ERRNO`) is raised. If `VAL` is too small an integer kind to contain the result, error `ERANGE` is raised; kind `LONG_KIND` is guaranteed to be big enough for all values.

```
SUBROUTINE time(itime,errno)
  INTEGER(time_kind),INTENT(OUT) :: itime
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

`ITIME` receives the operating system date/time in seconds since the Epoch.

```
INTEGER(KIND=clock_tick_kind) FUNCTION times(buffer)
  TYPE(tms),INTENT(OUT) :: buffer
```

This function returns the elapsed real time in clock ticks since an arbitrary point in the past, or -1 if the function is unavailable. `BUFFER` is filled in with CPU time information for the calling process and any terminated child processes.

If this function returns zero the values in `BUFFER` will still be correct but the elapsed-time timer was not available.

```
SUBROUTINE ttyname(lunit,s,lens,errno)
  INTEGER(*),INTENT(IN) :: lunit
  CHARACTER(*),OPTIONAL,INTENT(OUT) :: s
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: lens
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Accesses the name of the terminal connected to the logical unit identified by `LUNIT`. If `S` is present, it receives the text of the terminal name (blank-padded or truncated as appropriate, if the length of the terminal name differs from that of `S`). If `LENS` is present, it receives the length of the terminal name. If `ERRNO` is present, it receives the error status.

If `LUNIT` is not a valid logical unit number, or is not connected, error `EBADF` (see `F90_UNIX_ERRNO`) is raised.

```
SUBROUTINE uname(name,errno)
  TYPE(UTSNAME),INTENT(OUT) :: name
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Returns information about the operating system in `NAME`.

55 f90_unix_errno

This module contains part of a Fortran API to functions detailed in ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) [C Language].

The facilities in this module are from Section 2.4 Error Numbers.

55.1 Error Handling

Many procedures provided by the `F90_UNIX_*` modules have an optional `ERRNO` argument, declared as

```
INTEGER(error_kind),OPTIONAL,INTENT(OUT)
```

If this argument is provided it receives the error status from the procedure; zero indicates successful completion, otherwise it will be a non-zero error code, usually one of the ones listed in this module.

If the `ERRNO` argument is omitted and an error condition is raised, the program will be terminated with an informative error message.

If a procedure has no `ERRNO` argument it indicates that no error condition is possible - the procedure always succeeds.

55.2 Parameters

All parameters are of type `INTEGER` with kind `ERROR_KIND`. The following table lists the error message typically associated with each error code; for full details see ISO/IEC 9945-1:1990, either section 2.4 or the appropriate section for the function raising the error.

E2BIG	Arg list too long
EACCES	Permission denied
EAGAIN	Resource temporarily unavailable
EBADF	Bad file descriptor
EBUSY	Resource busy
ECHILD	No child process
EDEADLK	Resource deadlock avoided
EDOM	Domain error
EEXIST	File exists
EFAULT	Bad address
EFBIG	File too large
EINTR	Interrupted function call
EINVAL	Invalid argument
EIO	Input/Output error
EISDIR	Is a directory
EMFILE	Too many open files
EMLINK	Too many links
ENAMETOOLONG	Filename too long
ENFILE	Too many open files in system
ENODEV	No such device
ENOENT	No such file or directory
ENOEXEC	Exec format error
ENOLCK	No locks available
ENOMEM	Not enough space
ENOSPC	No space left on device
ENOSYS	Function not implemented
NOTDIR	Not a directory
NOTEMPTY	Directory not empty
NOTTY	Inappropriate I/O control operation
ENXIO	No such device or address
EPERM	Operation not permitted
EPIPE	Broken pipe
ERANGE	Result too large
EROFS	Read-only file system
ESPIPE	Invalid seek
ESRCH	No such process
EXDEV	Improper link

56 f90_unix_file

This module contains part of a Fortran API to functions detailed in ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) [C Language].

The functions in this module are from Section 5.6 File Characteristics.

Error handling is described in `F90_UNIX_ERRNO`. Note that for procedures with an optional `ERRNO` argument, if an error occurs and `ERRNO` is not present, the program will be terminated.

All the procedures in this module are generic; some may be specific but this is subject to change in a future release.

56.1 Parameters

`INTEGER(int32),PARAMETER :: F_OK`

Flag for requesting file existence check (see `ACCESS`).

`USE F90_UNIX_ENV, ONLY: ID_KIND`

See `F90_UNIX_ENV` for a description of this parameter.

```
USE F90_UNIX_DIR, ONLY: MODE_KIND
```

See `F90_UNIX_DIR` for a description of this parameter.

```
INTEGER(int32), PARAMETER :: R_OK
```

Flag for requesting file readability check (see `ACCESS`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IRGRP
```

File mode bit indicating group read permission (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IROTH
```

File mode bit indicating other read permission (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IRUSR
```

File mode bit indicating user read permission (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IRWXG
```

Mask to select the group accessibility bits from a file mode (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IRWXO
```

Mask to select the other accessibility bits from a file mode (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IRWXU
```

Mask to select the user accessibility bits from a file mode (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_ISGID
```

File mode bit indicating that the file is set-group-ID (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_ISUID
```

File mode bit indicating that the file is set-user-ID (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IWGRP
```

File mode bit indicating group write permission (see `STAT_T`).

```
INTEGER(MODE_KIND), PARAMETER :: S_IWOTH
```

File mode bit indicating other write permission (see `STAT_T`).

```
INTEGER(MODE_KIND),PARAMETER :: S_IWUSR
```

File mode bit indicating user write permission (see `STAT_T`).

```
INTEGER(MODE_KIND),PARAMETER :: S_IXGRP
```

File mode bit indicating group execute permission (see `STAT_T`).

```
INTEGER(MODE_KIND),PARAMETER :: S_IXOTH
```

File mode bit indicating other execute permission (see `STAT_T`).

```
INTEGER(MODE_KIND),PARAMETER :: S_IXUSR
```

File mode bit indicating user execute permission (see `STAT_T`).

```
USE F90_UNIX_ENV,ONLY :: TIME_KIND
```

See `F90_UNIX_ENV` for a description of this parameter.

```
INTEGER(int32),PARAMETER :: W_OK
```

Flag for requesting file writability check (see `ACCESS`).

```
INTEGER(int32),PARAMETER :: X_OK
```

Flag for requesting file executability check (see `ACCESS`).

56.2 Types

```
TYPE stat_t
  INTEGER(MODE_KIND) st_mode
  INTEGER(...) st_ino
  INTEGER(...) st_dev
  INTEGER(...) st_nlink
  INTEGER(id_kind) st_uid
  INTEGER(id_kind) st_gid
  INTEGER(...) st_size
  INTEGER(TIME_KIND) st_atime, st_mtime, st_ctime
END TYPE
```

Derived type holding file characteristics.

ST_MODE File mode (read/write/execute permission for user/group/other, plus set-group-ID and set-user-ID bits).

ST_INO File serial number.

ST_DEV ID for the device on which the file resides.

ST_NLINK The number of links (see `F90_UNIX_DIR`, `LINK` operation) to the file.

ST_UID User number of the file's owner.

ST_GID Group number of the file.

ST_SIZE File size in bytes (regular files only).

ST_ATIME Time of last access.

ST_MTIME Time of last modification.

ST_CTIME Time of last file status change.

```
TYPE UTIMBUF
  INTEGER(time_kind) actime, modtime
END TYPE
```

Data type holding time values for communication to **UTIME**. **ACTIME** is the new value for **ST_ATIME**, **MODTIME** is the new value for **ST_MTIME**.

56.3 Procedures

```
PURE SUBROUTINE access(path,amode,errno)
  CHARACTER(*),INTENT(IN) :: path
  INTEGER(*),INTENT(IN) :: amode
  INTEGER(error_kind),INTENT(OUT) :: errno
```

Checks file accessibility according to the value of **AMODE**; this should be **F_OK** or a combination of **R_OK**, **W_OK** and **X_OK**. In the latter case the values may be combined by addition or the intrinsic function **IOR**.

The result of the accessibility check is returned in **ERRNO**, which receives zero for success (i.e. the file exists for **F_OK**, or all the accesses requested by the **R_OK** et al combination are allowed) or an error code indicating the reason for access rejection. Possible rejection codes include **EACCES**, **ENAMETOOLONG**, **ENOENT**, **ENOTDIR** and **EROFS** (see **F90_UNIX_ERRNO**).

If the value of **AMODE** is invalid, error **EINVAL** is returned.

Note that most **ACCESS** enquiries are equivalent to an **INQUIRE** statement, in particular:

```
CALL ACCESS(PATH,F_OK,ERRNO)
  returns success (ERRNO==0) if and only if
  INQUIRE(FILE=PATH,EXIST=LVAR) would set LVAR to .TRUE.;

CALL ACCESS(PATH,R_OK,ERRNO)
  returns success (ERRNO==0) if and only if
  INQUIRE(FILE=PATH,READ=CHVAR) would set CHVAR to 'YES';

CALL ACCESS(PATH,W_OK,ERRNO)
  returns success (ERRNO==0) if and only if
  INQUIRE(FILE=PATH,WRITE=CHVAR) would set CHVAR to 'YES';

CALL ACCESS(PATH,IOR(W_OK,R_OK),ERRNO)
  returns success (ERRNO==0) if and only if
  INQUIRE(FILE=PATH,READWRITE=CHVAR) would set CHVAR to 'YES'.
```

The only differences being that **ACCESS** returns a reason for rejection, and can test file executability.

```
SUBROUTINE CHMOD(PATH,MODE,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  INTEGER(*),INTENT(IN) :: MODE
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Sets the file mode (**ST_MODE**) to **MODE**.

Possible errors include **EACCES**, **ENAMETOOLONG**, **ENOTDIR**, **EPERM** and **EROFS** (see **F90_UNIX_ERRNO**).

```
SUBROUTINE CHOWN(PATH,OWNER,GROUP,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  INTEGER(id_kind),INTENT(IN) :: OWNER, GROUP
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

Changes the owner (ST_UID) of file PATH to OWNER, and the group number (ST_GID) of the file to GROUP.

Possible errors include EACCES, EINVAL, ENAMETOOLONG, ENOTDIR, ENOENT, EPERM and EROFS (see F90_UNIX_ERRNO).

```
SUBROUTINE FSTAT(LUNIT,BUF,ERRNO)
  INTEGER(*),INTENT(IN) :: LUNIT
  TYPE(stat_t),INTENT(OUT) :: BUF
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

BUF receives the characteristics of the file connected to logical unit LUNIT.

If LUNIT is not a valid logical unit number or is not connected to a file, error EBADF is raised (see F90_UNIX_ERRNO).

```
PURE LOGICAL(word) FUNCTION s_isblk(mode)
  INTEGER(mode_kind),INTENT(IN) :: mode
```

Returns .TRUE. if and only if the MODE value indicates that the file is a “block device”.

```
PURE LOGICAL(word) FUNCTION s_ischr(mode)
  INTEGER(mode_kind),INTENT(IN) :: mode
```

Returns .TRUE. if and only if the MODE value indicates that the file is a “character device”.

```
PURE LOGICAL(word) FUNCTION s_isdir(mode)
  INTEGER(mode_kind),INTENT(IN) :: mode
```

Returns .TRUE. if and only if the MODE value indicates that the file is a directory (or folder).

```
PURE LOGICAL(word) FUNCTION s_isfifo(mode)
  INTEGER(mode_kind),INTENT(IN) :: mode
```

Returns .TRUE. if and only if the MODE value indicates that the file is a “FIFO” (named or unnamed pipe).

```
PURE LOGICAL(word) FUNCTION s_isreg(mode)
  INTEGER(mode_kind),INTENT(IN) :: mode
```

Returns .TRUE. if and only if the MODE value indicates that the file is a “regular” (i.e. normal) file.

```
SUBROUTINE STAT(PATH,BUF,ERRNO)
  CHARACTER(*),INTENT(IN) :: PATH
  TYPE(stat_t),INTENT(OUT) :: BUF
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: ERRNO
```

BUF receives the characteristics of the file PATH.

Possible errors include EACCES, ENAMETOOLONG, ENOENT and ENOTDIR (see F90_UNIX_ERRNO).

```
SUBROUTINE utime(path,times,errno)
  CHARACTER(*),INTENT(IN) :: path
  TYPE(utimbuf),OPTIONAL,INTENT(IN) :: times
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Set the access and modification times of the file named by `PATH` to those specified by the `ACTIME` and `MODTIME` components of `TIMES` respectively.

Possible errors include `EACCES`, `ENAMETOOLONG`, `ENOENT`, `ENOTDIR`, `EPERM` and `EROFS` (see `F90_UNIX_ERRNO`).

57 f90_unix_io

This module will contain part of a Fortran API to functions detailed in ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) [C Language].

In this release only the `FLUSH` procedure is supported. Users are advised to use the `ONLY` clause when using this module, as it will have further names added in later releases.

Error handling is described in `F90_UNIX_ERRNO`. Note that for procedures with an optional `ERRNO` argument, if an error occurs and `ERRNO` is not present, the program will be terminated.

All procedures in this module are generic and may also be specific (this may change in a future release).

57.1 Procedures

```
SUBROUTINE FLUSH(LUNIT,ERRNO)
  INTEGER(*),INTENT(IN) :: LUNIT
  INTEGER(ERROR_KIND),OPTIONAL,INTENT(OUT) :: ERRNO
```

Flushes the output buffer of logical unit `LUNIT`. The effect is similar to the Fortran 2003 `FLUSH` statement, except that the statement allows flushing a unit which is not connected but the procedure does not.

If `LUNIT` is not a valid unit number or is not connected to a file, error `EBADF` is raised (see `F90_UNIX_ERRNO`).

58 f90_unix_proc

This module contains part of a Fortran API to functions detailed in ISO/IEC 9945-1:1990 Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) [C Language].

The functions in this module are from Section 3: Process Primitives, excluding 3.3 Signals. The C language functions `abort`, `atexit`, `exit` and `system` are also provided by this module.

Error handling is described in `F90_UNIX_ERRNO`. Note that for procedures with an optional `ERRNO` argument, if an error occurs and `ERRNO` is not present, the program will be terminated.

All the procedures in this module are generic; some may be specific but this is subject to change in a future release.

58.1 Parameters

```
INTEGER(int32),PARAMETER :: atomic_int
```

Integer kind for “atomic” operations in signal handlers (see `ALARM`).

```
INTEGER(int32),PARAMETER :: atomic_log
```


Logical kind for “atomic” operations in signal handlers (see **ALARM**).

```
USE f90_unix_env, ONLY: pid_kind=>id_kind
```

Integer kind for representing process IDs; this has been superseded by **ID_KIND** from **F90_UNIX_ENV**.

```
USE f90_unix_env, ONLY: time_kind
```

Integer kind for representing times in seconds.

```
INTEGER(int32),PARAMETER :: wnohang
```

Option bit for **WAITPID** indicating that the calling process should not wait for the child process to stop or exit.

```
INTEGER(int32),PARAMETER :: wuntraced
```

Option bit for **WAITPID** indicating that status should be returned for stopped processes as well as terminated ones.

58.2 Procedures

```
SUBROUTINE abort(message)
  CHARACTER(*),OPTIONAL :: message
```

ABORT cleans up the i/o buffers and then terminates execution, producing a core dump on Unix systems. If **MESSAGE** is given it is written to logical unit 0 (zero) preceded by ‘**abort:**’.

```
SUBROUTINE alarm(seconds,subroutine,seclft,errno)
  INTEGER(*),INTENT(IN) :: seconds
  INTERFACE
    SUBROUTINE subroutine()
  END
  END INTERFACE
  OPTIONAL subroutine
  INTEGER(time_kind),OPTIONAL,INTENT(OUT) :: seclft
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Establishes an “alarm” call to the procedure **SUBROUTINE** to occur after **SECONDS** seconds have passed, or cancels an existing alarm if **SECONDS**==0. If **SUBROUTINE** is not present, any previous association of a subroutine with the alarm signal is left unchanged. If **SECLFT** is present, it receives the number of seconds that were left on the preceding alarm or zero if there were no existing alarm.

The subroutine invoked by the alarm call is only permitted to define **VOLATILE SAVED** variables that have the type **INTEGER(atomic_int)** or **LOGICAL(atomic_log)**; defining or referencing any other kind of variable may result in unpredictable behaviour, even program termination. Furthermore, it shall not perform any array or **CHARACTER** operations, input/output, or invoke any intrinsic function or module procedure.

If an alarm call is established with no handler (i.e. **SUBROUTINE** was not present on the first call) the process may be terminated when the alarm goes off.

Possible errors include **ENOSYS**.

```
SUBROUTINE atexit(subroutine,errno)
  INTERFACE
    SUBROUTINE subroutine()
  END
  END INTERFACE
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Registers an argumentless subroutine for execution on normal termination of the program.

If the program terminates normally, all subroutines registered with `ATEXIT` will be invoked in reverse order of their registration. Normal termination includes using the `F90_UNIX_PROC` procedure `EXIT`, executing a Fortran `STOP` statement or executing a main program `END` statement. `ATEXIT` subroutines are invoked before automatic file closure.

If the program terminates due to an error or by using the `F90_UNIX_PROC` procedure `FASTEXIT`, these subroutines will not be invoked.

Possible errors include `ENOMEM`.

Note: The list of `ATEXIT` procedures registered via Fortran is separate from those registered via C; the latter will be invoked after all files have been closed.

```
SUBROUTINE execl(path,arg0...,errno)
  CHARACTER(*),INTENT(IN) :: path
  CHARACTER(*),INTENT(IN) :: arg0...
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Executes a file instead of the current image, like `EXECV`. The arguments to the new program are specified by the dummy arguments which are named `ARG0`, `ARG1`, etc. up to `ARG20` (additional arguments may be provided in later releases). Note that these are not optional arguments, any actual argument that is itself an optional dummy argument must be present. This function is the same as `EXECV` except that the arguments are provided individually instead of via an array; and because they are provided individually, there is no need to provide the lengths (the lengths being taken from each argument itself).

Errors are the same as for `EXECV`.

```
SUBROUTINE execlp(file,arguments,,errno)
  CHARACTER(*),INTENT(IN) :: file
  CHARACTER(*),INTENT(IN) :: arguments
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Executes a file instead of the current image, like `EXECV`. The arguments to the new program are specified by the dummy arguments which are named `ARG0`, `ARG1`, etc. up to `ARG20` (additional arguments may be provided in later releases). Note that these are not optional arguments, any actual argument that is itself an optional dummy argument must be present. This function is the same as `EXECL` except that determination of the program to be executed follows the same rules as `EXECVP`.

Errors are the same as for `EXECV`.

```
SUBROUTINE execv(path,argv,lenargv,errno)
  CHARACTER(*),INTENT(IN) :: path
  CHARACTER(*),INTENT(IN) :: argv(:)
  INTEGER(*),INTENT(IN) :: lenargv(:)
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Executes the file `PATH` in place of the current process image; for full details see ISO/IEC 9945-1:1990 section 3.1.2. `ARGV` is the array of argument strings, `LENARGV` containing the desired length of each argument. If `ARGV` is not zero-sized, `ARGV(1)(:LENARGV(1))` is passed as argument zero (i.e. the program name).

If `LENARGV` is not the same shape as `ARGV`, error `EINVAL` is raised (see `F90_UNIX_ERRNO`). Other possible errors include `E2BIG`, `EACCES`, `ENAMETOOLONG`, `ENOENT`, `ENOTDIR`, `ENOEXEC` and `ENOMEM`.

```
SUBROUTINE EXECVE(path,argv,lenargv,env,lenenv,errno)
  CHARACTER(*),INTENT(IN) :: path
  CHARACTER(*),INTENT(IN) :: argv(:)
  INTEGER(*),INTENT(IN) :: lenargv(:)
  CHARACTER(*),INTENT(IN) :: env(:)
  INTEGER(*),INTENT(IN) :: lenenv(:)
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Similar to `EXECV`, with the environment strings specified by `ENV` and `LENENV` being passed to the new program; for full details see ISO/IEC 9945-1:1990 section 3.1.2.

If `LENARGV` is not the same shape as `ARGV` or `LENENV` is not the same shape as `ENV`, error `EINVAL` is raised (see `F90_UNIX_ERRNO`). Other errors are the same as for `EXECV`.

```
SUBROUTINE execvp(file,argv,lenargv,errno)
  CHARACTER(*),INTENT(IN) :: file
  CHARACTER(*),INTENT(IN) :: argv(:)
  INTEGER(*),INTENT(IN) :: lenargv(:)
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

The same as `EXECV` except that the program to be executed, `FILE`, is searched for using the `PATH` environment variable (unless it contains a slash character, in which case `EXECVP` is identical in effect to `EXECV`).

Errors are the same as for `EXECV`.

```
SUBROUTINE exit(status)
  INTEGER(int32),OPTIONAL,INTENT(IN) :: status
```

Terminate execution as if executing the `END` statement of the main program (or an unadorned `STOP` statement). If `STATUS` is given it is returned to the operating system (where applicable) as the execution status code.

```
SUBROUTINE fastexit(status)
  INTEGER,OPTIONAL,INTENT(IN) :: status
```

This provides the functionality of ISO/IEC 9945-1:1990 function `_exit` (section 3.2.2). There are two main differences between `FASTEXIT` and `EXIT`:

1. When `EXIT` is called all open logical units are closed (as if by a `CLOSE` statement). With `FASTEXIT` this is not done, nor are any file buffers flushed, thus the contents and status of any file connected at the time of calling `FASTEXIT` are undefined.
2. Subroutines registered with `ATEXIT` are not executed.

```
SUBROUTINE fork(pid,errno)
  INTEGER(id_kind),INTENT(OUT) :: pid
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Creates a new process which is an exact copy of the calling process. In the new process, the value returned in `PID` is zero; in the calling process the value returned in `PID` is the process ID of the new (child) process.

Possible errors include `EAGAIN`, `ENOMEM` and `ENOSYS` (see `F90_UNIX_ERRNO`).

```
SUBROUTINE pause(errno)
  INTEGER(error_kind),INTENT(OUT) :: errno
```

Suspends process execution until a signal is raised. If the action of the signal was to terminate the process, the process is terminated without returning from `PAUSE`. If the action of the signal was to invoke a signal handler (e.g. via `ALARM`), process execution continues after return from the signal handler.

If process execution is continued after a signal, `ERRNO` is set to `EINTR`.

If this functionality is not available, `ERRNO` is set to `ENOSYS`.

```
PURE SUBROUTINE sleep(seconds,seclft)
  INTEGER(*),INTENT(IN) :: seconds
  INTEGER(time_kind),OPTIONAL,INTENT(OUT) :: seclft
```

Suspends process execution for **SECONDS** seconds, or until a signal has been delivered. If **SECLEFT** is present, it receives the number of seconds remaining in the sleep time (zero unless the sleep was interrupted by a signal).

```
SUBROUTINE system(string,status,errno)
  CHARACTER(*),INTENT(IN) :: string
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: status,errno
```

Passes **STRING** to the command processor for execution. If **STATUS** is present it receives the completion status - this is the same status returned by **WAIT** and can be decoded with **WIFEXITED** etc. If **ERRNO** is present it receives the error status from the **SYSTEM** call itself.

Possible errors are those from **FORK** or **EXECV**.

```
SUBROUTINE wait(status,retpid,errno)
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: status
  INTEGER(id_kind),OPTIONAL,INTENT(OUT) :: retpid
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Wait for any child process to terminate (returns immediately if one has already terminated). See ISO/IEC 9945-1:1990 section 3.2.1 for full details.

If **STATUS** is present it receives the termination status of the child process. If **RETPID** is present it receives the process number of the child process.

Possible errors include **ECHILD** and **EINTR** (see **F90_UNIX_ERRNO**).

```
SUBROUTINE waitpid(pid,status,options,retpid,errno)
  INTEGER(id_kind),INTENT(IN) :: pid
  INTEGER(int32),OPTIONAL,INTENT(OUT) :: status
  INTEGER(int32),OPTIONAL,INTENT(IN) :: options
  INTEGER(id_kind),OPTIONAL,INTENT(OUT) :: retpid
  INTEGER(error_kind),OPTIONAL,INTENT(OUT) :: errno
```

Wait for a particular child process to terminate (or for any one if **PID==(-1)**). If **OPTIONS** is not present it is as if it were present with a value of 0. See ISO/IEC 9945-1:1990 section 3.2.1 for full details.

Possible errors include **ECHILD**, **EINTR** and **EINVAL** (see **F90_UNIX_ERRNO**).

```
PURE INTEGER(int32) FUNCTION wexitstatus(stat_val)
  INTEGER(int32),INTENT(IN) :: stat_val
```

If **WIFEXITED(STAT_VAL)** is **.TRUE.**, this function returns the low-order 8 bits of the status value supplied to **EXIT** or **FASTEXIT** by the child process. If the child process executed a **STOP** statement or main program **END** statement, the value will be zero. If **WIFEXITED(STAT_VAL)** is **.FALSE.**, the function value is undefined.

```
PURE LOGICAL(word) FUNCTION wifexited(stat_val)
  INTEGER(error_kind),INTENT(IN) :: stat_val
```

Returns **.TRUE.** if and only if the child process terminated by calling **FASTEXIT**, **EXIT**, by executing a **STOP** statement or main program **END** statement, or possibly by some means other than Fortran.

```
PURE LOGICAL(word) FUNCTION wifsignaled(stat_val)
  INTEGER(int32),INTENT(IN) :: stat_val
```

Returns **.TRUE.** if and only if the child process terminated by receiving a signal that was not caught.

```
PURE LOGICAL(word) FUNCTION wifstopped(stat_val)
    INTEGER(int32), INTENT(IN) :: stat_val
```

Returns `.TRUE.` if and only if the child process is stopped (and not terminated). Note that `WAITPID` must have been used with the `WUNTRACED` option to receive such a status value.

```
PURE INTEGER(int32) FUNCTION wstopsig(stat_val)
    INTEGER(int32), INTENT(IN) :: stat_val
```

If `WIFSTOPPED(STAT_VAL)` is `.TRUE.`, this function returns the signal number that caused the child process to stop. If `WIFSTOPPED(STAT_VAL)` is `.FALSE.`, the function value is undefined.

```
PURE INTEGER(int32) FUNCTION wtermsig(stat_val)
    INTEGER(int32), INTENT(IN) :: stat_val
```

If `WIFSIGNALED(STAT_VAL)` is `.TRUE.`, this function returns the signal number that caused the child process to terminate. If `WIFSIGNALED(STAT_VAL)` is `.FALSE.`, the function value is undefined.

High Performance Fortran Library Module

59 HPF_LIBRARY Module Summary

This section contains a quick reference guide to the `HPF_LIBRARY` module provided by the High Performance Fortran extension to Fortran 95. To access these functions, the `-hpf` option is needed in addition to ‘`USE HPF_LIBRARY`’ in the program unit.

In the following tables, arguments with names in italics are optional. For full information on these functions refer to the HPF Language Specification.

59.1 Elemental Functions

Function (Arguments)	Description
LEADZ (I)	Number of leading zero bits in the internal representation.
POPCNT (I)	Number of one bits in the internal representation.
POPPAR (I)	1 if POPCNT(I) is odd, 0 otherwise.

59.2 Reduction Functions

These are analogous to the `SUM`, `PRODUCT` et al intrinsic functions.

Function (Arguments)	Reduction Operation
IALL (ARRAY, DIM, MASK)	IAND intrinsic function
IANY (ARRAY, DIM, MASK)	IOR intrinsic function
IPARITY (ARRAY, DIM, MASK)	IEOR intrinsic function
PARITY (MASK, DIM)	.NEQV. operator

59.3 Scan Functions

Each scan function produces an array of the same shape as its first argument (the `MASK` argument for **LOGICAL** functions, the `ARRAY` argument for the others). Scanning combines the “current” value with each array element of this argument according to the scan function; e.g. `SUM_PREFIX` and `SUM_SUFFIX` add each array element. There are two scan functions for each operation; one forwards (`XXX_PREFIX`) and one backwards (`XXX_SUFFIX`).

Note that by definition, `XXX_SUFFIX(A(1:N))` is equal to `XXX_PREFIX(A(N:1:-1))`. Note also that the initial “current” value depends on the function as follows:

Function	Initial Value
ALL	.TRUE.
ANY and PARITY	.FALSE.
COPY	the value of the first element of the array
COUNT, IANY, IPARITY and SUM	zero
IALL	NOT(0)
MAXVAL	-HUGE(ARRAY)
MINVAL	HUGE(ARRAY)
PRODUCT	1

E.g., ‘`SUM_PREFIX((/ 1,2,3,4 /))`’ is equal to ‘`(/ 1,3,6,10 /)`’.

If the `DIM` argument is not present, the whole array is scanned in array element order. With a `DIM` argument, the array is scanned in that dimension in parallel.

If the `SEGMENT` argument is present, the “current” value is reset each time the element of `SEGMENT` corresponding to the scan changes value (from the last element of `SEGMENT` considered).

If the `MASK` optional argument is present (for non-**LOGICAL**-valued functions), only those elements of `ARRAY` for which the corresponding element of `MASK` is `.TRUE.` contribute to the value of the scan.

If the **EXCLUSIVE** argument is present and **.TRUE.**, the scan combining operation is performed after deciding the value for each element, thus shifting the value of the scan function one place and inserting, at the beginning, the initial current value. E.g., ‘**SUM_PREFIX**((/ 1,2,3 /), **EXCLUSIVE**=**.TRUE.**)’ is equal to ‘(/ 0,1,3 /)’.

Function (Arguments)	Combining Operation
ALL_PREFIX (<i>MASK</i> , <i>DIM</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	.AND. operator
ANY_PREFIX (<i>MASK</i> , <i>DIM</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	.OR. operator
COPY_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>SEGMENT</i>)	Copy
COUNT_PREFIX (<i>MASK</i> , <i>DIM</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	Count .TRUE. values
IALL_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	IAND intrinsic function
IANY_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	IOR intrinsic function
IPARITY_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	IEOR intrinsic function
MAXVAL_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	Maximum value
MINVAL_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	Minimum value
PARITY_PREFIX (<i>MASK</i> , <i>DIM</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	.NEQV. operator
PRODUCT_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	Multiplication
SUM_PREFIX (<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i> , <i>SEGMENT</i> , <i>EXCLUSIVE</i>)	Addition

59.4 Scatter Functions

Each scatter function produces an array of the same shape as the **BASE** argument. The default value of each element of the result is the same as that of the corresponding element of **BASE**.

The elements of **ARRAY** (or **MASK**, for **LOGICAL** functions) are scattered to selected elements of the result, where they are combined using the specified operation (see the relevant scan or reduction functions).

Each element is scattered by taking its index values and using the values from the corresponding **INDX_n** arguments to index into the result. Note that the number of **INDX_n** arguments must be the same as the rank of **ARRAY** (or **MASK**). The shape of **ARRAY** (or **MASK**) need not be the same as **BASE**.

Function	(Arguments)
ALL_SCATTER	(<i>MASK</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i>)
ANY_SCATTER	(<i>MASK</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i>)
COPY_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)
COUNT_SCATTER	(<i>MASK</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i>)
IALL_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)
IANY_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)
IPARITY_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)
MAXVAL_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)
MINVAL_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)
PARITY_SCATTER	(<i>MASK</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i>)
PRODUCT_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)
SUM_SCATTER	(<i>ARRAY</i> , <i>BASE</i> , <i>INDX1</i> , ..., <i>INDX_n</i> , <i>MASK</i>)

If the optional **MASK** is present (only for non-**LOGICAL** functions), only those elements of **ARRAY** for which the corresponding elements of **MASK** are **.TRUE.** are so scattered.

59.5 Sorting Functions

These functions produce permutation indices which can be used to index an array in a sorted sequence.

Function (Arguments)	Description
GRADE_DOWN (<i>ARRAY</i> , <i>DIM</i>)	Descending sort as array indices
GRADE_UP (<i>ARRAY</i> , <i>DIM</i>)	Ascending sort as array indices

59.6 Enquiry Functions

These functions enquire about the actual alignment, distribution and template values used at runtime in a multi-processor environment. Since the NAG Fortran Compiler is a single-processor implementation of HPF, these functions do not return very useful information.

Function	(Arguments)
HPF_ALIGNMENT	(ALIGNEE, <i>LB</i> , <i>UB</i> , <i>STRIDE</i> , <i>AXIS_MAP</i> , <i>IDENTITY_MAP</i> , <i>DYNAMIC</i> , <i>NCOPIES</i>)
HPF_DISTRIBUTION	(DISTRIBUTE, <i>AXIS_TYPE</i> , <i>AXIS_INFO</i> , <i>PROCESSORS_RANK</i> , <i>PROCESSORS_SHAPE</i>)
HPF_TEMPLATE	(ALIGNEE, <i>TEMPLATE_RANK</i> , <i>LB</i> , <i>UB</i> , <i>AXIS_TYPE</i> , <i>AXIS_INFO</i> , <i>NUMBER_ALIGNED</i> , <i>DYNAMIC</i>)

Standard Fortran 95

60 Fortran 95 Program Structure

This section contains a quick reference guide to the structure of a Fortran 95 program.

Expressions are described in the next section.

- Syntax element definition is represented by ‘::<=’.
- Names in italics (e.g., ‘*expr*’) are syntax element names.
- Items in ‘[]’ are optional.
- Ellipsis ‘...’ indicates optional repetition of the preceding item.
- ‘{’ and ‘}’ are for grouping only.
- Items separated by ‘|’ are a “one-of-many” selection.
- Statements end in ‘-*stmt*’ and are described in the section “Fortran 95 Statements”.

INCLUDE Line

The INCLUDE line has the form:

INCLUDE *char-literal-constant*

The *char-literal-constant* must not have a *kind-param* and must be the pathname of an accessible file; relative pathnames will be searched for in the current working directory and all directories named by *-I* options. The INCLUDE line is effectively replaced by the contents of the named file.

Program Unit Structure

program-unit ::= *main-program* | *module* | *procedure* | *block-data-subprogram*

main-program ::= [*program-stmt*]
body
end-program-stmt

body ::= *declaration-section*
[*executable-section*]
[*contains-section*]

Note: A *contains-section* cannot appear in an internal procedure (an internal procedure is a procedure in the *contains-section* of another procedure).

contains-section ::= *contains-stmt*
procedure...

module ::= *module-stmt*
[*declaration-section*]
[*contains-section*]
end-module-stmt

procedure ::= *function-subprogram* | *subroutine-subprogram*

function-subprogram ::= *function-stmt*
 body
 end-function-stmt

subroutine-subprogram ::= *subroutine-stmt*
 body
 end-subroutine-stmt

block-data-subprogram ::= *block-data-stmt*
 declaration-section
 end-block-data-stmt

Declaration Section

declaration-section ::= [*use-stmt...*]
 [*implicit-part...*]
 [*declaration...*]

implicit-part ::= *implicit-stmt* | *parameter-stmt* | *format-stmt* | *entry-stmt*

declaration ::= *derived-type-definition* | *interface-block* | *declarative* |
 format-stmt | *entry-stmt*

declarative ::= *allocatable-stmt* | *common-stmt* | *data-stmt* | *dimension-stmt* |
 equivalence-stmt | *external-stmt* | *intent-stmt* | *intrinsic-stmt* | *namelist-stmt* |
 optional-stmt | *parameter-stmt* | *pointer-stmt* | *private-stmt* | *public-stmt* |
 save-stmt | *statement-function-stmt* | *target-stmt* | *type-declaration-stmt*

derived-type-definition ::= *type-stmt*
 component-def-stmt...
 end-type-stmt

interface-block ::= *interface-stmt*
 { *interface-body* | *module-procedure-stmt* } ...
 end-interface-stmt

interface-body ::= *function-stmt declaration-section end-function-stmt* |
 subroutine-stmt declaration-section end-subroutine-stmt

Executable Section

executable-section ::= [*executable* | *executable-construct* | *data-stmt*]...

executable ::= *allocate-stmt* | *arithmetic-if-stmt* | *assignment-stmt* | *backspace-stmt* | *call-stmt* |
 close-stmt | *computed-goto-stmt* | *continue-stmt* | *cycle-stmt* | *deallocate-stmt*
 | *endfile-stmt* | *exit-stmt* | *forall-stmt* | *goto-stmt* | *if-stmt* | *inquire-stmt* |
 nullify-stmt | *open-stmt* | *pause-stmt* | *pointer-assignment-stmt* | *print-stmt* |
 read-stmt | *return-stmt* | *rewind-stmt* | *stop-stmt* | *where-stmt* | *write-stmt*

executable-construct ::= *select-case-construct* | *do-loop* | *forall-construct* | *if-construct* |
 where-construct

select-case-construct ::= *select-stmt*
 [*case-stmt executable-section*] ...
 end-select-stmt

if-construct ::= *if-then-stmt*
 [*executable-section*]
 [*elseif-stmt executable-section*] ...
 [*else-stmt executable-section*]
 endif-stmt

```
where-construct ::= where-construct-stmt  
                   where-body  
                   [ elsewhere-stmt where-body ]  
                   endwhere-stmt  
where-body ::= [ where-assignment-stmt | where-stmt | where-construct |  
                  elsewhere-mask-stmt ] ...  
forall-construct ::= forall-construct-stmt  
                     forall-body  
                     endforall-stmt  
forall-body ::= [ forall-assignment-stmt | where-stmt | where-construct | forall-stmt |  
                   forall-construct ] ...  
do-loop ::= do-stmt  
            executable-section  
            do-ending  
do-ending ::= enddo-stmt | executable
```

Note: If the *do-stmt* specifies a terminating label, the *do-ending* is the statement with that label and shall not be a *goto-stmt*, *return-stmt*, *stop-stmt*, *exit-stmt*, *cycle-stmt*, *arithmetic-if-stmt* or any form of **END** statement. If the *do-stmt* does not specify a terminating label, the *do-ending* shall be an *enddo-stmt*.

61 Fortran 95 Expressions

This section contains a quick reference guide to expression syntax and semantics in Fortran 95. For more details see Metcalf, Reid and Cohen “Fortran 95/2003 Explained” or the Fortran 95 standard “IS 1539-1:1997”.

- Syntax element definition is represented by ‘::=’.
- Names in italics (e.g., ‘*expr*’) are syntax element names.
- Items in ‘[]’ are optional.
- Ellipsis ‘...’ indicates optional repetition of the preceding item.
- ‘{’ and ‘}’ are for grouping only.
- Items separated by ‘|’ are a “one-of-many” selection.

General Form

The general form of an expression *expr* is as follows:

expr ::= [*unary-operator*] *operand* [*binary-operator operand*]...

operand ::= *literal-constant* | *object* | *array-constructor* |
structure-constructor | *function-reference* | (*expr*)

binary-operator ::= + | - | * | / | ** | == | /= | < | > | <= | >= | // |
.AND. | .OR. | .EQV. | .NEQV. | *user-operator*

unary-operator ::= + | - | .NOT. | *user-operator*

user-operator ::= .*letter*[*letter*...].

The operators ==, /=, <, >, <= and >= may also be represented by .EQ., .NE., .LT., .GT., .LE. and .GE. respectively.

Operators

Binary operators are associated left-to-right except for exponentiation, which associates right-to-left. The precedence of each operator is shown by the following table, from highest to lowest:

Operators	Binary/Unary
<i>user-operator</i>	Unary
**	Binary
* /	Binary
+ -	Unary
+ -	Binary
//	Binary
== /= < > <= >=	Binary
.NOT.	Unary
.AND.	Binary
.OR.	Binary
.EQV. .NEQV.	Binary
<i>user-operator</i>	Binary

Literal Constants

literal-constant ::= *integer-literal* | *real-literal* | *double-precision-literal* |
logical-literal | *complex-literal* | *character-literal*

integer-literal ::= *digit-string* [*_ kind-specifier*]

digit-string ::= { 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 }...

kind-specifier ::= *digit-string* | *name*

Note: *name* must be an INTEGER PARAMETER name.

Examples:

```
42
17_1
32767_int16
```

real-literal ::= *mantissa* [*E exponent*] [*- kind-specifier*] |
digit-string *E exponent* [*- kind-specifier*]

mantissa ::= *digit-string* . [*digit-string*] |
. *digit-string*

exponent ::= [+ | -] *digit-string*

Examples:

```
4.2
.7_1
327E+67_real64
```

double-precision-literal ::= { *mantissa* | *digit-string* } *D exponent*

Examples:

```
4d2
1.7D-10
```

logical-literal ::= { .TRUE. | .FALSE. } [*- kind-specifier*]

complex-literal ::= (*int-or-real-literal* , *int-or-real-literal*)

int-or-real-literal ::= *integer-literal* | *real-literal*

Examples:

```
(42,17_1)
(0,1d0)
```

character-literal ::= [*kind-specifier* -] { ' [*char...*] ' | " [*char...*] " }

Note: *char* is any character other than the quoting character, or the quoting character occurring twice (e.g., ''' is the same as '"').

Examples:

```
'42'
"Say 'Hello'."
```

Constants

constant ::= *literal-constant* | *name*

Named constants are declared by the PARAMETER statement or by a type-declaration statement with the PARAMETER attribute

Objects

```
object ::= object-ref [ % object-ref ]... [ ( substring-range ) ]
object-ref ::= name [ section-subscript-list ]
section-subscript-list ::= ( section-subscript [ , section-subscript ]... )
section-subscript ::= expr | triplet
triplet ::= [ expr ] : [ expr ] [ : expr ]
substring-range ::= [ expr ] : [ expr ]

name ::= letter [ letter | digit | _ ] ...

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

variable ::= object

constant-subobject ::= object | character-literal ( substring-range )
```

Note: For a *variable*, the name in the initial *object-ref* must be that of a variable.

For a *constant-subobject*, the name in the initial *object-ref* must be that of a constant (i.e., a **PARAMETER**).

An *object* is scalar if and only if all its *object-ref*'s are scalar. An *object-ref* with non-zero rank (i.e., an array) can only be followed by scalar non-pointer *object-refs*.

An **array element** is a scalar object whose final *object-ref* contains a *section-subscript-list* where each *section-subscript* is a scalar expression. For example

```
ARRAY(I,J)
SCALAR%ARRAY(I,J)
```

A **structure component** is an object with more than one *object-ref* whose final *object-ref* has no *section-subscript-list* and all preceding *object-refs* are scalar. For example

```
SCALAR%SCALAR
ARRAY(I,J)%SCALAR
```

A **substring** is a scalar object with a *substring-range*. For example

```
SCALAR(I:J)
ARRAY(I,J)(K:L)
```

A **whole array** is an array object with only one *object-ref*.

An **array section** is an array object where a non-final *object-ref* has non-zero rank, or the final *object-ref* has a *section-subscript-list* and one or more of the *section-subscripts* has non-zero rank or is a *triplet*. For example

```
ARRAY(2:M)
ARRAY%SCALAR
```

Array Constructor

```
array-constructor ::= (/ ac-item [ , ac-item ] ... /)
ac-item ::= expr | ac-implied-do
ac-implied-do ::= ( expr [ , expr ] ... , name = expr , expr [ , expr ] )
```

Note: *name* is an integer variable name, but its value is not affected by the execution of the array constructor. All expressions in an array constructor must have the same type and type parameters; in particular, for characters, they must have the same length.

Structure Constructor

```
structure-constructor ::= name ( expr [ , expr ] ... )
```

Note: *name* must be the name of a derived type. Each expression must be assignment-compatible with the corresponding component of the derived type.

Function Reference

function-reference ::= *name* ([*expr* [, *expr*] ...])

Note: *name* must be the name of a function.

62 Fortran 95 Statements

This section contains a quick reference guide to the statements in Fortran 95.

- Syntax element definition is represented by ‘::=’.
- Names in italics (e.g., ‘*expr*’) are syntax element names.
- Items in ‘[]’ are optional.
- Ellipsis ‘...’ indicates optional repetition of the preceding item.
- ‘{’ and ‘}’ are for grouping only.
- Items separated by ‘|’ are a “one-of-many” selection.

Type Specification

The following syntax is used by several declaration statements.

```
type-spec ::= numeric-type [ kind-specifier ] |  
              numeric-type * digit-string |  
              DOUBLE PRECISION |  
              CHARACTER [ char-specifier ]  
  
numeric-type ::= COMPLEX | INTEGER | LOGICAL | REAL  
kind-specifier ::= ( [ KIND = ] expr )  
  
char-specifier ::= * digit-string |  
                  * ( char-length ) |  
                  ( char-length [ , [ KIND = ] expr ] ) |  
                  ( LEN = char-length [ , KIND = expr ] ) |  
                  ( KIND = expr [ , LEN = char-length ] )
```

```
char-length ::= * | expr
```

Labels and Construct Names

```
label ::= digit-string
```

Note that the *digit-string* in a label must contain at most 5 digits, and at least one of them must be non-zero. Leading zeroes are not significant, but do count towards the limit of 5.

Although it is not shown in the syntax definitions, all statements may be labelled and the **FORMAT** statement must be labelled.

```
construct-name ::= name
```

Construct names are “class 1” names, and must not be the same as any other class 1 name in a subprogram; class 1 names includes variables, procedures, program unit names, et cetera.

Entity Declaration List

The following syntax is used by both Component Declaration and Type Declaration statements.

```
entity-decl-list ::= entity-decl [ , entity-decl ] ...  
entity-decl ::= name [ * char-length ] [ array-spec ] [ initial-value ]  
  
array-spec ::= explicit-shape | assumed-shape | deferred-shape | assumed-size  
  
explicit-shape ::= ( explicit-bound [ , explicit-bound ] ... )  
explicit-bound ::= [ expr : ] expr
```


assumed-shape ::= (*assumed-bound* [, *assumed-bound*] ...)
assumed-bound ::= [*expr*] :

deferred-shape ::= (*deferred-bound* [, *deferred-bound*] ...)
deferred-bound ::= :

assumed-size ::= ([*explicit-bound* ,]... *assumed-size-bound*)
assumed-size-bound ::= [*expr* :] *

initial-value ::= = *expression* | => NULL()

Allocatable Statement

ALLOCATABLE [::] *name* [*deferred-shape*] [, *name* [*deferred-shape*]] ...

Declares the listed entities to be allocatable arrays. If array bounds are present, they must be deferred.

Allocate Statement

ALLOCATE (*allocate-item* [, *allocate-item*] ... [, STAT = *variable*])
allocate-item ::= *variable* [*explicit-shape*]

Allocates a pointer or allocatable array. If the allocation fails and the STAT= clause is present, the STAT= variable will be assigned a non-zero value.

Arithmetic If Statement (*obsolescent*)

IF (*expr*) *label* , *label* , *label*

Branches to one of three labels depending on whether *expr* is negative, zero or positive respectively. The expression must be scalar and of type integer or real.

Assignment Statement

variable = *expr*

The expression is evaluated and assigned to the variable. For intrinsic assignment, it must be assignment-compatible with the variable, that is:

1. If the variable is scalar, the expression must be scalar. If the variable is an array, the expression may be scalar or an array of the same shape.
2. If the variable is of type LOGICAL, the expression must be LOGICAL but may be any KIND.
3. If the variable is of type CHARACTER, the expression must be of type CHARACTER and of the same KIND.
4. If the variable is of type INTEGER, REAL or COMPLEX, the expression must be of type INTEGER, REAL or COMPLEX.
5. If the variable is of derived type, the expression must be of that type.

For defined assignment, there must be a generic interface for ASSIGNMENT(=) which matches the types, kinds and ranks of the variable and expression; the user-specified assignment routine is called. Note that in the case of derived types, defined assignment is permitted to override the intrinsic assignment.

Backspace Statement

BACKSPACE *expr*

BACKSPACE (*position-spec-list*)

position-spec-list ::= *position-spec* [, *position-spec*]

position-spec ::= { [UNIT=] *expr* } | { IOSTAT= *variable* } | { ERR= *label* }

Note: A *position-spec-list* is required to have a `UNIT= position-spec`; the `UNIT=` keyword and equals sign may be omitted only if it is the first in the list.

Positions the file connected to the specified unit to the record preceding the current one. An error condition is raised if the file is not connected, or the unit does not support backspacing.

The effect of each *position-spec* is as below:

`UNIT=` Specifies the i/o unit; it must be a scalar integer expression identifying an external file.

`ERR=` Transfers control to the specified label if an error condition occurs during the i/o statement.

`IOSTAT=` Sets the *variable* to a positive number if an error occurs and to zero otherwise.

Block Data Statement

`BLOCK DATA [name]`

This is the first statement of a block data subprogram. All but one block data subprogram must be named.

Call Statement

`CALL name [([actual-arg-list])]`

actual-arg-list ::= *actual-arg* [, *actual-arg*]...

actual-arg ::= *expr* | * *label*

Calls the named subroutine.

Case Statement

`CASE DEFAULT [construct-name]`

`CASE (case-value-range [, case-value-range]...) [construct-name]`

case-value-range ::= *expr* [: [*expr*]] |
: *expr*

Marks the beginning of a `CASE` part (and the end of any preceding `CASE` part). Statements in this part are executed if the corresponding `SELECT` expression value satisfies the appropriate `CASE` condition:

`.EQ.` A *case-value-range* of the form (*expr*).

`.LE.` A *case-value-range* of the form (:*expr*).

`.GE.` A *case-value-range* of the form (*expr*:).

range A *case-value-range* of the form (*expr*:*expr*) is satisfied for values greater than or equal to the first expression, and less than or equal to the second expression.

DEFAULT The `CASE DEFAULT` clause is selected if the value does not satisfy any other `CASE` statements in that `SELECT` construct.

Note that within a `SELECT` construct, each `CASE` statement must have distinct conditions so that only one can be satisfied.

Close Statement

`CLOSE expr`

`CLOSE (position-spec-list)`

(See the `BACKSPACE` statement for the *position-spec-list* definition.)

Closes the specified unit.

Common Statement

```
COMMON [ / [ common-block-name ] / ] common-object-list  
       [ [ , ] / [ common-block-name ] / common-object-list ]...
```

common-object-list ::= *common-object* [, *common-object*]...

common-object ::= *name* [*array-spec*]

Declares a common block. If no *common-block-name* is specified, “blank common” is the common block declared. Multiple **COMMON** statements for the same common block act as if the *common-object-lists* were concatenated in a single statement.

Component Definition Statement

```
type-spec [ [ , component-attribute-list ] :: ] entity-decl-list
```

component-attribute-list ::= *component-attr* [, *component-attr*]...

component-attr ::= **DIMENSION** *array-spec* |
 POINTER

Declares one or more components of a derived type. Any *array-spec* in a component definition must be *deferred-shape* if the **POINTER** attribute is present, and must be *explicit-shape* otherwise. Any *initial-value* that is present defines the default value for that component of any new entities of the type.

Computed Goto Statement (*obsolescent*)

```
GOTO ( label [ , label ] ... ) expr
```

The (integer scalar) expression is evaluated; if it is less than one or greater than the number of labels in the list, control is transferred to the next statement. Otherwise control is transferred to the corresponding label.

Contains Statement

```
CONTAINS
```

This statement separates the declarations of a module from its contained procedures, and the declarations and executables of a main program or procedure from its contained procedure.

Continue Statement

```
CONTINUE
```

This is an executable statement that has no effect. If it has a label it may be used as the terminating statement of a **DO** construct or as the target of a **GOTO**, **computed-GOTO** or **assigned-GOTO**.

Cycle Statement

```
CYCLE [ construct-name ]
```

Begins the next iteration of either the specified **DO** construct, or if *construct-name* is omitted, the innermost enclosing **DO** construct.

Data Statement

```
DATA data-set [ , data-set ]...
```

data-set ::= *data-object-list* / *data-value-list* /

```
data-object-list ::= data-object [ , data-object ]...  
data-object ::= variable | data-implied-do  
data-implied-do ::= ( data-object [ , data-object ]... do-spec )  
  
data-value-list ::= data-value [ , data-value ]...  
data-value ::= [ data-repeat * ] data-constant  
data-repeat ::= constant | constant-subobject  
  
data-constant ::= literal-constant | NULL() | structure-constructor | object |  
                  { + | - } { real-literal | integer-literal }
```

Declares the initial value of the specified objects. This implicitly declares those objects to have the **SAVE** attribute.

Deallocate Statement

```
DEALLOCATE ( expr [ , expr ] [ , STAT = variable ] )
```

Deallocates the storage occupied by an allocatable array or pointer. An error is raised if an allocatable array to be deallocated is not allocated, or if a pointer to be deallocated is dissociated or is associated with an object that was not allocated with **ALLOCATE**.

Dimension Statement

```
DIMENSION [ :: ] name array-spec [ , name array-spec ]...
```

Declares the name(s) to be arrays with the specified bounds.

Do Statement

```
[ construct-name : ] DO [ label ] [ , ] [ loop-control ]  
  
loop-control ::= do-spec |  
                  WHILE ( logical-expr )  
do-spec ::= name = expr , expr [ , expr ]
```

The initial statement of a **DO** loop. If *label* is present, the loop ends on the statement with that label, which cannot be a **GOTO**, **RETURN**, **STOP**, **EXIT**, **CYCLE**, **END** or arithmetic **IF** statement. Nested **DO** loops can share the same ending statement, provided it is not an **ENDDO** statement. If the *loop-control* is missing, the **DO** loop terminates only if control is explicitly transferred outside the loop (e.g., by an **EXIT**, **GOTO** or **RETURN** statement).

If *construct-name* is present, the **DO** loop must end with an **ENDDO** statement identified with the same *construct-name*.

Else Statement

```
ELSE [ construct-name ]
```

Begins the **ELSE** part of an **IF-THEN** construct. Statements in this part are executed only if the **IF** condition is false and all **ELSEIF** conditions at the same level are false. If the **IF-THEN** statement had a *construct-name*, the **ELSE** statement may specify the same *construct-name*.

Elseif Statement

```
ELSE IF ( expr ) THEN [ construct-name ]
```

Begins a (new) **ELSEIF** part of an **IF-THEN** construct. Statements in this part are executed only if the **IF** condition is false, all preceding **ELSEIF** conditions at the same level are false, and this **ELSEIF** condition is true. If the **IF-THEN** statement had a *construct-name*, the **ELSEIF** statement may specify the same *construct-name*.

Elsewhere Statement

ELSEWHERE [*construct-name*]

Begins the ELSEWHERE part of a WHERE construct. The statements in this part are executed only for those elements for which the WHERE mask are false, and all ELSEWHERE masks at the same level are also false.

If the WHERE statement had a construct-name, the ELSEWHERE statement may specify the same construct-name.

Elsewhere Mask Statement

ELSEWHERE (*expr*) [*construct-name*]

Begins a masked ELSEWHERE part of a WHERE construct. The statements in this part are executed only for those elements for which the previous masks are false and this ELSEWHERE mask are true. (The previous masks are the WHERE mask and all preceding ELSEWHERE masks at the same level in this WHERE construct.) Note that the elements of the ELSEWHERE mask that do not correspond to false elements of the previous masks are not evaluated.

If the WHERE statement had a construct-name, the ELSEWHERE-mask statement may specify the same construct-name.

End Block Data Statement

END [BLOCK DATA [*name*]]

The last statement of a block data subprogram. If *name* is present, the BLOCK DATA statement at the beginning of the subprogram must have specified the same name.

Enddo Statement

END DO [*construct-name*]

Marks the end of a DO construct. The *construct-name* shall be present if and only if it were present on the DO statement, and must be the same construct-name if so. If the DO statement specifies an ending label, the ENDDO statement must be labelled with that label.

Endfile Statement

ENDFILE *expr*
ENDFILE (*position-spec-list*)

(See the BACKSPACE statement for the *position-spec-list* definition).

Writes an endfile record to the specified external file, truncating it at the current point.

End Function Statement

END [FUNCTION [*name*]]

The last statement of a function subprogram. If the function subprogram is a contained subprogram, the keyword FUNCTION must be present. If *name* is present, it must be the name of the function.

Endif Statement

END IF [*construct-name*]

Marks the end of an IF-THEN construct. The *construct-name* shall be present if and only if it were present on the IF-THEN statement, and must be the same construct-name if so.

End Interface Statement

END INTERFACE [*generic-spec*]

Marks the end of an interface block. If the `INTERFACE` statement had a *generic-spec*, it may appear on the `ENDINTERFACE` statement.

End Module Statement

```
END [ MODULE [ name ] ]
```

The final statement of a module subprogram. If *name* is present, it must match the name on the `MODULE` statement.

End Program Statement

```
END [ PROGRAM [ name ] ]
```

The final statement of a main program unit. If *name* is present, the main program must have a `PROGRAM` statement and the names must be the same.

End Select Statement

```
END SELECT [ construct-name ]
```

Marks the end of a `SELECT` construct. The *construct-name* shall be present if and only if it were present on the `SELECT` statement, and must be the same construct-name if so.

End Type Statement

```
END TYPE [ name ]
```

Marks the end of a derived type definition. If *name* is present it must be the name of the derived type.

Endwhere Statement

```
ENDWHERE [ construct-name ]
```

Marks the end of a `WHERE` construct. The *construct-name* shall be present if and only if it were present on the `WHERE` statement, and must be the same construct-name if so.

Equivalence Statement

```
EQUIVALENCE equivalence-set [ , equivalence-set ]...  
equivalence-set ::= ( variable { , variable }... )
```

Declares each object in an equivalence-set to occupy the same storage.

Entry Statement

```
ENTRY name [ ( [ arg-list ] ) ]
```

Declares an additional entry point to the enclosing subprogram (entry points are not allowed in block data, main program, module and internal subprograms).

External Statement

```
EXTERNAL name [ , name ]...
```

Declares the listed names to be external subprograms or block data subprograms.

Exit Statement

EXIT [*construct-name*]

Transfers control to the statement following named DO loop or, if *construct-name* is omitted, the innermost enclosing DO loop.

Forall Statement

FORALL (*triplet-spec* [, *triplet-spec*]... [, *expr*]) *forall-assignment-stmt*
triplet-spec ::= *name* = *expr* : *expr* [: *expr*]

The iteration space of a FORALL statement or construct is the cross-product of the sets of possible index values defined by each *triplet-spec* masked by the final *expr* (if present). Note that the scope of the index names is limited to the FORALL statement – a variable with the same name outside the FORALL statement is unaffected.

The FORALL statement executes the *forall-assignment* statement for each index value set in the iteration space.

Forall Assignment Statement

variable = *expr*
variable => *expr*

This is exactly like a normal assignment statement except that the *expr* is evaluated for each element of the iteration space before assignment or pointer assignment to each *variable*. Note that an assignment must not assign to the same element of an array more than once in the iteration space, and if the variable is scalar then the iteration space must be exactly one element.

Forall Construct Statement

FORALL (*triplet-spec* [, *triplet-spec*]... [, *expr*])

(See the FORALL statement for the *triplet-spec* definition and the explanation of the iteration space.)

Begins a FORALL construct.

Format Statement

FORMAT ([*format-list*])

format-list ::= *format-item* [, *format-item*]...

format-item ::= [*digit-string*] { *data-edit* | (*format-list*) } | *other-edit*

data-edit ::= { I | B | O | Z } *digit-string* [. *digit-string*] |
 { F | D } *digit-string* . *digit-string* |
 { E | EN | ES | G }
 digit-string . *digit-string* [E *digit-string*] |
 L *digit-string* |
 A [*digit-string*]

other-edit ::= *digit-string* { / | P | X } |
 { T | TR | TL } *digit-string* |
 character-literal |
 digit-string H *char...* |
 / | : | BN | BZ | S | SP | SS

Note: The *character-literal* must not have a *kind-specifier*. The H edit descriptor is followed by *digit-string chars*, which may be any character except end-of-line; this edit descriptor is *obsolescent* and the *character-literal* one should be used instead.

Declares an i/o format.

Note: The comma between *format-items* may be omitted as follows:

1. Between a ‘P’ descriptor and a following ‘D’, ‘E’, ‘EN’, ‘ES’, ‘F’ or ‘G’ descriptor,
2. Before a ‘/’ descriptor with no preceding *digit-string*,
3. After a ‘/’ descriptor and
4. Before or after a ‘:’ descriptor.

Function Statement

[*prefix*] FUNCTION *name* ([*name* [, *name*] ...]) [RESULT(*name*)]

prefix :: = { *type-spec* | RECURSIVE | PURE | ELEMENTAL }...

Note: At most one occurrence of each prefix item is allowed.

This is the first statement of a function subprogram. If no **RESULT** variable is specified the result variable has the same name as the function name (thus for direct recursion, a **RESULT** clause is necessary as well as the **RECURSIVE** keyword).

Goto Statement

GOTO *label*

Branches to the specified label, which must be on a branch target statement (i.e., the subprogram **END** statement, an executable statement, the first statement of an executable construct or the last statement of an enclosing executable construct).

If Statement

IF (*expr*) *executable*

Executes the sub-statement if and only if the condition is true. The sub-statement cannot itself be an **IF** statement.

If Then Statement

[*construct-name* :] IF (*expr*) THEN

Begins an **IF-THEN** construct and the **THEN** part thereof. Statements in this part are executed if and only if the condition is true. This statement may have a construct-name; if it does, the corresponding **ENDIF** statement shall have the same construct-name and intervening **ELSE** and **ELSEIF** statements at the same level may have the same construct-name.

Implicit Statement

IMPLICIT *implicit-spec* [, *implicit-spec*]...

implicit-spec ::= *type-spec* (*letter-spec* [, *letter-spec*] ...)

letter-spec ::= *letter* [- *letter*]

Alters the implicit type mapping from the default. The default map is

IMPLICIT REAL(A-H,O-Z),INTEGER(I-N)

in an external subprogram or interface body, and the same as the containing subprogram in a contained subprogram.

Implicit None Statement

IMPLICIT NONE

This statement sets the implicit type mapping for each letter to null, i.e., there are no implicit types. It must occur before any **PARAMETER** statements or other declarations (but after any **USE** statements).

Inquire Statement

```
INQUIRE ( IOLENGTH=object ) output-item [ , output-item ]...  
INQUIRE ( inquire-spec [ , inquire-spec ]... )
```

```
inquire-spec ::= [ UNIT= ] expr | ACCESS= variable | ACTION= variable | BLANK= variable | CONVERT= variable |  
DELIM= variable | DIRECT= variable | ERR= label | EXIST= variable | FILE= expr |  
FORM= variable | FORMATTED= variable | IOSTAT= variable | NAME= variable |  
NAMED= variable | NEXTREC= variable | NUMBER= variable | OPENED= variable |  
PAD= variable | POSITION= variable | READ= variable | READWRITE= variable |  
RECL= variable | SEQUENTIAL= variable | UNFORMATTED= variable
```

```
output-item ::= expr | ( { output-item , } ... do-spec )
```

The first form enquires as to the length needed to be specified for **RECL=** in the **OPEN** statement for an unformatted sequential file to be able to write records as large as the *output-item* list.

The second form enquires either by unit or by file; exactly one **UNIT=** or **FILE=** clause must be present (the **UNIT=** keyword can be omitted if it is the first *inquire-spec*). If the **FILE=** clause is used and that file is currently connected to a unit, the effect is as if that unit were specified.

The effect of each clause is as below:

- ACCESS=** Sets the scalar character object to 'SEQUENTIAL' if the unit is connected for sequential access, to 'DIRECT' if the unit is connected for direct access, and to 'UNDEFINED' if there is no connection.
- ACTION=** Sets the scalar character object to 'READ' if the unit is connected for input only, to 'WRITE' if the unit is connected for output only, to 'READWRITE' if the unit is connected for both input and output, and to 'UNDEFINED' if there is no connection.
- BLANK=** If the unit is connected for formatted i/o, sets the scalar character object to 'NULL' if blanks are treated as nulls on input and to 'ZERO' if they are treated as zeroes on input. Otherwise the object is set to 'UNDEFINED'.
- CONVERT=** Sets the scalar character object to 'UNKNOWN' if the file is not connected for unformatted input/output, and otherwise to the value of the **CONVERT=** specifier in the **OPEN** statement or as determined by the **FORT.CONVERT_n** environment variable. Possible values are 'NATIVE', 'BIG_NATIVE', 'LITTLE_NATIVE', 'BIG_IEEE', 'LITTLE_IEEE', 'BIG_IEEE_DD' and 'LITTLE_IEEE_DD'.
- DELIM=** Sets the scalar character object to 'APOSTROPHE' if the apostrophe is used to delimit character data in list-directed or namelist output for the unit, to 'QUOTE' if the quotation mark is to be so used, to 'NONE' if no delimiter is to be used, and to 'UNDEFINED' if there is no connection.
- DIRECT=** Sets the scalar character object to 'YES' if direct access is allowed for the unit, to 'NO' if direct access is not allowed, and to 'UNKNOWN' if the answer cannot be determined.
- ERR=** Transfers control to the specified label if an error condition occurs during execution of the enquiry.
- EXIST=** Sets the scalar default logical object to .TRUE. if the file or unit exists, and to .FALSE. otherwise.
- FORM=** Sets the scalar character object to 'FORMATTED' if the unit is connected for formatted i/o, to 'UNFORMATTED' if the unit is connected for unformatted i/o, and to 'UNDEFINED' if there is no connection.
- FORMATTED=**
Sets the scalar character object to 'YES' if formatted i/o is allowed for the unit, to 'NO' if it is not allowed, and to 'UNKNOWN' if the answer cannot be determined.
- IOSTAT=** Sets the scalar default integer object to a non-zero value to indicate an error condition occurring during the enquiry, and to zero otherwise.
- NAME=** Sets the scalar character object to the name of the file connected to the unit (undefined if there is no name).

- NAMED=** Sets the scalar default logical object to `.TRUE.` if the unit is connected to a named file, and to `.FALSE.` otherwise.
- NEXTREC=** If the unit is connected for direct access, sets the scalar default integer object to 1 greater than the number of the record last read or written (or to 1 if no record has been read or written); otherwise the object becomes undefined.
- NUMBER=** Sets the scalar default integer object to the number of the unit to which the file is connected, or to `-1` if the file is not connected to a unit.
- OPENED=** Sets the scalar default logical object to `.TRUE.` if the unit is connected to a file (or the file is connected to a unit), and to `.FALSE.` otherwise.
- PAD=** Sets the scalar character object to `'NO'` if the connection of the file to the unit included the `PAD='NO'` specifier; otherwise it is set to `'YES'`.
- POSITION=**
Sets the scalar character object to `'REWIND'` if the unit is positioned at the beginning of the file, to `'APPEND'` if it is positioned at the end of the file, to `'ASIS'` if the unit was connected with that specification (and no i/o or positioning has occurred since connection), to `'UNDEFINED'` if the unit is connected for direct access or there is no connection, and to a processor-dependent value otherwise.
- READ=** Sets the scalar character object to `'YES'` if input is allowed for the file or unit, to `'NO'` if input is not allowed for the file or unit, and to `'UNKNOWN'` if the answer cannot be determined.
- READWRITE=**
Sets the scalar character object to `'YES'` if both input and output are allowed for the file or unit, to `'NO'` if at least one of input or output is not allowed for the file or unit, and to `'UNKNOWN'` if the answer cannot be determined.
- RECL=** Sets the scalar default integer object to the record length of a file connected for direct access or to the maximum record length of a file connected for sequential access. If there is no connection the object is undefined.
- SEQUENTIAL=**
Sets the scalar character object to `'YES'` if sequential access is allowed for the unit, to `'NO'` if sequential access is not allowed, and to `'UNKNOWN'` if the answer cannot be determined.
- UNFORMATTED=**
Sets the scalar character object to `'YES'` if unformatted i/o is allowed for the unit, to `'NO'` if it is not allowed, and to `'UNKNOWN'` if the answer cannot be determined.
- WRITE=** Sets the scalar character object to `'YES'` if output is allowed for the file or unit, to `'NO'` if output is not allowed for the file or unit, and to `'UNKNOWN'` if the answer cannot be determined.

Intent Statement

`INTENT ({ IN | OUT | INOUT }) [::] name [, name]...`

Declares the specified names, which must be the names of dummy arguments, to have the specified intent. `INTENT(IN)` arguments cannot appear in any context where they will be modified, `INTENT(OUT)` arguments are undefined on entry to the procedure, and `INTENT(INOUT)` and `INTENT(OUT)` arguments can only be associated with modifiable actual arguments (e.g., not expressions).

Interface Statement

`INTERFACE`
`INTERFACE { name | ASSIGNMENT(=) | OPERATOR(operator) }`

The first form introduces an interface block, containing interface bodies which specify the interfaces to external or dummy procedures. The second form additionally defines a generic name or operator by which these procedures may be referenced, and its interface block may also contain `MODULE PROCEDURE` statements.

Intrinsic Statement

INTRINSIC *name* [, *name*]...

Declares the listed names to be intrinsic procedures.

Module Statement

MODULE *name*

This is the first statement of a module subprogram.

Module Procedure Statement

MODULE PROCEDURE *name* [, *name*]...

This statement is only allowed within generic interface blocks, where it declares the listed names as module procedures to be included in the generic.

Namelist Statement

NAMELIST *namelist-group* [[,] *namelist-group*]...

namelist-group ::= /*name*/ *name* [, *name*]...

Declares one or more i/o namelists. Multiple **NAMELIST** specifications for the namelist group /*name*/ are treated as if they were concatenated. The names in a namelist group must all be variables and not automatic, adjustable, allocatable, pointer, or contain a pointer.

Nullify Statement

NULLIFY (*object* [, *object*] ...)

Sets the pointer-association status of the listed objects, which must be pointers, to dissociated.

Open Statement

OPEN (*open-spec* [, *open-spec*]...)

open-spec ::= [UNIT=] *expr* | ACCESS= *expr* | ACTION= *expr* | BLANK= *expr* | CONVERT= *expr* | DELIM= *expr* |
ERR= *label* | FILE= *expr* | FORM= *expr* | IOSTAT= *object* | PAD= *expr* | POSITION= *expr* |
RECL= *expr* | STATUS= *expr*

Connects a file to a unit with the specified properties.

Optional Statement

OPTIONAL [::] *name* [, *name*]...

Declares the specified names, which must be the names of dummy arguments, to be optional dummy arguments.

Parameter Statement

PARAMETER (*name* = *expr* [, *name* = *expr*]...)

Declares the names to be named constants with the specified values. The expressions must be initialisation expressions and must be assignment compatible with the names.

Pause Statement (*deleted*)

PAUSE [*constant*]

Pauses program execution. If present, the constant must be a scalar character literal with no *kind-param* or a *digit-string* with at most 5 digits.

Pointer Statement

POINTER [::] *name* [*deferred-shape*] [, *name* [*deferred-shape*]]...

Declares the names to be pointers.

Pointer Assignment Statement

variable => *expr*

Associates the pointer *variable* with *expr*, which must be another pointer, a variable with the TARGET attribute, or a reference to a function that returns a pointer result.

Print Statement

PRINT *format* [, *output-item*]...

format ::= * | *label* | *expr*

Synonymous with a WRITE statement with 'UNIT=*' and a FMT=*format* clause.

The possibilities for *format* are:

'*' indicates list-directed formatting.

label must be the label of a FORMAT statement.

expr A character expression may be supplied; its value is interpreted as the text following the keyword FORMAT of a FORMAT statement. If the expression is array-valued the concatenation of all elements is interpreted in this way.

expr (*obsolescent*) A default scalar integer variable name may be supplied, in which case it must have been ASSIGNED the label of a FORMAT statement.

Private Statement

PRIVATE [[::] *access-id* [, *access-id*]...]

access-id ::= *name* | ASSIGNMENT(=) | OPERATOR(*operator*)

This statement can only occur in the declaration section of a module or before the component definitions in a type definition.

When this statement appears in a type definition, there can be no *access-ids*; it causes the components of the type to be inaccessible from outside the module in which the type is defined.

In a module's declaration section, this statement either sets the default accessibility of entities within the module to be PRIVATE, i.e., not accessible, or the accessibility of each *access-id* is set to be PRIVATE.

Program Statement

PROGRAM *name*

This is the first statement of a main program. It is optional.

Public Statement

```
PUBLIC [ [ :: ] access-id [ , access-id ]... ]
```

This statement can only occur in the declaration section of a module. With no *access-id* list, it confirms that the default accessibility of entities in the module is **PUBLIC**. With an *access-id* list, it explicitly sets the accessibility of those *access-ids* to **PUBLIC**.

Read Statement

```
READ format [ , input-item ]...
```

```
READ ( control-spec [ , control-spec ]... ) [ input-item [ , input-item ]... ]
```

```
input-item ::= variable | ( { input-item , }... do-spec )
```

```
control-spec ::= [ UNIT= ] { * | expr } |  
[ FMT= ] format | [ NML= ] name | ADVANCE= expr | END= label |  
EOR= label | ERR= label | IOSTAT= expr | REC= expr | SIZE= expr
```

(See the PRINT statement for *format* details.)

Reads one or more records (or partial records with **ADVANCE='NO'**) from the specified unit.

The effect of each control-specifier is as below:

UNIT=	Identifies the i/o unit; '*' indicates the default unit, a scalar integer expression indicates an external unit, and a character expression indicates an internal unit.
FMT=	Establishes the format; this is absent for namelist formatting or for unformatted i/o.
NML=	Specifies a namelist group name for namelist formatted i/o.
ADVANCE=	Indicates whether non-advancing (expression evaluates to 'NO') or the usual advancing (expression evaluates to 'YES') i/o is performed. This control-specifier is only allowed for formatted sequential i/o with an explicit format (i.e., not namelist or list-directed).
END=	Transfers control to the specified label if an end-of-file condition occurs during input (not allowed in WRITE).
EOR=	Transfers control to the specified label if an end-of-record condition occurs during input (not allowed in WRITE). ADVANCE='NO' must be specified.
ERR=	Transfers control to the specified label if an error condition occurs during i/o.
IOSTAT=	Sets the object to -1 if an end-of-file occurs, to -2 if an end-of-record occurs (non-advancing only), to a positive number if an error occurs, and to zero otherwise. Note that the negative values may vary on other compilers.
REC=	Specifies the record number for direct-access i/o.
SIZE=	Sets the object to the number of characters transferred by data edit descriptors (not allowed in WRITE). ADVANCE='NO' must be specified.

Return Statement

```
RETURN [ expr ]
```

Return immediately from the procedure. If the procedure is a subroutine with alternate return arguments (*obsolescent*), the scalar integer expression indicates to which label control is to be transferred on return (if the expression is less than one or greater than the number of alternate return arguments, execution continues with the statement following the subroutine reference).

Rewind Statement

```
REWIND expr
```

```
REWIND ( position-spec-list )
```

(See the BACKSPACE statement for the *position-spec-list* definition).

Positions an i/o unit, which must be connected to a rewindable file, to the beginning of the file.

Save Statement

```
SAVE [ [ :: ] save-item [ , save-item ]... ]  
save-item ::= variable-name | /common-block-name/
```

Specifies the **SAVE** attribute for the listed variables or common blocks, or, with no *save-item* list, specifies that all possible variables and common blocks in the current scoping unit should implicitly have the **SAVE** attribute by default.

Select Statement

```
[ construct-name : ] SELECT CASE ( expr )
```

The initial statement of a **SELECT CASE** construct. Control is transferred to the **CASE** statement satisfied by the expression's value, or to the **END SELECT** statement if no **CASE** is satisfied by the value.

Statement Function Statement (*obsolescent*)

```
name ( [ name [ , name ] ... ] ) = expr
```

Defines a statement function.

Stop Statement

```
STOP [ constant ]
```

Halts program execution. If present, the constant must be a scalar character literal with no *kind-param* or a *digit-string* with at most 5 digits.

Subroutine Statement

```
[ RECURSIVE | PURE | ELEMENTAL ]... SUBROUTINE name [ ( [ arg-list ] ) ]
```

(Note that at most one occurrence of each keyword is allowed).

```
arg-list ::= arg [ , arg ]...  
arg ::= name | *
```

This is the first statement of a subroutine subprogram. **RECURSIVE** must be specified if the subroutine calls itself, either directly or indirectly. If **PURE** is specified, the subroutine must satisfy the pure subroutine constraints and can then be called from a pure function. An *arg* that is '*' signifies an alternate return label; this is *obsolescent*.

Target Statement

```
TARGET [ :: ] name [ array-spec ] [ , name [ array-spec ] ]...
```

Declares that the specified entities have the **TARGET** attribute.

Type Statement

```
TYPE name
```

This statement marks the beginning of the definition of the derived type *name*.

Type Declaration Statement

type-spec [[, *attr-spec*] ... ::] *entity-decl-list*

attr-spec ::= ALLOCATABLE | DIMENSION *array-spec* | EXTERNAL | INTENT ({ IN | OUT | INOUT }) |
INTRINSIC | OPTIONAL | PARAMETER | POINTER | PRIVATE | PUBLIC | SAVE | TARGET

Declares the listed entities to be of the specified type with the specified attributes.

Use Statement

USE *name* [, *rename-list*]

USE *name*, ONLY: *only-list*

rename-list ::= *rename* [, *rename*]...

rename ::= *local-name* => *remote-name*

only-list ::= *only-item* [, *only-item*]...

only-item ::= *name* | *rename*

The **USE** statement accesses the named module. Multiple **USE** statements for the same module act as if all the *rename-lists* and *only-lists* were concatenated.

If all the **USE** statements in a scoping unit for a particular module have the **ONLY** clause, only those items listed in a *rename-list* or *only-list* are accessible.

A *rename* causes item *remote-name* in the referenced module to be accessible in the local scoping unit by *local-name*. An *only-item* that is not a *rename* causes the *name* in the referenced module to be accessible in the local scoping unit by the same name.

Where Assignment Statement

variable = *expr*

The expression is evaluated (and the object updated) only for those elements for which the current control mask is true.

Where Statement

WHERE (*expr*) *where-assignment-stmt*

Executes the Where Assignment statement with the provided expression as the control mask.

Where Construct Statement

[*construct-name* :] WHERE (*expr*)

Begins a Where Construct with the provided expression as the control mask.

Write Statement

WRITE (*control-spec* [, *control-spec*] ...) *output-item* [, *output-item*]...

(See the **READ** statement for *control-spec* details.)

Writes one or more records (or partial records with **ADVANCE='NO'**) to the specified unit.

63 Fortran 95 Intrinsic Procedures

This section provides a quick reference guide to the Fortran 90/95 intrinsic functions and subroutines.

Procedures marked with ‘*’ are non-generic versions of other intrinsics. Procedures marked with ‘*E*’ are elemental.

Procedures marked with ‘*I*’ are inquiry functions; these query characteristics of a variable other than its value.

Procedures marked with ‘*P*’ may be supplied as actual procedure arguments; when it is used as a procedure argument all arguments are scalar, no optional arguments are allowed, and it is not generic: if the name is normally generic the procedure argument version is default **INTEGER** if it begins with the letter **I** to **N**, and default **REAL** otherwise.

Arguments named ‘**KIND**’ must be initialisation expressions.

Arguments printed in *italics* (e.g., ‘*DIM*’) are optional.

Intrinsic Functions

Function Name	Flags	Description
ABS(<i>A</i>)	<i>EP</i>	Absolute value.
ACHAR(<i>I</i>)	<i>E</i>	Produce character from ASCII value.
ACOS(<i>X</i>)	<i>EP</i>	Arccosine.
ADJUSTL(<i>STRING</i>)	<i>E</i>	Adjust string to the left by moving leading blanks to the end.
ADJUSTR(<i>STRING</i>)	<i>E</i>	Adjust string to the right.
AIMAG(<i>Z</i>)	<i>EP</i>	The imaginary part of a complex number.
AINT(<i>A</i> , <i>KIND</i>)	<i>EP</i>	Truncate to a whole number.
ALL(<i>MASK</i> , <i>DIM</i>)		Reduce <i>MASK</i> with the .AND. operation.
ALLOCATED(<i>ARRAY</i>)	<i>I</i>	Whether an allocatable array is allocated.
ALOG(<i>X</i>)	* <i>P</i>	LOG function restricted to scalar default real.
ALOG10(<i>X</i>)	* <i>P</i>	LOG10 function restricted to scalar default real.
AMAX0(<i>A1</i> , <i>A2</i> , <i>A3</i> , ...)	*	REAL(MAX (<i>A1</i> , <i>A2</i> , <i>A3</i> , ...)) with scalar default integer args.
AMAX1(<i>A1</i> , <i>A2</i> , <i>A3</i> , ...)	*	MAX function restricted to scalar default real.
AMINO(<i>A1</i> , <i>A2</i> , <i>A3</i> , ...)	*	REAL(MIN (<i>A1</i> , <i>A2</i> , <i>A3</i> , ...)) with scalar default integer args.
AMIN1(<i>A1</i> , <i>A2</i> , <i>A3</i> , ...)	*	MIN function restricted to scalar default integer.
AMOD(<i>A</i> , <i>P</i>)	* <i>P</i>	MOD function restricted to scalar default real.
ANINT(<i>A</i> , <i>KIND</i>)	<i>EP</i>	Round to a whole number.
ANY(<i>MASK</i> , <i>DIM</i>)		Reduce <i>MASK</i> with the .OR. operation.
ASIN(<i>X</i>)	<i>EP</i>	Arcsine.
ASSOCIATED(<i>POINTER</i> , <i>TARGET</i>)	<i>I</i>	Whether a pointer is associated with a target.
ATAN(<i>X</i>)	<i>EP</i>	Arctangent.
ATAN2(<i>Y</i> , <i>X</i>)	<i>EP</i>	Arctangent.
BIT_SIZE(<i>I</i>)		Number of bits in an integer type.
BTEST(<i>I</i> , <i>POS</i>)	<i>E</i>	Whether a particular bit is 1.
CABS(<i>A</i>)	* <i>P</i>	ABS function restricted to scalar default complex.
CCOS(<i>X</i>)	* <i>P</i>	COS function restricted to scalar default complex.
CEILING(<i>A</i> , <i>KIND</i>)	<i>E</i>	Smallest integer greater than or equal to <i>A</i> .
CEXP(<i>X</i>)	* <i>P</i>	EXP function restricted to scalar default complex.
CHAR(<i>I</i> , <i>KIND</i>)	<i>E</i>	Produce character from native coded character set value.
CLOG(<i>X</i>)	* <i>P</i>	LOG function restricted to scalar default complex.
CMPLX(<i>X</i> , <i>Y</i> , <i>KIND</i>)	<i>E</i>	Convert to complex.
CONJG(<i>Z</i>)	<i>EP</i>	Complex conjugate.

Function Name	Flags	Description
COS(X)	<i>EP</i>	Cosine.
COSH(X)	<i>EP</i>	Hyperbolic cosine.
COUNT(MASK, DIM)		Reduce MASK by counting .TRUE. elements.
CSHIFT(ARRAY, SHIFT, DIM)		Circular shift of an array.
CSIN(X)	<i>*P</i>	SIN function restricted to scalar default complex.
CSQRT(X)	<i>*P</i>	SQRT function restricted to scalar default complex.
DABS(A)	<i>*P</i>	ABS function restricted to scalar double precision.
DACOS(X)	<i>*P</i>	ACOS function restricted to scalar double precision.
DASIN(X)	<i>*P</i>	ASIN function restricted to scalar double precision.
DATAN(X)	<i>*P</i>	ATAN function restricted to scalar double precision.
DATAN2(X)	<i>*P</i>	ATAN2 function restricted to scalar double precision.
DBLE(A)	<i>E</i>	Convert to double precision.
DCOS(X)	<i>*P</i>	COS function restricted to scalar double precision.
DCOSH(X)	<i>*P</i>	COSH function restricted to scalar double precision.
DDIM(X, Y)	<i>*P</i>	DIM function restricted to scalar double precision.
DEXP(X)	<i>*P</i>	EXP function restricted to scalar double precision.
DIGITS(X)	<i>I</i>	Number of mantissa digits in the model for X.
DIM(X, Y)	<i>EP</i>	Non-negative difference, MAX(0, X-Y).
DINT(A)	<i>*P</i>	AINT function restricted to scalar double precision.
DLOG(X)	<i>*P</i>	LOG function restricted to scalar double precision.
DLOG10(X)	<i>*P</i>	LOG10 function restricted to scalar double precision.
DMAX1(A1, A2, A3, ...)	<i>*</i>	MAX function restricted to scalar double precision.
DMIN1(A1, A2, A3, ...)	<i>*</i>	MIN function restricted to scalar double precision.
DMOD(A, P)	<i>*P</i>	MOD function restricted to scalar double precision.
DNINT(A)	<i>*P</i>	ANINT function restricted to scalar double precision.
DOT_PRODUCT(VECTOR_A, VECTOR_B)		Dot product.
DPROD(X, Y)	<i>EP</i>	Double precision result of X*Y.
DSIGN(A, B)	<i>*P</i>	SIGN function restricted to scalar double precision.
DSIN(X)	<i>*P</i>	SIN function restricted to scalar double precision.
DSINH(X)	<i>*P</i>	SINH function restricted to scalar double precision.
DSQRT(X)	<i>*P</i>	SQRT function restricted to scalar double precision.
DTAN(X)	<i>*P</i>	TAN function restricted to scalar double precision.
DTANH(X)	<i>*P</i>	TANH function restricted to scalar double precision.
EOSHIFT(ARRAY, SHIFT, BOUNDARY, DIM)		End-off array shift.
EPSILON(X)	<i>I</i>	Number almost negligible compared to 1.
EXP(X)	<i>EP</i>	Exponential (e^X).
EXPONENT(X)	<i>E</i>	The exponent part of a floating-point number.
FLOAT(A)	<i>*</i>	REAL function restricted to scalar default integer.
FLOOR(A, KIND)	<i>E</i>	Largest integer less than or equal to A.
FRACTION(X)	<i>E</i>	The mantissa of a floating-point number.
HUGE(X)	<i>I</i>	The largest number in the model for a real type.
IABS(A)	<i>*P</i>	ABS function restricted to scalar default integer.
IACHAR(C)	<i>E</i>	ASCII value for a character.
IAND(I, J)	<i>E</i>	Bitwise and.
IBCLR(I, POS)	<i>E</i>	Clear specified bit.
IBITS(I, POS, LEN)	<i>E</i>	Extract a group of bits.
IBSET(I, POS)	<i>E</i>	Set a bit.
ICHAR(C)	<i>E</i>	Native coded character set value for a character.
IDIM(X, Y)	<i>*P</i>	DIM function restricted to scalar default integer.
IDINT(A)	<i>*</i>	INT function restricted to scalar double precision.
IDNINT(A)	<i>*P</i>	NINT function restricted to scalar double precision.

Function Name	Flags	Description
IEOR(I, J)	<i>E</i>	Bitwise exclusive or.
IFIX(A)	*	INT function restricted to scalar default real.
INDEX(STRING, SUBSTRING, BACK)	<i>EP</i>	Search for substring.
INT(A, KIND)	<i>E</i>	Convert to integer.
IOR(I, J)	<i>E</i>	Bitwise inclusive or.
ISHFT(I, SHIFT)	<i>E</i>	Shift bits.
ISHFTC(I, SHIFT, SIZE)	<i>E</i>	Shift bits circularly.
ISIGN(A, B)	* <i>P</i>	SIGN function restricted to scalar default integer.
KIND(X)	<i>I</i>	Kind type parameter of X.
LBOUND(ARRAY, DIM)	<i>I</i>	Lower bound(s) of an array.
LEN(STRING)	<i>IP</i>	Length of a character string.
LEN_TRIM(STRING)	<i>E</i>	Length of a character string ignoring trailing blanks.
LGE(STRING_A, STRING_B)	<i>E</i>	Comparison (\geq) using ASCII collating sequence.
LGT(STRING_A, STRING_B)	<i>E</i>	Comparison ($>$) using ASCII collating sequence.
LLE(STRING_A, STRING_B)	<i>E</i>	Comparison (\leq) using ASCII collating sequence.
LLT(STRING_A, STRING_B)	<i>E</i>	Comparison ($<$) using ASCII collating sequence.
LOG(X)	<i>EP</i>	Natural logarithm.
LOG10(X)	<i>EP</i>	Common logarithm.
LOGICAL(L, KIND)	<i>E</i>	Convert to a specific logical kind.
MATMUL(MATRIX_A, MATRIX_B)		Matrix multiplication.
MAX(A1, A2, A3, ...)	<i>E</i>	Maximum value.
MAX0(A1, A2, A3, ...)	*	MAX function restricted to scalar default integer.
MAX1(A1, A2, A3, ...)	*	INT(MAX(A1, A2, A3, ...)) with scalar default real args.
MAXEXPONENT(X)	<i>I</i>	Maximum model exponent value for X.
MAXLOC(ARRAY, MASK)		Position of maximum value in an array.
MAXLOC(ARRAY, DIM, MASK)		Dimensional reduction of maximum value positions.
MAXVAL(ARRAY, MASK)		Reduce a (masked) array with the MAX intrinsic.
MAXVAL(ARRAY, DIM, MASK)		Dimensional reduction by the MAX intrinsic.
MERGE(TSOURCE, FSOURCE, MASK)	<i>E</i>	Choose value depending on logical value.
MIN(A1, A2, A3, ...)	<i>E</i>	Minimum value.
MIN0(A1, A2, A3, ...)	*	MIN function restricted to scalar default integer.
MIN1(A1, A2, A3, ...)	*	INT(MIN(A1, A2, A3, ...)) with scalar default real args.
MINEXPONENT(X)	<i>I</i>	Minimum model exponent value for X.
MINLOC(ARRAY, MASK)		Position of minimum value in an array.
MINLOC(ARRAY, DIM, MASK)		Dimensional reduction of minimum value positions.
MINVAL(ARRAY, MASK)		Reduce a (masked) array with the MIN intrinsic.
MINVAL(ARRAY, DIM, MASK)		Dimensional reduction by the MIN intrinsic.
MOD(A, P)	<i>EP</i>	Remainder; sign(result) = sign(A).
MODULO(A, P)	<i>E</i>	Modulo; sign(result) = sign(P).
NEAREST(X, S)	<i>E</i>	Nearest machine-representable number.
NINT(A, KIND)	<i>EP</i>	Round and convert to integer.
NOT(I)	<i>E</i>	Bitwise complement.
NULL(MOLD)		Null (disassociated) pointer.
PACK(ARRAY, MASK, VECTOR)		Pack an array into a vector.
PRECISION(X)	<i>I</i>	Decimal model precision for X.
PRESENT(A)	<i>I</i>	Whether an optional argument is present.
PRODUCT(ARRAY, MASK)		Reduce a (masked) array by multiplication.
PRODUCT(ARRAY, DIM, MASK)		Dimensional reduction by multiplication.
RADIX(X)	<i>I</i>	Model radix for X.
RANGE(X)	<i>I</i>	Decimal model exponent range for X.
REAL(A, KIND)	<i>E</i>	Convert to real.

Function Name	Flags	Description
REPEAT(<i>STRING</i> , <i>NCOPIES</i>)		Concatenate a string with itself.
RESHAPE(<i>SOURCE</i> , <i>SHAPE</i> , <i>PAD</i> , <i>ORDER</i>)		Reshape an array.
RRSPACING(<i>X</i>)	<i>E</i>	Reciprocal relative model spacing near <i>X</i> .
SCALE(<i>X</i> , <i>I</i>)	<i>E</i>	$X * RADIX(X) ** I$.
SCAN(<i>STRING</i> , <i>SET</i> , <i>BACK</i>)	<i>E</i>	Look for characters in a set.
SELECTED_INT_KIND(<i>R</i>)		Integer kind with at least <i>R</i> decimal digits.
SELECTED_REAL_KIND(<i>P</i> , <i>R</i>)		Real kind with at least <i>P</i> decimal precision and/or at least <i>R</i> decimal exponent range.
SET_EXPONENT(<i>X</i> , <i>I</i>)	<i>E</i>	$X * RADIX(X) ** (I - EXPONENT(X))$.
SHAPE(<i>SOURCE</i>)	<i>I</i>	Shape of an array.
SIGN(<i>A</i> , <i>B</i>)	<i>EP</i>	<i>A</i> with the sign of <i>B</i> .
SIN(<i>X</i>)	<i>EP</i>	Sine.
SINH(<i>X</i>)	<i>EP</i>	Hyperbolic sine.
SIZE(<i>ARRAY</i> , <i>DIM</i>)	<i>I</i>	Size of an array or dimension.
SNGL(<i>A</i>)	<i>*</i>	REAL function restricted to scalar double precision.
SPACING(<i>X</i>)	<i>E</i>	Spacing of model numbers near <i>X</i> .
SPREAD(<i>SOURCE</i> , <i>DIM</i> , <i>NCOPIES</i>)		Replicate an array in a given dimension by copying.
SQRT(<i>X</i>)	<i>EP</i>	Square root.
SUM(<i>ARRAY</i> , <i>MASK</i>)		Reduce a (masked) array by addition.
SUM(<i>ARRAY</i> , <i>DIM</i> , <i>MASK</i>)		Dimensional reduction of a (masked) array by addition.
TAN(<i>X</i>)	<i>EP</i>	Tangent.
TANH(<i>X</i>)	<i>EP</i>	Hyperbolic tangent.
TINY(<i>X</i>)	<i>I</i>	Smallest model number for the kind of <i>X</i> .
TRANSFER(<i>SOURCE</i> , <i>MOLD</i> , <i>SIZE</i>)		Copy internal representation.
TRANSPOSE(<i>MATRIX</i>)		Transpose an array.
TRIM(<i>STRING</i>)		Character string with trailing blanks removed.
UBOUND(<i>ARRAY</i> , <i>DIM</i>)	<i>I</i>	Upper bound(s) of an array.
UNPACK(<i>VECTOR</i> , <i>MASK</i> , <i>FIELD</i>)		Unpack a vector into an array.
VERIFY(<i>STRING</i> , <i>SET</i> , <i>BACK</i>)	<i>E</i>	Look for characters not in set.

Intrinsic Subroutines

Subroutine Name	Flags	Description
CPU_TIME(<i>TIME</i>)		CPU execution time.
DATE_AND_TIME(<i>DATE</i> , <i>TIME</i> , <i>ZONE</i> , <i>VALUES</i>)		Date and time information.
MVBITS(<i>FROM</i> , <i>FROMPOS</i> , <i>LEN</i> , <i>TO</i> , <i>TOPOS</i>)	<i>E</i>	Move or copy bitfield.
RANDOM_NUMBER(<i>HARVEST</i>)		Return pseudo-random number(s).
RANDOM_SEED(<i>SIZE</i> , <i>PUT</i> , <i>GET</i>)		Control pseudo-random number generator.
SYSTEM_CLOCK(<i>COUNT</i> , <i>COUNT_RATE</i> , <i>COUNT_MAX</i>)		Real-time clock information.

Fortran 2003 Extensions

64 Fortran 2003 Overview

This part of the manual describes those parts of the Fortran 2003 language which are not in Fortran 95, and indicates which features are currently supported by the NAG Fortran Compiler.

Features marked in the section heading as ‘[5.3]’ are newly available in release 5.3, those marked ‘[5.2]’ were available in release 5.2, those marked ‘[5.1]’ were available in release 5.1 (and in some cases earlier), and those marked ‘[n/a]’ are not yet available.

Fortran 2003 is a major advance over Fortran 95: the new language features can be grouped as follows:

- object-oriented programming features,
- allocatable attribute extensions,
- other data-oriented enhancements,
- interoperability with C,
- IEEE arithmetic support,
- input/output enhancements and
- miscellaneous enhancements.

The basic object-oriented features are type extension, polymorphic variables, and type selection; these provide inheritance and the ability to program ad-hoc polymorphism in a type-safe manner. The advanced features are typed allocation, cloning, type-bound procedures, type-bound generics, and object-bound procedures. Type-bound procedures provide the mechanism for dynamic dispatch (methods).

The **ALLOCATABLE** attribute is extended to allow it to be used for dummy arguments, function results, structure components, and scalars (not just arrays). An intrinsic procedure has been added to transfer an allocation from one variable to another. Finally, in intrinsic assignment, allocatable variables or components are automatically reallocated with the correct size if they have a different shape or type parameter value from that of the expression. This last feature, together with deferred character length, provides the user with true varying-length character variables.

There are two other major data enhancements: the addition of type parameters to derived types, and finalisation (by final subroutines). Other significant data enhancements are the **PROTECTED** attribute, pointer bounds specification and rank remapping, procedure pointers, and individual accessibility control for structure components.

Interoperability with the C programming language consists of allowing C procedures to be called from Fortran, Fortran procedures to be called from C, and for the sharing of global variables between C and Fortran. This can only happen where C and Fortran facilities are equivalent: an intrinsic module provides derived types and named constants for mapping Fortran and C types, and the **BIND(C)** syntax is added for declaring Fortran entities that are to be shared with C. Additionally, C style enumerations have been added.

Support for IEEE arithmetic is provided by three intrinsic modules. Use of the **IEEE_FEATURES** module requests IEEE compliance for specific Fortran features, the **IEEE_EXCEPTIONS** module provides access to IEEE modes and exception handling, and the **IEEE_ARITHMETIC** module provides enquiry functions and utility functions for determining the extent of IEEE conformance and access to IEEE-conformant facilities.

The input/output facilities have had three major new features: asynchronous input/output, stream input/output, and user-defined procedures for derived-type input/output (referred to as “defined input/output”). Additionally, the input/output specifiers have been regularised so that where they make sense: all specifiers that can be used on an **OPEN** statement can also be used on a **READ** or **WRITE** statement, and vice versa. Access to input/output error messages is provided by the new **IOMSG=** specifier, and processor-dependent constants for input/output (e.g. the unit number for the standard input file) are provided in a new intrinsic module.

Finally, there are a large number of miscellaneous improvements in almost every aspect of the language. Some of the more significant of these are the **IMPORT** statement (provides host association into interface blocks), the **VALUE** and

VOLATILE attributes, the ability to use all intrinsic functions in constant expressions, and extensions to the syntax of array and structure constructors.

65 Object-oriented programming

65.1 Type Extension

Type extension provides the first phase of object orientation: inheritance and polymorphic objects.

65.1.1 Extending Types [5.0]

Any derived type can be extended using the EXTENDS keyword, except for SEQUENCE types and BIND(C) types. (The latter types are “non-extensible”, as are intrinsic types, whereas all other derived types are “extensible”.) The extended type inherits all the components of the parent type and may add extra components.

For example:

```
TYPE point
  REAL x,y
END TYPE
TYPE,EXTENDS(point) :: point_3d
  REAL z
END TYPE
```

The type `point_3d` has `x`, `y` and `z` components. Additionally, it has a `point` component which refers to the inherited part; this “parent component” is “inheritance-associated” with the inherited components, so that the `point%x` component is identical to the `x` component et cetera.

However, when extending a type it is not required to add any new components; for example,

```
TYPE,EXTENDS(point) :: newpoint
END TYPE
```

defines a new type `newpoint` which has exactly the same components as `point` (plus the associated parent component). Similarly, it is no longer necessary for a type to contain any components:

```
TYPE empty_type
END TYPE
```

declares the extensible (but not extended) type `empty_type` which has no components at all.

65.1.2 Polymorphic Variables [5.0]

A polymorphic variable is a pointer, allocatable array or dummy argument that is declared using the CLASS keyword instead of the TYPE keyword. A `CLASS(typename)` variable can assume any type in the class of types consisting of `TYPE(typename)` and all extensions of *typename*.

For example:

```
REAL FUNCTION bearing(a)
  CLASS(point) a
  bearing = atan2(a%y,a%x)
END
```

The function `bearing` may be applied to a `TYPE(point)` object or to a `TYPE(point_3d)` object, or indeed to an object of any type that is an extension of `TYPE(point)`.

65.1.3 Type Selection [5.0]

The **SELECT TYPE** construct provides both a means of testing the dynamic type of a polymorphic variable and access to the extended components of that variable.

For example:

```
CLASS(t) x
...
SELECT TYPE(p=>x)
TYPE IS (t1)
!
! This section is executed only if X is exactly of TYPE(t1), not an
! extension thereof. P is TYPE(t1).
!
TYPE IS (t2)
!
! This section is executed only if X is exactly of TYPE(t2), not an
! extension thereof. P is TYPE(t2).
!
CLASS IS (t3)
!
! This section is executed if X is of TYPE(t3), or of some extension
! thereof, and if it is not caught by a more specific case. P is CLASS(t3).
!
END SELECT
```

Note that ‘**SELECT TYPE(x)**’ is short for ‘**SELECT TYPE(x=>x)**’.

65.1.4 Unlimited polymorphism [5.2]

A variable that is ‘**CLASS(*)**’ is an unlimited polymorphic variable. It has no type, but can assume any type including non-extensible types and intrinsic types (and kinds). Apart from allocation, deallocation and pointer assignment, to perform any operation on an unlimited polymorphic you first have to discover its type using **SELECT TYPE**. For example:

```
CLASS(*),POINTER :: x
CHARACTER(17),TARGET :: ch
x => ch
SELECT TYPE(x)
TYPE IS (COMPLEX(KIND=KIND(0d0)))
PRINT *,x+1
TYPE IS (CHARACTER(LEN=*))
PRINT *,LEN(x)
END SELECT
```

Note that in the case of **CHARACTER** the length must be specified as ‘*’ and is automatically assumed from whatever the polymorphic is associated with.

In the case of a non-extensible (i.e. **BIND(C)** or **SEQUENCE**) type, **SELECT TYPE** cannot be used to discover the type; instead, an unsafe pointer assignment is allowed, for example:

```
TYPE t
  SEQUENCE
  REAL x
END TYPE
CLASS(*),POINTER :: x
TYPE(t),POINTER :: y
```

```
...  
y => x ! Unsafe - the compiler cannot tell whether X is TYPE(t).
```

65.1.5 Ad hoc type comparison [5.3]

Two new intrinsic functions are provided for comparing the dynamic types of polymorphic objects. These are

```
EXTENDS_TYPE_OF(A,MOLD)  
SAME_TYPE_AS(A,B)
```

The arguments must be objects of extensible types (though they need not be polymorphic). `SAME_TYPE_AS` returns `.TRUE.` if and only if both `A` and `B` have the same dynamic type. `EXTENDS_TYPE_OF` returns `.TRUE.` if and only if the dynamic type of `A` is the same as, or an extension of, the dynamic type of `MOLD`. Note that if `MOLD` is an unallocated unlimited polymorphic (`CLASS(*)`), the result will be true regardless of the state of `A`.

The arguments are permitted to be unallocated or disassociated, but they are not permitted to be pointers with an undefined association status.

It is recommended that where possible these intrinsic functions be avoided, and that `SELECT TYPE` be used for type checking instead.

65.2 Typed allocation [5.1]

The `ALLOCATE` statement now accepts a *type-spec*; this can be used to specify the dynamic type (and type parameters, if any) of an allocation. The *type-spec* appears before the allocation list, and is separated from it by a double colon.

For example, if `T` is an extensible type and `ET` is an extension of `T`,

```
CLASS(t),POINTER :: a(:)  
ALLOCATE(et::a(100))
```

allocates `A` to have dynamic type `ET`. Note that the *type-spec* in an `ALLOCATE` statement omits the `TYPE` keyword for derived types, similarly to the `TYPE IS` and `CLASS IS` statements.

An unlimited polymorphic object can be allocated to be any type including intrinsic types: for example

```
CLASS(*),POINTER :: c,d  
ALLOCATE(DOUBLE PRECISION::c)  
READ *,n  
ALLOCATE(CHARACTER(LEN=n)::d)
```

allocates `C` to be double precision real, and `D` to be of type `CHARACTER` with length `N`.

Typed allocation is only useful for allocating polymorphic variables and `CHARACTER` variables with deferred length (`LEN=:`). For a non-polymorphic variable, the *type-spec* must specify the declared type and, if it is type `CHARACTER` but not deferred-length, to have the same character length. The character length must not be specified as an asterisk (`CHARACTER(LEN=*)`) unless the allocate-object is a dummy argument with an asterisk character length (and vice versa).

Finally, since there is only one *type-spec* it must be compatible with all the items in the allocation list.

65.3 Sourced allocation (cloning) [5.1]

The `ALLOCATE` statement now accepts the `SOURCE=` specifier. The dynamic type and value of the allocated entity is taken from the expression in the specifier. If the derived type has type parameters (q.v.), the value for any deferred type parameter is taken from the source expression, and the values for other type parameters must agree. This is not just applicable to derived types: if the entity being allocated is type `CHARACTER` with deferred length (`LEN=:`), the character length is taken from the source expression.

Only one entity can be allocated when the **SOURCE=** specifier is used. Note that when allocating an array the array shape is not taken from the source expression but must be specified in the usual way. If the source expression is an array, it must have the same shape as the array being allocated.

For example,

```
CLASS(*),POINTER :: a,b
...
ALLOCATE(a,SOURCE=b)
```

The allocated variable A will be a “clone” of B, whatever the current type of B happens to be.

65.4 Type-bound procedures [5.1]

Type-bound procedures provide a means of packaging operations on a type with the type itself, and also for dynamic dispatch to a procedure depending on the dynamic type of a polymorphic variable.

65.4.1 The type-bound procedure part

The type-bound procedure part of a type definition is separated from the components by the **CONTAINS** statement. The default accessibility of type-bound procedures is public even if the components are private; this may be changed by using the **PRIVATE** statement after the **CONTAINS**.

65.4.2 Specific type-bound procedures

The syntax of a specific, non-deferred, type-bound procedure declaration is:

```
PROCEDURE [[binding-attr-list::] binding-name [=>procedure-name]
```

The name of the type-bound procedure is *binding-name*, and the name of the actual procedure which implements it is *procedure-name*. If the optional **=>*procedure-name*** is omitted, the actual procedure has the same name as the binding.

A type-bound procedure is invoked via an object of the type, e.g.

```
CALL variable(i)%tbp(arguments)
```

Normally, the invoking variable is passed as an extra argument, the “passed-object dummy argument”; by default this is the first dummy argument of the actual procedure and so the first argument in the argument list becomes the second argument, etc. The passed-object dummy argument may be changed by declaring the type-bound procedure with the **PASS(*argument-name*)** attribute, in which case the variable is passed as the named argument. The **PASS** attribute may also be used to confirm the default (as the first argument), and the **NOPASS** attribute prevents passing the object as an argument at all. The passed-object dummy argument must be a polymorphic scalar variable of that type, e.g. **CLASS(t) self**.

When a type is extended, the new type either inherits or **overrides** each type-bound procedure of the old type. An overriding procedure must be compatible with the old procedure; in particular, each dummy argument must have the same type except for the passed-object dummy argument which must have the new type. A type-bound procedure that is declared to be **NON-OVERRIDABLE** cannot be overridden during type extension.

When a type-bound procedure is invoked, it is the dynamic type of the variable which determines which actual procedure to call.

The other attributes that a type-bound procedure may have are **PUBLIC**, **PRIVATE**, and **DEFERRED** (the latter only for abstract types, which are described later).

65.4.3 Generic type-bound procedures

A generic type-bound procedure is a set of specific type-bound procedures, in the same way that an ordinary generic procedure is a set of specific ordinary procedures. It is declared with the **GENERIC** statement, e.g.

```
GENERIC :: generic_name => specific_name_1, specific_name_2, specific_name_3
```

Generic type-bound procedures may also be operators or assignment, e.g.

```
GENERIC :: OPERATOR(+) => add_t_t, add_t_r, add_r_t
```

Such type-bound generic operators cannot have the **NOPASS** attribute; the dynamic type of the passed-object dummy argument determines which actual procedure is called.

When a type is extended, the new type inherits all the generic type-bound procedures without exception, and the new type may extend the generic with additional specific procedures. To override procedures in the generic, simply override the specific type-bound procedure. For example, in

```
TYPE mycomplex
...
CONTAINS
  PROCEDURE :: myc_plus_r => myc1_plus_r
  PROCEDURE,PASS(B) :: r_plus_myc => r_plus_myc1
  GENERIC :: OPERATOR(+) => myc_plus_r, r_plus_myc
END TYPE
...
TYPE,EXTENDS(mycomplex) :: mycomplex_2
...
CONTAINS
  PROCEDURE :: myc_plus_r => myc2_plus_r
  PROCEDURE,PASS(B) :: r_plus_myc => r_plus_myc2
END TYPE
```

the type **mycomplex_2** inherits the generic operator '+'; invoking the generic (+) invokes the specific type-bound procedure, which for entities of type **mycomplex_2** will invoke the overriding actual procedure (**myc2_plus_r** or **r_plus_myc2**).

65.5 Abstract derived types [5.1]

An extensible derived type can be declared to be **ABSTRACT**, e.g.

```
TYPE, ABSTRACT :: mytype
```

An abstract type cannot be instantiated; i.e. it is not allowed to declare a non-polymorphic variable of abstract type, and a polymorphic variable of abstract type must be allocated to be a non-abstract extension of the type.

Abstract type may contain **DEFERRED** type-bound procedures, e.g.

```
...
CONTAINS
  PROCEDURE(interface_name),DEFERRED :: tbpname
```

No binding ("**=> name**") is allowed or implied by a deferred procedure binding. The **interface_name** must be the name of an abstract interface or a procedure with an explicit interface, and defines the interface of the deferred type-bound procedure.

When extending an abstract type, the extended type must also be abstract unless it overrides all of the deferred type-bound procedures with normal bindings.

65.6 Object-bound procedures [5.2]

These are procedure pointer components, and act similarly to type-bound procedures except that the binding is per-object not per-type. The syntax of a procedure pointer component declaration is:

```
PROCEDURE( [proc-interface] ) , proc-component-attr-spec-list :: proc-decl-list
```

where

- each *proc-component-attr-spec* is one of `NOPASS`, `PASS`, `PASS(arg-name)`, `POINTER`, `PRIVATE` or `PUBLIC`, and
- each *proc-decl* is a component name optionally followed by default initialisation to null (`'=> NULL()'`).

The `POINTER` attribute is required.

Note that object-bound procedures have a passed-object dummy argument just like type-bound procedures; if this is not wanted, the `NOPASS` attribute must be used (and this is required if the interface is implicit, i.e. when `proc-interface` is missing or is a type specification).

The following example demonstrates using a list of subroutines with no arguments.

```
TYPE action_list
  PROCEDURE(),NOPASS,POINTER :: action => NULL()
  TYPE(action_list),POINTER :: next
END TYPE
TYPE(t),TARGET :: top
TYPE(t),POINTER :: p
EXTERNAL sub1,sub2
top%action = sub1
ALLOCATE(top%next)
top%next%action = sub2
...
p => top
DO WHILE (ASSOCIATED(p))
  IF (ASSOCIATED(p%action)) CALL p%action
  p => p%next
END DO
```

66 ALLOCATABLE extensions

In Fortran 2003 the `ALLOCATABLE` attribute is permitted not just on local variables but also on components, dummy variables, and function results. These are the same as described in the ISO Technical Report ISO/IEC TR 15581:1999.

Also, the `MOVE_ALLOC` intrinsic subroutine has been added, as well as automatic reallocation on assignment.

66.1 Allocatable Dummy Arrays [4.x]

A dummy argument can be declared to be an allocatable array, e.g.

```
SUBROUTINE s(dum)
  REAL,ALLOCATABLE :: dum(:, :)
  ...
END SUBROUTINE
```

Having an allocatable dummy argument means that there must be an explicit interface for any reference: i.e. if the procedure is not an internal or module procedure there must be an accessible interface block in any routine which references that procedure.

Any actual argument that is passed to an allocatable dummy array must itself be an allocatable array; it must also have the same type, kind type parameters, and rank. For example:

```
REAL,ALLOCATABLE :: x(:, :)
CALL s(x)
```

The actual argument need not be allocated before calling the procedure, which may itself allocate or deallocate the argument. For example:

```
PROGRAM example2
  REAL,ALLOCATABLE :: x(:, :)
  OPEN(88,FILE='myfile',FORM='unformatted')
  CALL read_matrix(x,88)
  !
  ... process x in some way
  !
  REWIND(88)
  CALL write_and_delete_matrix(x,88)
END
!
MODULE module
CONTAINS
  !
  ! This procedure reads the size and contents of an array from an
  ! unformatted unit.
  !
  SUBROUTINE read_matrix(variable,unit)
    REAL,ALLOCATABLE,INTENT(OUT) :: variable(:, :)
    INTEGER,INTENT(IN) :: unit
    INTEGER dim1,dim2
    READ(unit) dim1,dim2
    ALLOCATE(variable(dim1,dim2))
    READ(unit) variable
    CLOSE(unit)
  END SUBROUTINE
  !
  ! This procedure writes the size and contents of an array to an
  ! unformatted unit, and then deallocates the array.
  !
  SUBROUTINE write_and_delete_matrix(variable,unit)
    REAL,ALLOCATABLE,INTENT(INOUT) :: variable(:, :)
    INTEGER,INTENT(IN) :: unit
    WRITE(unit) SIZE(variable,1),SIZE(variable,2)
    WRITE(unit) variable
    DEALLOCATE(variable)
  END SUBROUTINE
END
```

66.2 Allocatable Function Results [4.x]

The result of a function can be declared to be an allocatable array, e.g.

```
FUNCTION af() RESULT(res)
  REAL,ALLOCATABLE :: res
```

On invoking the function, the result variable will be unallocated. It must be allocated before returning from the function. For example:

```
!  
! The result of this function is the original argument with adjacent  
! duplicate entries deleted (so if it was sorted, each element is unique).  
!  
FUNCTION compress(array)  
  INTEGER,ALLOCATABLE :: compress(:)  
  INTEGER,INTENT(IN) :: array(:)  
  IF (SIZE(array,1)==0) THEN  
    ALLOCATE(compress(0))  
  ELSE  
    N = 1  
    DO I=2,SIZE(array,1)  
      IF (array(I)/=array(I-1)) N = N + 1  
    END DO  
    ALLOCATE(compress(N))  
    N = 1  
    compress(1) = array(1)  
    DO I=2,SIZE(array,1)  
      IF (array(I)/=compress(N)) THEN  
        N = N + 1  
        compress(N) = array(I)  
      END IF  
    END DO  
  END IF  
END
```

The result of an allocatable array is automatically deallocated after it has been used.

66.3 Allocatable Structure Components [4.x]

A structure component can be declared to be allocatable, e.g.

```
MODULE matrix_example  
  TYPE MATRIX  
    REAL,ALLOCATABLE :: value(:, :)  
  END TYPE  
END MODULE
```

An allocatable array component is initially not allocated, just like allocatable array variables. On exit from a procedure containing variables with allocatable components, all the allocatable components are automatically deallocated. This is in contradistinction to pointer components, which are not automatically deallocated. For example:

```
SUBROUTINE sub(n,m)  
  USE matrix_example  
  TYPE(matrix) a,b,c  
  !  
  ! a%value, b%value and c%value are all unallocated at this point.  
  !  
  ALLOCATE(a%value(n,m),b%value(n,m))  
  !  
  ... do some computations, then  
  !  
  RETURN  
  !  
  ! Returning from the procedure automatically deallocates a%value, b%value,  
  ! and c%value (if they are allocated).  
  !  
END
```

Deallocating a variable that has an allocatable array component deallocates the component first; this happens recursively so that all ALLOCATABLE subobjects are deallocated with no memory leaks.

Any allocated allocatable components of a function result are automatically deallocated after the result has been used.

```
PROGRAM deallocation_example
  TYPE inner
    REAL,ALLOCATABLE :: ival(:)
  END TYPE
  TYPE outer
    TYPE(inner),ALLOCATABLE :: ovalue(:)
  END TYPE
  TYPE(outer) x
  !
  ! At this point, x%ovalue is unallocated
  !
  ALLOCATE(x%ovalue(10))
  !
  ! At this point, x%ovalue(i)%ival are unallocated, i=1,10
  !
  ALLOCATE(x%ovalue(2)%ival(1000),x%ovalue(5)%ival(9999))
  !
  ! Only x%ovalue(2)%ival and x%ovalue(5)%ival are allocated
  !
  DEALLOCATE(x%ovalue)
  !
  ! This has automatically deallocated x%ovalue(2)%ival and x%ovalue(5)%ival
  !
END
```

In a structure constructor for such a type, the expression corresponding to an allocatable array component can be

- the NULL() intrinsic, indicating an unallocated array,
- an allocatable array which may be allocated or unallocated, or
- any other array expression, indicating an allocated array.

```
SUBROUTINE constructor_example
  USE matrix_example
  TYPE(matrix) a,b,c
  REAL :: array(10,10) = 1
  REAL,ALLOCATABLE :: alloc_array(:,:)
  a = matrix(NULL())
  !
  ! At this point, a%value is unallocated
  !
  b = matrix(array*2)
  !
  ! Now, b%value is a (10,10) array with each element equal to 2.
  !
  c = matrix(alloc_array)
  !
  ! Now, c%value is unallocated (because alloc_array was unallocated).
  !
END
```

Intrinsic assignment of such types does a “deep copy” of the allocatable array components; it is as if the allocatable array component were deallocated (if necessary), then if the component in the expression was allocated, the variable’s component is allocated to the right size and the value copied.

```
SUBROUTINE assignment_example
  USE matrix_example
  TYPE(matrix) a,b
  !
  ! First we establish a value for a
  !
  ALLOCATE(a%value(10,20))
  a%value(3,:) = 30
  !
  ! And a value for b
  !
  ALLOCATE(b%value(1,1))
  b%value = 0
  !
  ! Now the assignment
  !
  b = a
  !
  ! The old contents of b%value have been deallocated, and b%value now has
  ! the same size and contents as a%value.
  !
END
```

66.4 Allocatable Component Example

This example shows the definition and use of a simple module that provides polynomial arithmetic. To do this it makes use of intrinsic assignment for allocatable components, the automatically provided structure constructors and defines the addition (+) operator. A more complete version of this module would provide other operators such as multiplication.

```
!
! Module providing a single-precision polynomial arithmetic facility
!
MODULE real_poly_module
  !
  ! Define the polynomial type with its constructor.
  ! We will use the convention of storing the coefficients in the normal
  ! order of highest degree first, thus in an N-degree polynomial, COEFF(1)
  ! is the coefficient of X**N, COEFF(N) is the coefficient of X**1, and
  ! COEFF(N+1) is the scalar.
  !
  TYPE,PUBLIC :: real_poly
    REAL,ALLOCATABLE :: coeff(:)
  END TYPE
  !
  PUBLIC OPERATOR(+)
  INTERFACE OPERATOR(+)
    MODULE PROCEDURE rp_add_rp,rp_add_r,r_add_rp
  END INTERFACE
  !
CONTAINS
  TYPE(real_poly) FUNCTION rp_add_r(poly,real)
    TYPE(real_poly),INTENT(IN) :: poly
    REAL,INTENT(IN) :: real
    INTEGER isize
    IF (.NOT.ALLOCATED(poly%coeff)) STOP 'Undefined polynomial value in +'
    isize = SIZE(poly%coeff,1)
    rp_add_r%coeff(isize) = poly%coeff(isize) + real
  END FUNCTION
```

```

END FUNCTION
TYPE(real_poly) FUNCTION r_add_rp(real,poly)
  TYPE(real_poly),INTENT(IN) :: poly
  REAL,INTENT(IN) :: real
  r_add_rp = rp_add_r(poly,real)
END FUNCTION
TYPE(real_poly) FUNCTION rp_add_rp(poly1,poly2)
  TYPE(real_poly),INTENT(IN) :: poly1,poly2
  INTEGER I,N,N1,N2
  IF (.NOT.ALLOCATED(poly1%coeff).OR..NOT.ALLOCATED(poly2%coeff)) &
    STOP 'Undefined polynomial value in +'
  ! Set N1 and N2 to the degrees of the input polynomials
  N1 = SIZE(poly1%coeff) - 1
  N2 = SIZE(poly2%coeff) - 1
  ! The result polynomial is of degree N
  N = MAX(N1,N2)
  ALLOCATE(rp_add_rp%coeff(N+1))
  DO I=0,MIN(N1,N2)
    rp_add_rp%coeff(N-I+1) = poly1%coeff(N1-I+1) + poly2%coeff(N2-I+1)
  END DO
  ! At most one of the next two DO loops is ever executed
  DO I=N1+1,N
    rp_add_rp%coeff(N-I+1) = poly2%coeff(N2-I+1)
  END DO
  DO I=N2+1,N
    rp_add_rp%coeff(N-I+1) = poly1%coeff(N1-I+1)
  END DO
END FUNCTION
END MODULE
!
! Sample program
!
PROGRAM example
  USE real_poly_module
  TYPE(real_poly) p,q,r
  p = real_poly((/1.0,2.0,4.0/)) ! x**2 + 2x + 4
  q = real_poly((/1.0,-5.5/))    ! x - 5.5
  r = p + q                     ! x**2 + 3x - 1.5
  print 1,'The coefficients of the answer are:',r%coeff
1 format(1x,A,3F8.2)
END

```

When executed, the above program prints:

```
The coefficients of the answer are:    1.00    3.00   -1.50
```

66.5 The MOVE_ALLOC intrinsic subroutine [5.2]

This subroutine moves an allocation from one allocatable variable to another. This can be used to expand an allocatable array with only one copy operation, and allows full control over where in the new array the values should go. For example:

```

REAL,ALLOCATABLE :: a(:),tmp(:)
...
ALLOCATE(a(n))
...
! Here we want to double the size of A, without losing any of the values
! that are already stored in it.

```

```
ALLOCATE(tmp(size(a)*2))
tmp(1:size(a)) = a
CALL move_alloc(from=tmp,to=a)
! TMP is now deallocated, and A has the new size and values.
```

To have the values end up somewhere different, just change the assignment statement, for example to move them all to the end:

```
tmp(size(a)+1:size(a)*2) = a
```

66.6 Allocatable scalars [5.2]

The `ALLOCATABLE` attribute may now be applied to scalar variables and components, not just arrays. This is most useful in conjunction with polymorphism (`CLASS`) and/or deferred type parameters (e.g. `CHARACTER(:)`); for more details see the “Typed allocation”, “Sourced allocation” and “Automatic reallocation” sections.

66.7 Automatic reallocation [5.2]

If, in an assignment to a whole allocatable array, the expression being assigned is an array of a different size or shape, the allocatable array is reallocated to have the correct shape (in Fortran 95 this assignment would have been an error). For example:

```
ALLOCATE(a(10))
...
a = (/ (i,i=1,100) /)
! A is now size 100
```

Similarly, if an allocatable variable has a deferred type parameter (these are described in a later section), and is either unallocated or has a value different from that of the expression, the allocatable variable is reallocated to have the same value for that type parameter. This allows for true varying-length character variables:

```
CHARACTER(:),ALLOCATABLE :: name
...
name = 'John Smith'
! LEN(name) is now 10, whatever it was before.
name = '?'
! LEN(name) is now 1.
```

Note that since a subobject of an allocatable object is not itself allocatable, this automatic reallocation can be suppressed by using substrings (for characters) or array sections (for arrays), e.g.

```
name(:) = '?'           ! Normal assignment with truncation/padding.
a(:) = (/ (i,i=1,100) /) ! Asserts that A is already of size 100.
```

67 Other data-oriented enhancements

67.1 Parameterised derived types [n/a]

Derived types may now have type parameters. Like those of the intrinsic types, they come in two flavours: “kind”-like ones which must be known at compile time (called “kind” type parameters), and ones like character length which may vary at runtime (called “length” type parameters).

67.1.1 Basic Syntax and Semantics

A derived type which has type parameters must list them in the type definition, give them a type, and specify whether they are “kind” or “length” parameters. For example,

```
TYPE real_matrix(kind,n,m)
  INTEGER,KIND :: kind
  INTEGER(int64),LENGTH :: n,m
```

All type parameters must be explicitly specified to be of type `INTEGER`, but the kind of integer may vary. Type parameters are always scalar, never arrays. Within the type definition, “kind” type parameters may be used in constant expressions, and any type parameter may be used in a specification expression (viz array bound, character length, or “length” type parameter value). For example, the rest of the above type definition might look like:

```
    REAL(kind) value(n,m)
END TYPE real_matrix
```

When declaring entities of such a derived type, the type parameters must be given after the name. For example,

```
TYPE(real_matrix(KIND(0d0),100,200)) :: my_real_matrix_variable
```

Similarly, the type parameters must be given when constructing values of such a type; for example,

```
my_real_matrix_variable = &
  real_matrix(kind(0d0),100,200)((/ (i*1.0d0,i=1,20000) /))
```

To examine the value of a derived type parameter from outside the type definition, the same notation is used as for component accesses, e.g.

```
print *, 'Columns =', my_real_matrix_variable%m
```

Thus type parameter names are in the same class as component names and type-bound procedure names. However, a type parameter cannot be changed by using its specifier on the left-hand-side of an assignment. Furthermore, the intrinsic type parameters may also be examined using this technique, for example:

```
REAL :: array(:, :)
CHARACTER(*), INTENT(IN) :: ch
PRINT *, array%kind, ch%len
```

prints the same values as for `KIND(array)` and `LEN(ch)`. Note that a kind parameter enquiry is always scalar, even if the object is an array.

A derived type parameter does not actually have to be used at all within the type definition, and a kind type parameter might only be used within specification expressions. For example,

```
TYPE fixed_byte(n)
  INTEGER,KIND :: n
  INTEGER(1) :: value(n)
END TYPE
TYPE numbered_object(object_number)
  INTEGER,LENGTH :: object_number
END TYPE
```

Even though the `fixed_byte` parameter `n` is not used in a constant expression, a constant value must always be specified for it because it has been declared to be a “kind” type parameter. Similarly, even though `object_number` has not been used at all, a value must always be specified for it. This is not quite as useless as it might seem: each `numbered_object` has a single value for `object_number` even if the `numbered_object` is an array.

67.1.2 More Semantics

A derived type with type parameters can have default values for one or more of them; in this case the parameters with default values may be omitted from the type specifiers. For example,

```
TYPE char_with_maxlen(maxlen,kind)
  INTEGER,LENGTH :: maxlen = 254
  INTEGER,KIND   :: kind = SELECTED_CHAR_KIND('ascii')
  INTEGER        :: len = 0
  CHARACTER(len=maxlen,kind=kind) :: value
END TYPE
...
TYPE(char_with_maxlen) temp
TYPE(char_with_maxlen(80)) card(1000)
TYPE(char_with_maxlen(kind=SELECTED_CHAR_KIND('iso 10646')))) ucs4_temp
```

Note that although kind type parameters can be used in constant expressions and thus in default initialisation, components that are variable-sized (because they depend on length type parameters) cannot be default-initialised at all. Thus `value` in the example above cannot be default-initialised.

Further note that unlike intrinsic types, there are no automatic conversions for derived type assignment with different type parameter values, thus given the above declarations,

```
card(1) = card(2) ! This is ok, maxlen==80 for both sides.
temp = card       ! This is not allowed - maxlen 254 vs. maxlen 80.
```

67.1.3 Assumed type parameters

Assumed type parameters for derived types work similarly to character length, except that they are only allowed for dummy arguments (not for named constants). For example, the following subroutine works on any `char_with_maxlen` variable.

```
SUBROUTINE stars(x)
  TYPE(char_with_maxlen(*)) x
  x%value = REPEAT('*',x%maxlen)
END SUBROUTINE
```

67.1.4 Deferred type parameters

Deferred type parameters are completely new to Fortran 2003; these are available both for `CHARACTER` and for parameterised derived types, and work similarly to deferred array bounds. A variable with a deferred type parameter must have the `ALLOCATABLE` or `POINTER` attribute. The value of a deferred type parameter for an allocatable variable is that determined by allocation (either by a typed allocation, or by an intrinsic assignment with automatic reallocation). For a pointer, the value of a deferred type parameter is the value of the type parameter of its target. For example, using the type `real_matrix` defined above,

```
TYPE(real_matrix(KIND(0.0),100,200)),TARGET :: x
TYPE(real_matrix(KIND(0.0),:,:)),POINTER :: y, z
ALLOCATE(real_matrix(KIND(0.0),33,44) :: y) ! Typed allocation.
z => x ! Assumes from the target.
PRINT *,y%n,z%n ! Prints 33 and 100.
```

Note that it is not allowed to reference the value of a deferred type parameter of an unallocated allocatable or of a pointer that is not associated with a target.

If a dummy argument is allocatable or a pointer, the actual argument must have deferred exactly the same type parameters as the dummy. For example,

```
SUBROUTINE sub(rm_dble_ptr)
  TYPE(real_matrix(KIND(0d0),*,:)),POINTER :: rm_dble_ptr
  ...
  TYPE(real_matrix(KIND(0d0),100,200)),POINTER :: x
  TYPE(real_matrix(KIND(0d0),100,:)),POINTER :: y
  TYPE(real_matrix(KIND(0d0),:,:)),POINTER :: z
  CALL sub(x) ! Invalid - X%M is not deferred (but must be).
  CALL sub(y) ! This is ok.
  CALL sub(z) ! Invalid - X%N is deferred (but must not be).
```

67.2 Finalisation [5.3]

An extensible derived type can have “final subroutines” associated with it; these subroutines are automatically called whenever an object of the type is about to be destroyed, whether by deallocation, procedure return, being on the left-hand-side of an intrinsic assignment, or being passed to an `INTENT(OUT)` dummy argument.

A final subroutine of a type must be a subroutine with exactly one argument, which must be an ordinary dummy variable of that type (and must not be `INTENT(OUT)`). It may be scalar or an array, and when an object of that type is destroyed the final subroutine whose argument has the same rank as the object is called. The final subroutine may be elemental, in which case it will handle any rank of object that has no other subroutine handling it. Note that if there is no final subroutine for the rank of an object, no subroutine will be called.

Final subroutines are declared in the type definition after the `CONTAINS` statement, like type-bound procedures. They are declared by a `FINAL` statement, which has the syntax

```
FINAL [ :: ] name [ , name ]...
```

where each *name* is a subroutine that satisfies the above rules.

A simple type with a final subroutine is as follows.

```
TYPE flexible_real_vector
  LOGICAL :: value_was_allocated = .FALSE.
  REAL,POINTER :: value(:) => NULL()
CONTAINS
  FINAL destroy_frv
END TYPE
...
ELEMENTAL SUBROUTINE destroy_frv(x)
  TYPE(flexible_real_vector),INTENT(INOUT) :: x
  IF (x%value_was_allocated) DEALLOCATE(x%value)
END SUBROUTINE
```

If an object being destroyed has finalisable components, any final subroutine for the object-as-a-whole will be called before finalising any components. If the object is an array, each component will be finalised separately (and any final subroutine called will be the one for the rank of the component, not the rank of the object).

For example, in

```
TYPE many_vectors
  TYPE(flexible_real_vector) scalar
  TYPE(flexible_real_vector) array(2,3)
CONTAINS
  FINAL :: destroy_many_vectors_1
END TYPE
...
SUBROUTINE destroy_many_vectors_1(array1)
  TYPE(many_vectors) array1(:)
  PRINT *, 'Destroying a', SIZE(array1), 'element array of many vectors'
```

```
END SUBROUTINE
...
TYPE(many_vector) mv_object(3)
```

when `mv_object` is destroyed, firstly ‘`destroy_many_vectors_1`’ will be called with `mv_object` as its argument; this will print

```
Destroying a 3 element array of many vectors
```

Secondly, for each element of the array, both the `scalar` and `array` components will be finalised by calling `destroy_frv` on each of them. These may be done in any order (or, since they are elemental, potentially in parallel).

Note that final subroutines are not inherited through type extension; instead, when an object of extended type is destroyed, first any final subroutine it has will be called, then any final subroutine of the parent type will be called on the parent component, and so on.

67.3 The PROTECTED attribute [5.0]

The `PROTECTED` attribute may be specified by the `PROTECTED` statement or with the `PROTECTED` keyword in a type declaration statement. It protects a module variable against modification from outside the module.

67.3.1 Syntax

The syntax of the `PROTECTED` statement is:

```
PROTECTED [ :: ] name [ , name ] ...
```

The `PROTECTED` attribute may only be specified for a variable in a module.

67.3.2 Semantics

Variables with the `PROTECTED` attribute may only be modified within the defining module. Outside of that module they are not allowed to appear in a variable definition context (e.g. on the left-hand-side of an assignment statement), similar to `INTENT(IN)` dummy arguments.

This allows the module writer to make the values of some variables generally available without relinquishing control over their modification.

67.3.3 Example

```
MODULE temperature_module
  REAL,PROTECTED :: temperature_c = 0, temperature_f = 32
CONTAINS
  SUBROUTINE set_temperature_c(new_value_c)
    REAL,INTENT(IN) :: new_value_c
    temperature_c = new_value_c
    temperature_f = temperature_c*(9.0/5.0) + 32
  END SUBROUTINE
  SUBROUTINE set_temperature_f(new_value_f)
    REAL,INTENT(IN) :: new_value_f
    temperature_f = new_value_f
    temperature_c = (temperature_f - 32)*(5.0/9.0)
  END SUBROUTINE
END
```

The `PROTECTED` attribute allows users of `temperature_module` to read the temperature in either Farenheit or Celsius, but the variables can only be changed via the provided subroutines which ensure that both values agree.

67.4 Pointer enhancements

67.4.1 INTENT for pointers [5.1]

A `POINTER` dummy argument may now have the `INTENT` attribute. This attribute applies to the pointer association status, not to the target of the pointer.

An `INTENT(IN)` pointer can be assigned to, but cannot be pointer-assigned, nullified, allocated or deallocated. An `INTENT(OUT)` pointer receives an undefined association status on entry to the procedure. An `INTENT(INOUT)` pointer has no restrictions on its use, but the actual argument must be a pointer variable, not a pointer function reference.

67.4.2 Pointer bounds specification [5.2]

The bounds of a pointer can be changed (from the default) in a pointer assignment by including them on the left-hand-side. For example,

```
REAL, TARGET :: x(-100:100, -10:10)
REAL, POINTER :: p(:, :)
p(1:, 1:) => x
```

The upper bound is formed by adding the extent (minus 1) to the lower bound, so in the above example, the bounds of `P` will be `1:201, 1:21`. Note that when setting the lower bound of any rank in a pointer assignment, the values must be explicitly specified (there is no default of 1 like there is in array declarators) and they must be specified for all dimensions of the pointer.

67.4.3 Rank-remapping Pointer Assignment [5.0]

This feature allows a multi-dimensional pointer to point to a single-dimensional object. For example:

```
REAL, POINTER :: diagonal(:), matrix(:, :), base(:)
...
ALLOCATE(base(n*n))
matrix(1:n, 1:n) => base
diagonal => base(::n+1)
!
! DIAGONAL now points to the diagonal elements of MATRIX.
!
```

Note that when rank-remapping, the values for both the lower and upper bounds must be explicitly specified for all dimensions, there are no defaults.

67.5 Individual component accessibility [5.1]

It is now possible to set the accessibility of individual components in a derived type. For example,

```
TYPE t
  LOGICAL, PUBLIC :: flag
  INTEGER, PRIVATE :: state
END TYPE
```

The structure constructor for the type is not usable from outside the defining module if there is any private component that is not inherited, allocatable or default-initialised (see Structure constructor syntax enhancements).

67.6 Public entities of private type [5.1]

It is now possible to export entities (named constants, variables, procedures) from a module even if they have private type or (for procedures) have arguments of private type. For example,

```
MODULE m
  TYPE, PRIVATE :: hidden_type
    CHARACTER(6) :: code
  END TYPE
  TYPE(hidden_type), PUBLIC, PARAMETER :: code_green = hidden_type('green')
  TYPE(hidden_type), PUBLIC, PARAMETER :: code_yellow = hidden_type('yellow')
  TYPE(hidden_type), PUBLIC, PARAMETER :: code_red = hidden_type('red')
END
```

68 C interoperability [mostly 5.1]

68.1 The ISO_C_BINDING module

The intrinsic module ISO_C_BINDING contains

- for each C type (e.g. `float`), a named constant for use as a `KIND` parameter for the corresponding Fortran type,
- types `C_PTR` and `C_FUNPTR` for interoperating with C object pointers and function pointers,
- procedures for manipulating Fortran and C pointers.

68.1.1 The kind parameters

The kind parameter names are for using with the corresponding Fortran types; for example, `INTEGER` for integral types and `REAL` for floating-point types. This is shown in the table below. Note that only `c_int` is guaranteed to be available; if there is no compatible type the value will be negative.

C type	Fortran type and kind
<code>_Bool</code>	<code>LOGICAL(c_bool)</code>
<code>char</code>	<code>CHARACTER(c_char)</code> — For characters as text.
<code>double</code>	<code>REAL(c_double)</code>
<code>double _Complex</code>	<code>COMPLEX(c_double_complex)</code> or <code>COMPLEX(c_double)</code>
<code>float</code>	<code>REAL(c_float)</code>
<code>float _Complex</code>	<code>COMPLEX(c_float_complex)</code> or <code>COMPLEX(c_float)</code>
<code>int</code>	<code>INTEGER(c_int)</code>
<code>int16_t</code>	<code>INTEGER(c_int16_t)</code>
<code>int32_t</code>	<code>INTEGER(c_int32_t)</code>
<code>int64_t</code>	<code>INTEGER(c_int64_t)</code>
<code>int8_t</code>	<code>INTEGER(c_int8_t)</code>
<code>int_fast16_t</code>	<code>INTEGER(c_int_fast16_t)</code>
<code>int_fast32_t</code>	<code>INTEGER(c_int_fast32_t)</code>
<code>int_fast64_t</code>	<code>INTEGER(c_int_fast64_t)</code>
<code>int_fast8_t</code>	<code>INTEGER(c_int_fast8_t)</code>
<code>int_least16_t</code>	<code>INTEGER(c_int_least16_t)</code>
<code>int_least32_t</code>	<code>INTEGER(c_int_least32_t)</code>
<code>int_least64_t</code>	<code>INTEGER(c_int_least64_t)</code>
<code>int_least8_t</code>	<code>INTEGER(c_int_least8_t)</code>
<code>intmax_t</code>	<code>INTEGER(c_intmax_t)</code>
<code>intptr_t</code>	<code>INTEGER(c_intptr_t)</code>
<code>long</code>	<code>INTEGER(c_long)</code>
<code>long double</code>	<code>REAL(c_long_double)</code>
<code>long double _Complex</code>	<code>COMPLEX(c_long_double_complex)</code> or <code>COMPLEX(c_long_double)</code>
<code>long long</code>	<code>INTEGER(c_long_long)</code>
<code>short</code>	<code>INTEGER(c_short)</code>
<code>signed char</code>	<code>INTEGER(c_signed_char)</code> — For characters as integers.
<code>size_t</code>	<code>INTEGER(c_size_t)</code>

68.1.2 Using C_PTR and C_FUNPTR

These are derived type names, so you use them as `Type(c_ptr)` and `Type(c_funptr)`. `Type(c_ptr)` is essentially equivalent to the C `void *`; i.e. it can contain any object pointer. `Type(c_funptr)` does the same thing for function pointers.

For C arguments like `'int *'`, you don't need to use `Type(c_ptr)`, you can just use a normal dummy argument (in this case of type `Integer(c_int)`) without the `VALUE` attribute. However, for more complicated pointer arguments such as pointer to pointer, or for variables or components that are pointers, you need to use `Type(c_ptr)`.

Null pointer constants of both `Type(c_ptr)` and `Type(c_funptr)` are provided: these are named `C_NULL_PTR` and `C_NULL_FUNPTR` respectively.

To create a `Type(c_ptr)` value, the function `C_LOC(X)` is used on a Fortran object `X` (and `X` must have the `TARGET` attribute). Furthermore, the Fortran object cannot be polymorphic, a zero-sized array, an assumed-size array, or an array pointer. To create a `Type(c_funptr)` value, the function `C_FUNLOC` is used on a procedure; this procedure must have the `BIND(C)` attribute.

To test whether a `Type(c_ptr)` or `Type(c_funptr)` is null, the `C_ASSOCIATED(C_PTR_1)` function can be used; it returns `.TRUE.` if and only if `C_PTR_1` is not null. Two `Type(c_ptr)` or two `Type(c_funptr)` values can be compared using `C_ASSOCIATED(C_PTR_1, C_PTR_2)` function; it returns `.TRUE.` if and only if `C_PTR_1` contains the same C address as `C_PTR_2`.

The subroutine `C_F_POINTER(CPTR, FPTR)` converts the `TYPE(C_PTR)` value `CPTR` to the scalar Fortran pointer `FPTR`; the latter can have any type (including non-interoperable types) but must not be polymorphic. The subroutine `C_F_POINTER(CPTR, FPTR, SHAPE)` converts a `TYPE(C_PTR)` value into the Fortran array pointer `FPTR`, where `SHAPE` is an integer array of rank 1, with the same number of elements as the rank of `FPTR`; the lower bounds of the resultant `FPTR` will all be 1.

The subroutine `C_F_PROCPTR(CPTR, FPTR)` is provided. This converts the `TYPE(C_FUNPTR)` `CPTR` to the Fortran procedure pointer `FPTR`.

Note that in all the conversion cases it is up to the programmer to use the correct type and other information.

68.2 BIND(C) types

Derived types corresponding to C `struct` types can be created by giving the type the `BIND(C)` attribute, e.g.

```
TYPE,BIND(C) :: mytype
```

The components of a `BIND(C)` type must have types corresponding to C types, and cannot be pointers or allocatables. Furthermore, a `BIND(C)` type cannot be a `SEQUENCE` type (it already acts like a `SEQUENCE` type), cannot have type-bound procedures, cannot have final procedures, and cannot be extended.

68.3 BIND(C) variables

Access to C global variables is provided by giving the Fortran variable the `BIND(C)` attribute. Such a variable can only be declared in a module, and cannot be in a `COMMON` block. By default, the C name of the variable is the Fortran name converted to all lowercase characters; a different name may be specified with the `NAME=` clause, e.g.

```
INTEGER,BIND(C,NAME="StrangelyCapiTalisedCName") :: x
```

Within Fortran code, the variable is referred to by its Fortran name, not its C name.

68.4 BIND(C) procedures

A Fortran procedure that can be called from C can be defined using the `BIND(C)` attribute on the procedure heading. By default its C name is the Fortran name converted to lowercase; a different name may be specified with the `NAME=` clause. For example

```
SUBROUTINE sub() BIND(C,NAME='Sub')  
...
```

Again, the C name is for use only from C, the Fortran name is used from Fortran. If the C name is all blanks (or a zero-length string), there is no C name. Such a procedure can still be called from C via a procedure pointer (i.e. by assigning it to a `TYPE(C_FUNPTR)` variable).

All `BIND(C)` procedures must have an explicit interface.

A `BIND(C)` procedure may be a subroutine or a scalar function with a type corresponding to a C type. Each dummy argument must be a variable whose type corresponds to a C type, and cannot be allocatable, assumed-shape, optional or a pointer. If the dummy argument does not have the `VALUE` attribute, it corresponds to a C dummy argument that is a pointer.

Here is an example of a Fortran procedure together with its reference from C:

```
SUBROUTINE find_minmax(x,n,max,min) BIND(C,NAME='FindMinMax')  
  USE iso_c_binding  
  REAL(c_double) x(*),max,min  
  INTEGER(c_int),VALUE :: n  
  INTRINSIC maxval,minval  
  max = MAXVAL(x(:n))  
  min = MINVAL(x(:n))  
END  
  
extern void FindMinMax(double *x,int n,double *maxval,double *minval);  
double x[100],xmax,xmin;  
int n;
```



```
...  
FindMinMax(x,n,&xmax,&xmin);
```

This also allows C procedures to be called from Fortran, by describing the C procedure to be called in an interface block. Here is an example:

```
/* This is the prototype for a C library function from 4.3BSD. */  
int getloadavg(double loadavg[],int nelem);  
  
PROGRAM show_loadavg  
  USE iso_c_binding  
  INTERFACE  
    FUNCTION getloadavg(loadavg,nelem) BIND(C)  
      IMPORT c_double,c_int  
      REAL(c_double) loadavg(*)  
      INTEGER(c_int),VALUE :: nelem  
      INTEGER(c_int) getloadavg  
    END FUNCTION  
  END INTERFACE  
  REAL(c_double) averages(3)  
  IF (getloadavg(averages,3)/=3) THEN  
    PRINT *,'Unexpected error'  
  ELSE  
    PRINT *,'Load averages:',averages  
  END IF  
END
```

68.5 Enumerations

An enumeration defines a set of integer constants of the same kind, and is equivalent to the C `enum` declaration. For example,

```
ENUM,BIND(C)  
  ENUMERATOR :: open_door=4, close_door=17  
  ENUMERATOR :: lock_door  
END ENUM
```

is equivalent to

```
enum {  
  open_door=4, close_door=17, lock_door  
};
```

If a value is not given for one of the enumerators, it will be one greater than the previous value (or zero if it is the first enumerator in the list). The kind used for a particular set of enumerators can be discovered by using the `KIND` intrinsic on one of the enumerators.

Note that the `BIND(C)` clause is required; the standard only defines enumerations for interoperating with C.

69 IEEE arithmetic support [4.x except as otherwise noted]

69.1 Introduction

Three intrinsic modules are provided to support use of IEEE arithmetic, these are: `IEEE_ARITHMETIC`, `IEEE_EXCEPTIONS` and `IEEE_FEATURES`. This extension is small superset of the one described by the ISO Technical Report ISO/IEC TR 15580:1999.

69.2 Exception flags, modes and information flow

The model of IEEE arithmetic used by Fortran 2003 is that there is a global set of flags that indicate whether particular floating-point exceptions (such as overflow) have occurred, several operation modes which affect floating-point operations and exception handling, and a set of rules for propagating the flags and modes to obtain good performance and reliability.

The propagation rule for the exception flags is that the information “flows upwards”. Thus each procedure starts with the flags clear, and when it returns any flag that is set will cause the corresponding flag in the caller to become set. This ensures that procedures that can be executed in parallel will not interfere with each other via the IEEE exception flags. When the computer hardware only supports a single global set of flags, this model needs enforcement on in procedures that themselves examine the flags (by `IEEE_GET_FLAG` or `IEEE_GET_STATUS`).

The propagation rule for modes, is that the mode settings “flow downwards”. This enables code motion optimisations within all routines, and the only cost is that procedures which modify the modes must restore them (to the state on entry) when they return.

The modes that are available are:

- separately, whether each floating point exception terminates the program or allows execution to continue (providing a default result and raising an exception flag);
- rounding mode for floating-point operations;
- underflow mode.

69.3 Procedures in the modules

All the procedures provided by these modules are generic procedures, and not specific: that means that they cannot be passed as actual arguments.

In the descriptions of the procedures, where it says `REAL(*)` it means any kind of `REAL` (this is not standard Fortran syntax). Conversely, where it says `LOGICAL` it means default `LOGICAL` only, not any other kind of `LOGICAL`.

The functions whose names begin ‘`IEEE_SUPPORT_`’ are all enquiry functions. Many of these take a `REAL(*)` argument `X`; only the kind of `X` is used by the enquiry function, so `X` is permitted to be undefined, unallocated, disassociated, or an undefined pointer.

Note that a procedure must not be invoked on a data type that does not support the feature the procedure uses; the “support” enquiry functions can be used to detect this.

69.4 The `IEEE_FEATURES` module

This module defines the derived type `IEEE_FEATURES_TYPE`, and up to 11 constants of that type representing IEEE features: these are as follows.

<code>IEEE_DATATYPE</code>	whether any IEEE datatypes are available
<code>IEEE_DENORMAL</code>	whether IEEE denormal values are available*
<code>IEEE_DIVIDE</code>	whether division has the accuracy required by IEEE*
<code>IEEE_HALTING</code>	whether control of halting is supported
<code>IEEE_INEXACT_FLAG</code>	whether the inexact exception is supported*
<code>IEEE_INF</code>	whether IEEE positive and negative infinities are available*
<code>IEEE_INVALID_FLAG</code>	whether the invalid exception is supported*
<code>IEEE_NAN</code>	whether IEEE NaNs are available*
<code>IEEE_ROUNDING</code>	whether all IEEE rounding modes are available*
<code>IEEE_SQRT</code>	whether <code>SQRT</code> conforms to the IEEE standard*
<code>IEEE_UNDERFLOW_FLAG</code>	whether the underflow flag is supported*

(*) for at least one kind of `REAL`.

Those feature types which are required by the user procedure should be explicitly referenced by the `USE` statement with an `ONLY` clause, e.g.

```
USE,INTRINSIC :: IEEE_FEATURES,ONLY:IEEE_SQRT
```

This ensures that if the feature specified is not available the compilation will fail.

The type `IEEE_FEATURES_TYPE` is not in itself useful.

69.5 IEEE_EXCEPTIONS

Provides data types, constants and generic procedures for handling IEEE floating-point exceptions.

69.5.1 Types and constants

```
TYPE IEEE_STATUS_TYPE
```

Variables of this type can hold a floating-point status value; it combines all the mode settings and flags.

```
TYPE IEEE_FLAG_TYPE
```

Values of this type specify individual IEEE exception flags; constants for these are available as follows.

<code>IEEE_DIVIDE_BY_ZERO</code>	division by zero flag
<code>IEEE_INEXACT</code>	inexact result flag
<code>IEEE_INVALID</code>	invalid operation flag
<code>IEEE_OVERFLOW</code>	overflow flag
<code>IEEE_UNDERFLOW</code>	underflow flag

In addition, two array constants are available for indicating common combinations of flags:

```
TYPE(IEEE_FLAG_TYPE),PARAMETER :: &  
  IEEE_USUAL(3) = (/ IEEE_DIVIDE_BY_ZERO,IEEE_INVALID,IEEE_OVERFLOW /), &  
  IEEE_ALL(5) = (/ IEEE_DIVIDE_BY_ZERO,IEEE_INVALID,IEEE_OVERFLOW, &  
                 IEEE_UNDERFLOW,IEEE_INEXACT /)
```

69.5.2 Procedures

The procedures provided by `IEEE_EXCEPTIONS` are as follows.

```
ELEMENTAL SUBROUTINE IEEE_GET_FLAG(FLAG,FLAG_VALUE)  
  TYPE(IEEE_FLAG_TYPE),INTENT(IN) :: FLAG  
  LOGICAL,INTENT(OUT) :: FLAG_VALUE
```

Sets `FLAG_VALUE` to `.TRUE.` if the exception flag indicated by `FLAG` is currently set, and to `.FALSE.` otherwise.

```
ELEMENTAL SUBROUTINE IEEE_GET_HALTING_MODE(FLAG,HALTING)  
  TYPE(IEEE_FLAG_TYPE),INTENT(IN) :: FLAG  
  LOGICAL,INTENT(OUT) :: HALTING
```

Sets `HALTING` to `.TRUE.` if the program will be terminated on the occurrence of the floating-point exception designated by `FLAG`, and to `.FALSE.` otherwise.

```
PURE SUBROUTINE IEEE_GET_STATUS(STATUS_VALUE)  
  TYPE(IEEE_STATUS_TYPE),INTENT(OUT) :: STATUS_VALUE
```

Sets `STATUS_VALUE` to the current floating-point status; this contains all the current exception flag and mode settings.

```
PURE SUBROUTINE IEEE_SET_FLAG(FLAG,FLAG_VALUE)
  TYPE(IEEE_FLAG_TYPE),INTENT(IN) :: FLAG
  LOGICAL,INTENT(IN) :: FLAG_VALUE
```

Sets the exception flag designated by `FLAG` to `FLAG_VALUE`. `FLAG` may be an array of any rank, as long as it has no duplicate values, in which case `FLAG_VALUE` may be scalar or an array with the same shape.

```
PURE SUBROUTINE IEEE_SET_HALTING_MODE(FLAG,HALTING)
  TYPE(IEEE_FLAG_TYPE),INTENT(IN) :: FLAG
  LOGICAL,INTENT(IN) :: HALTING
```

Sets the halting mode for the exception designated by `FLAG` to `HALTING`. `FLAG` may be an array of any rank, as long as it has no duplicate values, in which case `HALTING` may be scalar or an array with the same shape.

```
PURE SUBROUTINE IEEE_SET_STATUS(STATUS_VALUE)
  TYPE(IEEE_STATUS_TYPE),INTENT(IN) :: STATUS_VALUE
```

Sets the floating-point status to that stored in `STATUS_VALUE`. This must have been previously obtained by calling `IEEE.GET_STATUS`.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_FLAG(FLAG)
  TYPE(IEEE_FLAG_TYPE),INTENT(IN) :: FLAG
```

Returns whether the exception flag designated by `FLAG` is supported for all kinds of `REAL`.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_FLAG(FLAG,X)
  TYPE(IEEE_FLAG_TYPE),INTENT(IN) :: FLAG
  REAL(*),INTENT(IN) :: X
```

Returns whether the exception flag designated by `FLAG` is supported for `REAL` with the kind of `X`.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_HALTING(FLAG)
  TYPE(IEEE_FLAG_TYPE),INTENT(IN) :: FLAG
```

Returns whether control of the “halting mode” for the exception designated by `FLAG` is supported.

69.6 IEEE_ARITHMETIC module

Provides additional functions supporting IEEE arithmetic: it includes the entire contents of `IEEE_EXCEPTIONS`.

69.6.1 IEEE datatype selection

The `IEEE_SELECTED_REAL_KIND` function is similar to the `SELECTED_REAL_KIND` intrinsic function, but selects among IEEE-compliant `REAL` types ignoring any that are not compliant.

69.6.2 Enquiry functions

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_DATATYPE()
```

Returns whether all real variables `X` satisfy the conditions for `IEEE_SUPPORT_DATATYPE(X)`.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_DATATYPE(X)
  REAL(*),INTENT(IN) :: X
```

Returns whether variables with the kind of **X** satisfy the following conditions:

- the numbers with absolute value between **TINY(X)** and **HUGE(X)** are exactly those of an IEEE floating-point format;
- the **+**, **-** and ***** operations conform to IEEE for at least one rounding mode;
- the functions **IEEE_COPY_SIGN**, **IEEE_LOGB**, **IEEE_NEXT_AFTER**, **IEEE_REM**, **IEEE_SCALB** and **IEEE_UNORDERED** may be used.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_DENORMAL()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_DENORMAL(X)  
  REAL(*), INTENT(IN) :: X
```

Returns whether for all real kinds, or variables with the kind of **X**, subnormal values (with absolute value between zero and **TINY**) exist as required by IEEE and operations on them conform to IEEE.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_DIVIDE()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_DIVIDE(X)  
  REAL(*), INTENT(IN) :: X
```

Returns whether intrinsic division (**/**) conforms to IEEE, for all real kinds or variables with the kind of **X** respectively.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_INF()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_INF(X)  
  REAL(*), INTENT(IN) :: X
```

Returns whether for all real kinds, or variables with the kind of **X**, positive and negative infinity values exist and behave in conformance with IEEE.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_IO()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_IO(X)  
  REAL(*), INTENT(IN) :: X
```

[5.2] Returns whether for all real kinds, or variables with the kind of **X**, conversion to and from text during formatted input/output conforms to IEEE, for the input/output rounding modes **ROUND='DOWN'**, **'NEAREST'**, **'UP'** and **'ZERO'** (and the corresponding edit descriptors **RD**, **RN**, **RU** and **RZ**).

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_NAN()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_NAN(X)  
  REAL(*), INTENT(IN) :: X
```

Returns whether for all real kinds, or variables with the kind of **X**, positive and negative “Not-a-Number” values exist and behave in conformance with IEEE.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_ROUNDING(ROUND_VALUE)  
  TYPE(IEEE_ROUND_TYPE), INTENT(IN) :: ROUND_VALUE
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_ROUNDING(ROUND_VALUE, X)  
  TYPE(IEEE_ROUND_TYPE), INTENT(IN) :: ROUND_VALUE  
  REAL(*), INTENT(IN) :: X
```

Returns whether for all real kinds, or variables with the kind of **X**, the rounding mode designated by **ROUND_VALUE** may be set using **IEEE_SET_ROUNDING_MODE** and conforms to IEEE.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_SQRT()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_SQRT(X)  
  REAL(*), INTENT(IN) :: X
```

Returns whether the intrinsic function **SQRT** conforms to IEEE, for all real kinds or variables with the kind of **X** respectively.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_STANDARD()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_STANDARD(X)  
  REAL(*), INTENT(IN) :: X
```

Returns whether for all real kinds, or variables with the kind of **X**, all the facilities described by the IEEE modules except for input/output conversions (see **IEEE_SUPPORT_IO**) are supported and conform to IEEE.

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_UNDERFLOW_CONTROL()
```

```
PURE LOGICAL FUNCTION IEEE_SUPPORT_UNDERFLOW_CONTROL(X)  
  REAL(*), INTENT(IN) :: X
```

[5.2] Returns whether for all real kinds, or variables with the kind of **X**, the underflow mode can be controlled with **IEEE_SET_UNDERFLOW_MODE**.

69.6.3 Rounding mode

```
TYPE IEEE_ROUND_TYPE
```

Values of this type specify the IEEE rounding mode. The following predefined constants are provided.

IEEE_DOWN	round down (towards minus infinity)
IEEE_NEAREST	round to nearest (ties to even)
IEEE_TO_ZERO	round positive numbers down, negative numbers up
IEEE_UP	round up (towards positive infinity)
IEEE_OTHER	any other rounding mode

```
PURE SUBROUTINE IEEE_GET_ROUNDING_MODE(ROUND_VALUE)  
  TYPE(IEEE_ROUND_TYPE), INTENT(OUT) :: ROUND_VALUE
```

Set **ROUND_VALUE** to the current rounding mode.

```
PURE SUBROUTINE IEEE_SET_ROUNDING_MODE(ROUND_VALUE)  
  TYPE(IEEE_ROUND_TYPE), INTENT(IN) :: ROUND_VALUE
```

Set the rounding mode to that designated by **ROUND_VALUE**.

69.6.4 Underflow mode

The underflow mode is either “gradual underflow” as specified by the IEEE standard, or “abrupt underflow”.

With gradual underflow, the space between **-TINY(X)** and **TINY(X)** is filled with equally-spaced “subnormal” numbers; the spacing of these numbers is equal to the spacing of model numbers above **TINY(X)** (and equal to the smallest

subnormal number). This gradually reduces the precision of the numbers as they get closer to zero: the smallest number has only one bit of precision, so any calculation with such a number will have a very large relative error.

With abrupt underflow, the only value between $-\text{TINY}(X)$ and $\text{TINY}(X)$ is zero. This kind of underflow is nearly universal in non-IEEE arithmetics and is widely provided by hardware even for IEEE arithmetic. Its main advantage is that it can be much faster.

```
SUBROUTINE IEEE_GET_UNDERFLOW_MODE(GRADUAL)
  LOGICAL, INTENT(OUT) :: GRADUAL
```

Sets `GRADUAL` to `.TRUE.` if the current underflow mode is gradual underflow, and to `.FALSE.` if it is abrupt underflow.

```
SUBROUTINE IEEE_SET_UNDERFLOW_MODE(GRADUAL)
  LOGICAL, INTENT(IN) :: GRADUAL
```

Sets the underflow mode to gradual underflow if `GRADUAL` is `.TRUE.`, and to abrupt underflow if it is `.FALSE.`.

69.6.5 Number Classification

```
TYPE IEEE_CLASS_TYPE
```

Values of this type indicate the IEEE class of a number; this is equal to one of the provided constants:

<code>IEEE_NEGATIVE_DENORMAL</code>	negative subnormal number x , in the range $-\text{TINY}(x) < x < 0$
<code>IEEE_NEGATIVE_INF</code>	$-\infty$ (negative infinity)
<code>IEEE_NEGATIVE_NORMAL</code>	negative normal number x , in the range $-\text{HUGE}(x) \leq x \leq -\text{TINY}(x)$
<code>IEEE_NEGATIVE_ZERO</code>	-0 (zero with the sign bit set)
<code>IEEE_POSITIVE_DENORMAL</code>	positive subnormal number x , in the range $0 < x < \text{TINY}(x)$
<code>IEEE_POSITIVE_INF</code>	$+\infty$ (positive infinity)
<code>IEEE_POSITIVE_NORMAL</code>	positive normal number x , in the range $\text{TINY}(x) \leq x \leq \text{HUGE}(x)$
<code>IEEE_POSITIVE_ZERO</code>	$+0$ (zero with the sign bit clear)
<code>IEEE_QUIET_NAN</code>	Not-a-Number (usually the result of an invalid operation)
<code>IEEE_SIGNALING_NAN</code>	Not-a-Number which raises the invalid signal on reference
[5.2] <code>IEEE_OTHER_VALUE</code>	any value that does not fit one of the above categories

The comparison operators `.EQ.` (`=`) and `.NE.` (`/=`) are provided for comparing values of this type.

```
ELEMENTAL TYPE(IEEE_CLASS_TYPE) FUNCTION IEEE_CLASS(X)
  REAL(*), INTENT(IN) :: X
```

returns the classification of the value of `X`.

```
ELEMENTAL REAL(*) FUNCTION IEEE_VALUE(X, CLASS)
  REAL(*), INTENT(IN) :: X
  TYPE(IEEE_CLASS_TYPE), INTENT(IN) :: CLASS
```

Returns a “sample” value with the kind of `X` and the classification designated by `CLASS`.

69.6.6 Test functions

The following procedures are provided for testing IEEE values.

```
ELEMENTAL LOGICAL FUNCTION IEEE_IS_FINITE(X)
  REAL(*), INTENT(IN) :: X
```

Returns whether `X` is “finite”, i.e. not an infinity, NaN, or `IEEE_OTHER_VALUE`.

```
ELEMENTAL LOGICAL FUNCTION IEEE_IS_NAN(X)
  REAL(*), INTENT(IN) :: X
```

Returns whether X is a NaN.

```
ELEMENTAL LOGICAL FUNCTION IEEE_IS_NEGATIVE(X)
  REAL(*), INTENT(IN) :: X
```

Returns whether X is negative; it differs the comparison $X < 0$ only in the case of negative zero, where it returns `.TRUE..`

```
ELEMENTAL LOGICAL FUNCTION IEEE_UNORDERED(X,Y)
  REAL(*), INTENT(IN) :: X,Y
```

Returns the value of 'IEEE_IS_NAN(X) .OR. IEEE_IS_NAN(Y)'.

69.6.7 Arithmetic functions

```
ELEMENTAL REAL(*) FUNCTION IEEE_COPY_SIGN(X,Y)
  REAL(*), INTENT(IN) :: X,Y
```

Returns X with the sign bit of Y.

```
ELEMENTAL REAL(*) FUNCTION IEEE_LOGB(X)
  REAL(*), INTENT(IN) :: X
```

If X is zero, returns $-\infty$ if infinity is supported and `-HUGE(X)` otherwise. For nonzero X, returns `EXPONENT(X)-1`

```
ELEMENTAL REAL(*) FUNCTION IEEE_NEXT_AFTER(X,Y)
  REAL(*), INTENT(IN) :: X,Y
```

Returns the nearest number to X that is closer to Y, or X if X and Y are equal.

```
ELEMENTAL REAL(*) FUNCTION IEEE_REM(X,Y)
  REAL(*), INTENT(IN) :: X,Y
```

Returns the exact remainder resulting from the division of X by Y.

```
ELEMENTAL REAL(*) FUNCTION IEEE_RINT(X)
  REAL(*), INTENT(IN) :: X
```

Returns X rounded to an integer according to the current rounding mode.

```
ELEMENTAL REAL(*) FUNCTION IEEE_SCALB(X,I)
  REAL(*), INTENT(IN) :: X
  INTEGER(*), INTENT(IN) :: I
```

Returns `SCALE(X,I)`, i.e. $X \cdot 2^I$, without computing 2^I separately.

70 Input/output Features

70.1 Stream input/output [5.1]

A stream file is a file that is opened with the `ACCESS='STREAM'` specifier. A stream file is either a formatted stream or an unformatted stream.

A formatted stream file is equivalent to a C text stream; this acts much like an ordinary sequential file, except that there is no limit on the length of a record. Just as in C, when writing to a formatted stream, an embedded newline character in the data causes a new record to be created. The new intrinsic enquiry function `NEW_LINE(A)` returns this character for the kind (character set) of `A`; if the character set is ASCII, this is equal to `IACHAR(10)`. For example,

```
OPEN(17,FORM='formatted',ACCESS='stream',STATUS='new')
WRITE(17,'(A)'), 'This is record 1.'//NEW_LINE('A')// 'This is record 2.'
```

An unformatted stream file is equivalent to a C binary stream, and has no record boundaries. This makes it impossible to `BACKSPACE` an unformatted stream. Data written to an unformatted stream is transferred to the file with no formatting, and data read from an unformatted stream is transferred directly to the variable as it appears in the file.

When reading or writing a stream file, the `POS=` specifier may be used to specify where in the file the data is to be written. The first character of the file is at position 1. The `POS=` specifier may also be used in an `INQUIRE` statement, in which case it returns the current position in the file. When reading or writing a formatted stream, the `POS=` in a `READ` or `WRITE` shall be equal to 1 (i.e. the beginning of the file) or to a value previously discovered through `INQUIRE`.

Note that unlike unformatted sequential files, writing to an unformatted stream file at a position earlier than the end of the file does not truncate the file. (However, this truncation does happen for formatted streams.)

Finally, the `STREAM=` specifier has been added to the `INQUIRE` statement. This specifier takes a scalar default character variable, and assigns it the value `'YES'` if the file may be opened for stream input/output (i.e. with `ACCESS='STREAM'`), the value `'NO'` if the file cannot be opened for stream input/output, and the value `'UNKNOWN'` if it is not known whether the file may be opened for stream input/output.

70.2 The `BLANK=` and `PAD=` specifiers [5.1]

The `BLANK=` and `PAD=` specifiers, which previously were only allowed on an `OPEN` statement (and `INQUIRE`), are now allowed on a `READ` statement. These change the `BLANK=` or `PAD=` mode for the duration of that `READ` statement only.

70.3 Decimal Comma [5.1]

It is possible to read and write numbers with a decimal comma instead of a decimal point. Support for this is provided by the `DECIMAL=` specifier and the `DC` and `DP` edit descriptors. The `DECIMAL=` specifier may appear in `OPEN`, `READ`, `WRITE` and `INQUIRE` statements; possible values are `'POINT'` (the default) and `'COMMA'`. For an unconnected or unformatted unit, `INQUIRE` returns `'UNDEFINED'`. The `DC` edit descriptor temporarily sets the mode to `DECIMAL='COMMA'`, and the `DP` edit descriptor temporarily sets the mode to `DECIMAL='POINT'`.

When the mode is `DECIMAL='COMMA'`, all floating-point output will produce a decimal comma instead of a decimal point, and all floating-point input will expect a decimal comma. For example,

```
PRINT '(1X,"Value cest ",DC,F0.2)',1.25
```

will produce the output

```
Value cest 1,25
```

Additionally, in this mode, a comma cannot be used in list-directed input to separate items; instead, a semi-colon may be used.

70.4 The DELIM= specifier [5.1]

The DELIM= specifier, which previously was only allowed on an OPEN statement (and INQUIRE), is now allowed on a WRITE statement. It changes the DELIM= mode for the duration of that WRITE statement; note that this only has any effect if the WRITE statement uses list-directed or namelist output. For example,

```
WRITE(*,*,DELIM='QUOTE') "That's all folks!"
```

will produce the output

```
'That''s all folks!'
```

70.5 The ENCODING= specifier [5.1]

The ENCODING= specifier is permitted on OPEN and INQUIRE statements. Standard values for this specifier are 'UTF-8' and the default value of 'DEFAULT'; the 'UTF-8' value is only allowed on compilers that support a Unicode character kind (see SELECTED_CHAR_KIND). This release of the NAG Fortran Compiler supports 'DEFAULT' and 'ASCII'.

70.6 The IOMSG= specifier [5.1]

The IOMSG= specifier has been added to all input/output statements. This takes a scalar default character variable, which in the event of an error is assigned an explanatory message. (Note that this is only useful if the statement contains an IOSTAT= or ERR= specifier, otherwise the program will be terminated on error anyway.) If no error occurs, the value of the variable remains unchanged.

70.7 The IOSTAT= specifier [5.1]

This now accepts any kind of integer variable (previously this was required to be default integer).

70.8 The SIGN= specifier [5.1]

The SIGN= specifier has been added to the OPEN, WRITE and INQUIRE statements; possible values are 'PLUS', 'SUPPRESS' and 'PROCESSOR_DEFINED' (the default). For the NAG Fortran Compiler, SIGN='PROCESSOR_DEFINED' has the same effect as SIGN='SUPPRESS'.

The effect of SIGN='PLUS' is the same as the SP edit descriptor, the effect of SIGN='SUPPRESS' is the same as the SS edit descriptor, and the effect of SIGN='PROCESSOR_DEFINED' is the same as the S edit descriptor.

70.9 Intrinsic functions for testing IOSTAT= values [5.1]

The intrinsic functions IS_IOSTAT_END and IS_IOSTAT_EOR test IOSTAT= return values, determining whether a value indicates an end-of-file condition or an end-of-record condition. These are equivalent to testing the IOSTAT= return value against the named constants IOSTAT_END and IOSTAT_EOR respectively; these constants are available through the ISO_FORTRAN_ENV module.

70.10 Input/output of IEEE infinities and NaNs [5.1]

Input and output of IEEE infinities and NaNs is possible: the output format is

- -Infinity (or -Inf if it will not fit) for negative infinity;
- Infinity (or Inf if it will not fit) for positive infinity, or +Infinity (or +Inf) with SP or SIGN='PLUS' mode;

- NaN for a NaN.

Furthermore, the output is right-justified within the output field. For list-directed output the output field is the minimum size to hold the result.

Input of IEEE infinities and NaNs is now possible; these take the same form as the output described above, except that:

- case is not significant,
- a NaN may be preceded by a sign, which is ignored, and
- a NaN may be followed by alphanumeric characters enclosed in parentheses (the NAG Fortran Compiler also ignores these).

The result of reading a NaN value in NAG Fortran is always a quiet NaN, never a signalling one.

70.11 Output of floating-point zero [5.1]

List-directed and namelist output of floating-point zero is now done using F format instead of E format. (The Fortran 90 and 95 standards both specified E format.)

70.12 NAMELIST and internal files [5.1]

Namelist input/output is now permitted to/from internal files.

70.13 Recursive input/output [5.2]

Input/output to internal files is now permitted while input/output to another internal file or an external file is in progress. This occurs when a function in an input/output list executes an internal file READ or WRITE statement.

Input/output to an external file while external file input/output is already in progress remains prohibited, except for the case of nested data transfer (see “Defined input/output”).

70.14 Asynchronous input/output

70.14.1 Basic syntax [5.1]

Asynchronous input/output syntax is accepted; this consists of the **ASYNCHRONOUS=** specifier on **OPEN**, **READ**, **WRITE** and **INQUIRE**, the **ID=** specifier on **READ**, **WRITE** and **INQUIRE**, and the **PENDING=** specifier on **INQUIRE**.

Except for the **INQUIRE** statement, the **ASYNCHRONOUS=** specifier takes a scalar default character expression; this must evaluate either to **'YES'** or **'NO'**, treating lowercase the same as uppercase. In the **READ** and **WRITE** statements, this character expression must be a constant expression, and the statement must refer to an external file whose connection allows asynchronous input/output; if **ASYNCHRONOUS='YES'** is specified, the data transfer may occur asynchronously. In the **OPEN** statement, this specifier determines whether asynchronous data transfer is allowed for that file: the default setting is **'NO'**. For the **INQUIRE** statement, the **ASYNCHRONOUS=** specifier takes a scalar default character variable, and sets it to **'YES'** if the file is currently connected for asynchronous input/output, **'NO'** if the current connection does not allow asynchronous input/output and **'UNKNOWN'** if the file is not connected.

For the **READ** and **WRITE** statements, the **ID=** specifier takes a scalar integer variable. This specifier is only permitted if **ASYNCHRONOUS='YES'** also appears. The integer variable is assigned the “identifier” of the asynchronous data transfer that the **READ** or **WRITE** initiates; this value can be used in **INQUIRE** and **WAIT** statements to track the progress of the asynchronous data transfer.

For the **INQUIRE** statement, the **ID=** specifier takes a scalar integer expression whose value must be that returned from **ID=** on a **READ** or **WRITE** statement for that file, and is only permitted in conjunction with the **PENDING=** specifier.

The `PENDING=` specifier takes a scalar default logical variable and sets it to `.TRUE.` if the specified asynchronous data transfer is still underway and to `.FALSE.` if it has completed. If `PENDING=` is used without `ID=`, the enquiry is about all outstanding asynchronous data transfer on that file.

After initiating an asynchronous data transfer, the variables affected must not be referenced or defined until after the transfer is known to have finished. For an asynchronous `WRITE` of a local variable, this means not returning from the procedure until after ensuring the transfer is complete. An asynchronous data transfer is known to have been finished if there is a subsequent synchronous data transfer, an `INQUIRE` statement which returns `.FALSE.` for `PENDING=`, or a `WAIT` statement has been executed; in each case, for that file.

70.14.2 Basic Example

The following example uses two buffers, `buf1` and `buf2`, alternately reading into one while processing the other, while there is still data in the file to process (each dataset being followed by a single logical value which indicates whether more is to come).

```
REAL :: buf1(n,m),buf2(n,m)
LOGICAL more,stillwaiting
READ (unit) buf1
DO
  READ (unit) more ! Note: synchronous
  IF (more) READ (unit,ASYNCHRONOUS='YES') buf2
  CALL process(buf1)
  IF (.NOT.more) EXIT
  READ (unit) more ! Note: synchronous
  IF (more) READ (unit,ASYNCHRONOUS='YES') buf1
  CALL process(buf2)
  IF (.NOT.more) EXIT
END DO
```

Note that the synchronous `READ` statements automatically “wait” for any outstanding asynchronous data transfer to complete before reading the logical value; this ensures that the dataset will have finished being read into its buffer and is safe to process.

70.14.3 The `ASYNCHRONOUS` attribute [5.2]

A `READ` or `WRITE` statement with `ASYNCHRONOUS='YES'` automatically gives the `ASYNCHRONOUS` attribute to any variable that appears in its input/output list, in a `SIZE=` specifier, or which is part of a namelist specified by `NML=`. This is adequate for asynchronous data transfers that initiate and complete within a single procedure. However, it is inadequate for transfers to/from module variables or dummy arguments if the procedure returns while the transfer is still underway.

The `ASYNCHRONOUS` attribute may be explicitly specified in a type declaration statement or in an `ASYNCHRONOUS` statement. The latter has the syntax

```
ASYNCHRONOUS [::] variable-name [ , variable-name ]...
```

If a variable with the `ASYNCHRONOUS` attribute is a dummy array and is not an assumed-shape array or array pointer, any associated actual argument cannot be an array section, an assumed-shape array or array pointer. Furthermore, if a dummy argument has the `ASYNCHRONOUS` attribute the procedure must have an explicit interface. Both of these restrictions apply whether the attribute was given explicitly or implicitly.

70.14.4 The `WAIT` statement [5.2]

The `WAIT` statement provides the ability to wait for an asynchronous data transfer to finish without performing any other input/output operation. It has the form

```
WAIT ( wait-spec [ , wait-spec ]... )
```

where *wait-spec* is one of the following:

```
UNIT = file-unit-number  
END = label  
EOR = label  
ERR = label  
ID = scalar-integer-variable  
IOMSG = scalar-default-character-variable  
IOSTAT = scalar-integer-variable
```

The UNIT= specifier must appear, but the 'UNIT =' may be omitted if it is the first specifier in the list. The ID= specifier takes a scalar integer expression whose value must be that returned from ID= on a READ or WRITE statement for the file; if the ID= specifier does not appear the WAIT statement refers to all pending asynchronous data transfer for that file.

On completion of execution of the WAIT statement the specified asynchronous data transfers have been completed. If the specified file is not open for asynchronous input/output or is not connected, the WAIT statement has no effect (it is not an error unless the ID= specifier was used with an invalid value).

Here is an example of using the WAIT statement.

```
REAL array(1000,1000,10),xferid(10)  
! Start reading each segment of the array  
DO i=1,10  
  READ (unit,id=xfer(i)) array(:,:,i)  
END DO  
...  
! Now process each segment of the array  
DO i=1,10  
  WAIT (unit,id=xfer(i))  
  CALL process(array(:,:,i))  
END DO
```

70.14.5 Execution Semantics

At this time all actual input/output operations remain synchronous, as allowed by the standard.

70.15 Scale factor followed by repeat count [5.1]

The comma that was previously required between a scale factor (*nP*) and a repeat count (e.g. the '3' in 3E12.2), is now optional. This trivial extension was part of Fortran 66 that was removed in Fortran 77, and reinstated in Fortran 2003.

70.16 FLUSH statement [5.2]

Execution of a FLUSH statement causes the data written to a specified file to be made available to other processes, or causes data placed in that file by another process to become available to the program. The syntax of the FLUSH statement is similar to the BACKSPACE, ENDFILE and REWIND statements, being one of the two possibilities

```
FLUSH file-unit-number  
FLUSH ( flush-spec [ , flush-spec ]... )
```

where *file-unit-number* is the logical unit number (a scalar integer expression), and *flush-spec* is one of the following:

```

UNIT = file-unit-number
IOSTAT = scalar-integer-variable
IOMSG = scalar-default-character-variable
ERR = label

```

The UNIT= specifier must appear, but the ‘UNIT =’ may be omitted if it is the first *flush-spec* in the list.

Here is an example of the use of a FLUSH statement.

```

WRITE (pipe) my_data
FLUSH (pipe)

```

70.17 Defined input/output [n/a]

The generic identifiers READ(FORMATTED), READ(UNFORMATTED), WRITE(FORMATTED) and WRITE(UNFORMATTED) provide the ability to replace normal input/output processing for an item of derived type. In the case of formatted input/output, the replacement will occur for list-directed formatting, namelist formatting, and for an explicit format with the DT edit descriptor: it does not affect any other edit descriptors in an explicit format. Using defined input/output, it is possible to perform input/output on derived types containing pointer components and allocatable components, since the user-defined procedure will be handling it.

Here is a type definition with defined input/output procedures.

```

TYPE tree
  TYPE(tree_node),POINTER :: first
CONTAINS
  PROCEDURE :: fmtread=>tree_fmtread
  PROCEDURE :: fmtwrite=>tree_fmtwrite
  GENERIC,PUBLIC :: READ(formatted)=>fmtread, WRITE(formatted)=>fmtwrite
END TYPE

```

Given the above type definition, whenever a TYPE(tree) object is an effective item in a formatted input/output list, the module procedure `tree_fmtread` will be called (in a READ statement) or the module procedure `tree_fmtwrite` will be called (in a WRITE statement) to perform the input or output of that object. Note that a generic interface block may also be used to declare procedures for defined input/output; this is the only option for sequence or BIND(C) types but is not recommended for extensible types.

The procedures associated with each input/output generic identifier must have the same characteristics as the ones listed below, where *type-declaration* is CLASS(*derived-type-spec*) for an extensible type and TYPE(*derived-type-spec*) for a sequence or BIND(C) type. Note that if the derived type has any length type parameters, they **must** be “assumed” (specified as ‘*’).

```

SUBROUTINE formatted_read(var,unit,iotype,vlist,iostat,iomsg)

```

```

  type-declaration,INTENT(INOUT) :: var
  INTEGER,INTENT(IN) :: unit
  CHARACTER(*),INTENT(IN) :: iotype
  INTEGER,INTENT(IN) :: vlist(:)
  INTEGER,INTENT(OUT) :: iostat
  CHARACTER(*),INTENT(INOUT) :: iomsg

```

```

SUBROUTINE unformatted_read(var,unit,iostat,iomsg)

```

```

  type-declaration,INTENT(INOUT) :: var
  INTEGER,INTENT(IN) :: unit
  INTEGER,INTENT(OUT) :: iostat
  CHARACTER(*),INTENT(INOUT) :: iomsg

```

```

SUBROUTINE formatted_write(var,unit,iotype,vlist,iostat,iomsg)

```

```

  type-declaration,INTENT(IN) :: var
  INTEGER,INTENT(IN) :: unit
  CHARACTER(*),INTENT(IN) :: iotype

```

```
INTEGER,INTENT(IN) :: vlist(:)
INTEGER,INTENT(OUT) :: iostat
CHARACTER(*),INTENT(INOUT) :: iomsg

SUBROUTINE unformatted_write(var,unit,iostat,iomsg)
  type-declaration,INTENT(IN) :: var
  INTEGER,INTENT(IN) :: unit
  INTEGER,INTENT(OUT) :: iostat
  CHARACTER(*),INTENT(INOUT) :: iomsg
```

In each procedure, `unit` is either a normal unit number if the parent input/output statement used a normal unit number, a negative number if the parent input/output statement is for an internal file, or a processor-dependent number (which might be negative) for the ‘*’ unit. The `iostat` argument **must** be assigned a value before returning from the defined input/output procedure: either zero to indicate success, the negative number `IOSTAT_EOR` (from the intrinsic module `ISO_FORTRAN_ENV`) to signal an end-of-record condition, the negative number `IOSTAT_END` to signal an end-of-file condition, or a positive number to indicate an error condition. The `iomsg` argument **must** be left alone if no error occurred, and must be assigned an explanatory message if `iostat` is set to a nonzero value.

For the formatted input/output procedures, the `iotype` argument will be set to ‘LISTDIRECTED’ if list-directed formatting is being done, ‘NAMELIST’ if namelist formatting is being done, and ‘DT’ concatenated with the *character-literal* if the DT edit descriptor is being processed. The `vlist` argument contains the list of values in the DT edit descriptor if present, and is otherwise a zero-sized array. Note that the syntax of the DT edit descriptor is:

$$DT [\textit{character-literal}] [(\textit{value} [, \textit{value}] \dots)]$$

where blanks are insignificant, *character-literal* is a default character literal constant with no kind parameter, and each *value* is an optionally signed integer literal constant with no kind parameter. For example, ‘DT’, ‘DT’z8,i4,e10.2’, ‘DT(100,-3,+4.666)’ and ‘DT”silly example”(0)’ are all syntactically correct DT edit descriptors: it is up to the user-defined procedure to interpret what they might mean.

During execution of a defined input/output procedure, there must be no input/output for an external unit (other than for the `unit` argument), but input/output for internal files is permitted. No file positioning commands are permitted. For unformatted input/output, all input/output occurs within the current record, no matter how many separate data transfer statements are executed by the procedure; that is, file positioning both before and after “nested” data transfer is suppressed. For formatted input/output, this effect is approximately equivalent to the nested data transfer statements being considered to be nonadvancing; explicit record termination (using the slash (/) edit descriptor, or transmission of newline characters to a stream file) is effective, and record termination may be performed by a nested list-directed or namelist input/output statement.

If `unit` is associated with an external file (i.e. non-negative, or equal to one of the constants `ERROR_UNIT`, `INPUT_UNIT` or `OUTPUT_UNIT` from the intrinsic module `ISO_FORTRAN_ENV`), the current settings for the pad mode, sign mode, etc., can be discovered by using `INQUIRE` with `PAD=`, `SIGN=`, etc. on the `unit` argument. Note that `INQUIRE` must not be used if `unit` is associated with an internal file.

Finally, defined input/output is not compatible with asynchronous input/output; all input/output statements involved with defined input/output must be synchronous.

71 Miscellaneous Fortran 2003 Features

71.1 Abstract interfaces and the PROCEDURE statement [5.1]

Abstract interfaces have been added, together with the procedure declaration statement. An abstract interface is defined in an interface block that has the `ABSTRACT` keyword, i.e.

```
ABSTRACT INTERFACE
```

Each interface body in an abstract interface block defines an abstract interface instead of declaring a procedure. The name of an abstract interface can be used in the procedure declaration statement to declare a specific procedure with that interface, e.g.

```
PROCEDURE(aname) :: spec1, spec2
```

declares `SPEC1` and `SPEC2` to be procedures with the interface (i.e. type, arguments, etc.) defined by the abstract interface `ANAME`.

The procedure declaration statement can also be used with the name of any procedure that has an explicit interface, e.g.

```
PROCEDURE(x) y
```

declares `Y` to have the same interface as `X`. Also, procedures with implicit interfaces can be declared by using `PROCEDURE` with a type specification instead of a name, or by omitting the name altogether.

The following attributes can be declared at the same time on the procedure declaration statement: `BIND(C...)`, `INTENT(intent)`, `OPTIONAL`, `POINTER`, `PRIVATE`, `PUBLIC`, `SAVE`. For example,

```
PROCEDURE(aname),PRIVATE :: spec3
```

Note that `POINTER` declares a procedure pointer (see next section), and that `INTENT` and `SAVE` are only allowed for procedure pointers not for ordinary procedures. The NAG Fortran Compiler also allows the `PROTECTED` attribute to be specified on the procedure declaration statement: this is an extension to the published Fortran 2003 standard.

71.2 Named procedure pointers [5.2]

A procedure pointer is a procedure with the `POINTER` attribute; it may be a named pointer or a structure component (the latter are described elsewhere). The usual way of declaring a procedure pointer is with the procedure declaration statement, by including the `POINTER` clause in that statement: for example,

```
PROCEDURE(aname),POINTER :: p => NULL()
```

declares `P` to be a procedure pointer with the interface `ANAME`, and initialises it to be a disassociated pointer.

A named procedure pointer may also be declared by specifying the `POINTER` attribute in addition to its normal procedure declaration: for example, a function declared by a type declaration statement will be a function pointer if the `POINTER` attribute is included in the type declaration:

```
REAL, EXTERNAL, POINTER :: funptr
```

The `POINTER` statement can also be used to declare a procedure pointer, either in conjunction with an interface block, an `EXTERNAL` statement, or a type declaration statement, for example:

```
INTERFACE
  SUBROUTINE sub(a,b)
    REAL,INTENT(INOUT) :: a,b
  END SUBROUTINE
END INTERFACE
POINTER sub
```

Procedure pointers may also be stored in derived types as procedure pointer components. The syntax and effects are slightly different, making them act like “object-bound procedures”, and as such are described in the object-oriented programming section.

71.3 Intrinsic modules [4.x]

The Fortran 2003 standard classifies modules as either intrinsic or non-intrinsic. A non-intrinsic module is the normal kind of module (i.e. user-defined); an intrinsic module is one that is provided as an intrinsic part of the Fortran compiler.

There are five **standard** modules in Fortran 2003: IEEE_ARITHMETIC, IEEE_EXCEPTIONS, IEEE_FEATURES, ISO_C_BINDING and ISO_FORTRAN_ENV.

A program is permitted to have a non-intrinsic module with the same name as that of an intrinsic module: to this end, the USE statement has been extended: 'USE,INTRINSIC ::' specifies that an intrinsic module is required, whereas 'USE,NON_INTRINSIC ::' specifies that a non-intrinsic module is required. If these are not used, the compiler will select an intrinsic module only if no user-defined module is found. For example,

```
USE,INTRINSIC :: iso_fortran_env
```

uses the standard intrinsic module ISO_FORTRAN_ENV, whereas

```
USE,NON_INTRINSIC :: iso_fortran_env
```

uses a user-defined module with that name. Note that the double-colon '::' is required if either specifier is used.

71.4 Renaming user-defined operators on the USE statement [5.2]

It is now possible to rename a user-defined operator on the USE statement, similarly to how named entities can be renamed. For example,

```
USE my_module, OPERATOR(.localid.)=>OPERATOR(.remotename.)
```

would import everything from MY_MODULE, but the .REMOTENAME. operator would have its name changed to .LOCALID..

Note that this is only available for user-defined operator names; the intrinsic operators .AND. et al cannot have their names changed in this way, nor can ASSIGNMENT(=) be renamed. The local name must be an operator if and only if the remote (module entity) name is an operator: that is, both of

```
USE my_module, something=>OPERATOR(.anything.)
USE my_module, OPERATOR(.something.)=>anything
```

are invalid (a syntax error will be produced).

71.5 The ISO_FORTRAN_ENV module [5.1]

The standard intrinsic module ISO_FORTRAN_ENV is now available. It contains the following default INTEGER named constants.

CHARACTER_STORAGE_SIZE

size of a character storage unit in bits.

ERROR_UNIT

logical unit number for error reporting ("stderr").

FILE_STORAGE_SIZE

size of the file storage unit used by RECL= in bits.

INPUT_UNIT

default (*) unit number for READ.

IOSTAT_END

IOSTAT= return value for end-of-file.

IOSTAT_EOR

IOSTAT= return value for end-of-record.

NUMERIC_STORAGE_SIZE

size of a numeric storage unit in bits.

OUTPUT_UNIT

unit used by PRINT, the same as the '*' unit for WRITE.

71.6 The IMPORT statement [5.1]

The **IMPORT** statement has been added. This has the syntax

```
IMPORT [ [ :: ] name [ , name ]... ]
```

and is only allowed in an interface body, where it imports the named entities from the host scoping unit (normally, these entities cannot be accessed from an interface body). If no names are specified, normal host association rules are in effect for this interface body.

The **IMPORT** statement must follow any **USE** statements and precede all other declarations, in particular, **IMPLICIT** and **PARAMETER** statements. Anything imported with **IMPORT** must have been declared prior to the interface body.

71.7 Length of names and statements

Names are now ([4.x]) permitted to be 63 characters long (instead of 31), and statements are now ([5.2]) permitted to have 255 continuation lines (instead of 39).

71.8 Array constructor syntax enhancements

Square brackets ([]) can now ([5.1]) be used in place of the parenthesis-slash pairs ((/ /)) for array constructors. This allows expressions to be more readable when array constructors are being mixed with ordinary parentheses.

```
RESHAPE((/i/2.0,i=1,100/),(/2,3/))  ! Old way
RESHAPE([(i/2.0,i=1,100)], [2,3])    ! New way
```

Array constructors may now ([5.2]) begin with a type specification followed by a double colon (::); this makes zero-sized constructors easy (and eliminates potential ambiguity with character length), and also provides assignment conversions thus eliminating the need to pad all character strings to the same length.

```
[ Logical :: ]                ! Zero-sized logical array
[ Double Precision :: 17.5, 0, 0.1d0 ] ! Conversions
[ Character(200) :: 'Alf', 'Bernadette' ] ! Padded to length 200
```

71.9 Structure constructor syntax enhancements [5.3]

There are three enhancements that have been made to structure constructors in Fortran 2003:

1. component names can be used as keywords, the same way that dummy argument names can be used as argument keywords;
2. values can be omitted for components that have default initialisation; and
3. type names can be the same as generic function names, and references are resolved by choosing a suitable function (if the syntax matches the function's argument list) and treating as a structure constructor only if no function matches the actual arguments.

A fourth enhancement is made in the Fortran 2008 standard: a value can be omitted for a component that is allocatable.

This makes structure constructors more like built-in generic functions that can be overridden when necessary. Here is an example showing all three enhancements.

```
TYPE quaternion
  REAL x=0,ix=0,jx=0,kx=0
END TYPE
...
```

```
INTERFACE quaternion
  MODULE PROCEDURE quat_from_complex
END INTERFACE
...
TYPE(quaternion) FUNCTION quat_from_complex(c) RESULT(r)
  COMPLEX c
  r%x = REAL(c)
  r%y = AIMAG(c)
  r%z = 0
  r%a = 0
END FUNCTION
...
COMPLEX c
TYPE(quaternion) q
q = quaternion(3.14159265) ! Structure constructor, value (~pi,0,0,0).
q = quaternion(jx=1)      ! Structure constructor, value (0,0,1,0).
q = quaternion(c)         ! "Constructor" function quat_from_complex.
```

Also, if the type is an extended type an ancestor component name can be used to provide a value for all those inherited components at once.

These extensions mean that even if a type has a private component, you can use the structure constructor if

- the component is allocatable (it will be unallocated),
- the component is default-initialised (it will have the default value), or
- the component is inherited and you use an ancestor component name to provide a value for it and the other components inherited from that ancestor.

71.10 Deferred character length [5.2]

The length of a character pointer or allocatable variable can now be declared to be deferred, by specifying the length as a colon: for example,

```
CHARACTER(LEN=:),POINTER :: ch
```

The length of a deferred-length pointer (or allocatable variable) is determined when it is allocated (see next section) or pointer-associated; for example

```
CHARACTER,TARGET :: t1*3,t2*27
CHARACTER(:),POINTER :: p
p => t1
PRINT *,LEN(p)
p => t2
PRINT *,LEN(p)
```

will first print 3 and then 27. It is not permitted to ask for the `LEN` of a disassociated pointer that has deferred length.

Note that deferred length is most useful in conjunction with the new features of typed allocation, sourced allocation, scalar allocatables and automatic reallocation.

71.11 The `ERRMSG=` specifier [5.1]

The `ALLOCATE` and `DEALLOCATE` statements now accept the `ERRMSG=` specifier. This specifier takes a scalar default character variable, which in the event of an allocation or deallocation error being detected will be assigned an explanatory message. If no error occurs the variable is left unchanged. Note that this is useless unless the `STAT=` specifier is also used, as otherwise the program will be terminated on error anyway.

For example,

```
ALLOCATE(w(n),STAT=ierror,ERRMSG=message)
IF (ierror/=0) THEN
  PRINT *, 'Error allocating W: ', TRIM(message)
  RETURN
END IF
```

71.12 Intrinsic functions in constant expressions [5.2 partial; 5.3 complete]

It is now allowed to use any intrinsic function with constant arguments in a constant expression. (In Fortran 95 real and complex intrinsic functions were not allowed.) For example,

```
MODULE m
  REAL, PARAMETER :: e = EXP(1.0)
END
```

71.13 Access to the command line [5.1]

The intrinsic procedures `COMMAND_ARGUMENT_COUNT`, `GET_COMMAND` and `GET_COMMAND_ARGUMENT` have been added. These duplicate functionality previously only available via the procedures `IARGC` and `GETARG` from the `F90_UNIX_ENV` module.

```
INTEGER FUNCTION command_argument_count()
```

Returns the number of command-line arguments. Unlike `IARGC` in the `F90_UNIX_ENV` module, this returns 0 even if the command name cannot be retrieved.

```
SUBROUTINE get_command(command,length,status)
  CHARACTER(*), INTENT(OUT), OPTIONAL :: command
  INTEGER, INTENT(OUT), OPTIONAL :: length, status
```

Accesses the command line which invoked the program. This is formed by concatenating the command name and the arguments separated by blanks. This might differ from the command the user actually typed, and should be avoided (use `GET_COMMAND_ARGUMENT` instead).

If `COMMAND` is present, it receives the command (blank-padded or truncated as appropriate). If `LENGTH` is present, it receives the length of the command. If `STATUS` is present, it is set to `-1` if `COMMAND` is too short to hold the whole command, a positive number if the command cannot be retrieved, and zero otherwise.

```
SUBROUTINE get_command_argument(number,value,length,status)
  INTEGER, INTENT(IN) :: number
  CHARACTER(*), INTENT(OUT), OPTIONAL :: value
  INTEGER, INTENT(OUT), OPTIONAL :: length, status
```

Accesses command-line argument number `NUMBER`, where argument zero is the program name. If `VALUE` is present, it receives the argument text (blank-padded or truncated as appropriate if the length of the argument differs from that of `VALUE`). If `LENGTH` is present, it receives the length of the argument. If `STATUS` is present, it is set to zero for success, `-1` if `VALUE` is too short, and a positive number if an error occurs.

Note that it is an error for `NUMBER` to be less than zero or greater than the number of arguments (returned by `COMMAND_ARGUMENT_COUNT`).

71.14 Access to environment variables [5.1]

The intrinsic procedure `GET_ENVIRONMENT_VARIABLE` has been added. This duplicates the functionality previously only available via the procedure `GETENV` in the `F90_UNIX_ENV` module.

```
SUBROUTINE get_environment_variable(name,value,length,status,trim_name)
  CHARACTER(*),INTENT(IN) :: name
  CHARACTER(*),INTENT(OUT),OPTIONAL :: value
  INTEGER,INTENT(OUT),OPTIONAL :: length,status
  LOGICAL,INTENT(IN),OPTIONAL :: trim_name
END
```

Accesses the environment variable named by **NAME**; trailing blanks in **NAME** are ignored unless **TRIM_NAME** is present with the value **.FALSE..** If **VALUE** is present, it receives the text value of the variable (blank-padded or truncated as appropriate if the length of the value differs from that of **VALUE**). If **LENGTH** is present, it receives the length of the value. If **STATUS** is present, it is assigned the value 1 if the environment variable does not exist, **-1** if **VALUE** is too short, and zero for success. Other positive values might be assigned for unusual error conditions.

71.15 Character kind selection [5.1]

The intrinsic function **SELECTED_CHAR_KIND** has been added. At this time the only character set supported is **'ASCII'**.

71.16 Argument passing relaxation [5.1]

A **CHARACTER** scalar actual argument may now be passed to a routine which expects to receive a **CHARACTER** array, provided the array is explicit-shape or assumed-size (i.e. not assumed-shape, allocatable, or pointer). This is useful for C interoperability.

71.17 The **MAXLOC** and **MINLOC** intrinsic functions [5.1]

The **MAXLOC** and **MINLOC** intrinsic functions now return zeroes for empty set locations, as required by Fortran 2003 (Fortran 95 left this result processor-dependent).

71.18 The **VALUE** attribute [4.x]

The **VALUE** attribute specifies that an argument should be passed by value.

71.18.1 Syntax

The **VALUE** attribute may be specified by the **VALUE** statement or with the **VALUE** keyword in a type declaration statement.

The syntax of the **VALUE** statement is:

```
VALUE [ :: ] name [ , name ] ...
```

The **VALUE** attribute may only be specified for a scalar dummy argument.

Procedures with a **VALUE** dummy argument must have an explicit interface.

71.18.2 Semantics

A dummy argument with the **VALUE** attribute is “passed by value”; this means that a local copy is made of the argument on entry to the routine and so modifications to the dummy argument do not affect the associated actual argument and vice versa.

A **VALUE** dummy argument may be **INTENT(IN)** but cannot be **INTENT(INOUT)** or **INTENT(OUT)**.

71.18.3 Example

```
PROGRAM value_example
  INTEGER :: i = 3
  CALL s(i)
  PRINT *,i ! This will print the value 3
CONTAINS
  SUBROUTINE s(j)
    INTEGER,VALUE :: j
    j = j + 1 ! This changes the local J without affecting the actual argument
    PRINT *,j ! This will print the value 4
  END SUBROUTINE
END
```

This example is not intended to be particularly useful, just to illustrate the functionality.

71.19 The VOLATILE attribute [5.0]

This is a horrible attribute which specifies that a variable can be modified by means outside of Fortran. Its semantics are basically the same as that of the C ‘volatile’ type qualifier; essentially it disables optimisation for access to that variable.

71.20 Enhanced complex constants [5.2]

The real or imaginary part may now be a named constant, it is not limited to being a literal constant. For example:

```
REAL,PARAMETER :: minusone = -1.0
COMPLEX,PARAMETER :: c = (0,minusone)
```

This is not particularly useful, since the same effect can be achieved by using the `CMPLX` intrinsic function.

71.21 The ASSOCIATE construct [5.2]

The `ASSOCIATE` construct establishes a temporary association between the “associate names” and the specified variables or values, during execution of a block. Its syntax is

```
ASSOCIATE ( association [ , association ]... )
  block
END ASSOCIATE
```

where *block* is a sequence of executable statements and constructs, and *association* is one of

```
name => expression
name => variable
name
```

The last of those is short for ‘*name* => *name*’. The scope of each “associate name” is the *block* of the `ASSOCIATE` construct. An associate name is never allocatable or a pointer, but otherwise has the same attributes as the *variable* or *expression* (and it has the `TARGET` attribute if the *variable* or *expression* is a pointer). If it is being associated with an expression, the expression is evaluated on execution of the `ASSOCIATE` statement and its value does not change during execution of the *block* — in this case, the associate name is not permitted to appear on the left-hand-side of an assignment or any other context which might change its value. If it is being associated with a variable, the associate name can be treated as a variable.

The type of the associate name is that of the expression or variable with which it is associated. For example, in

```
ASSOCIATE(zoom=>NINT(SQRT(a+b)), alt=>state%mapval(:,i)%altitude)
  alt%x = alt%x*zoom
  alt%y = alt%y*zoom
END ASSOCIATE
```

ALT is associated with a variable and therefore can be modified whereas ZOOM cannot. The expression for ZOOM is of type INTEGER and therefore ZOOM is also of type INTEGER.

71.22 Binary, octal and hexadecimal constants [5.2]

In Fortran 95 these were restricted to DATA statements, but in Fortran 2003 these are now allowed to be arguments of the intrinsic functions CMPLX, DBLE, INT and REAL. The interpretation is processor-dependent, but the intent is that this specifies the internal representation of the complex or real value. The NAG Fortran compiler requires these constants to have the correct length for the specified kind of complex or real, viz 32 or 64 bits as appropriate.

For example, on a machine where default REAL is IEEE single precision,

```
REAL(z"41280000")
```

has the value 10.5.

71.23 Character sets [5.1; 5.3]

The support for multiple character sets, especially for Unicode (ISO 10646) has been improved.

The default character set is now required to include lowercase letters and all the 7-bit ASCII printable characters.

The ENCODING= specifier for the OPEN and INQUIRE statements is described in the input/output section.

A new intrinsic function `SELECTED_CHAR_KIND(NAME)` has been added: this returns the character kind for the named character set, or `-1` if there is no kind for that character set. Standard character set names are 'DEFAULT' for the default character kind, 'ASCII' for the 7-bit ASCII character set and 'ISO_10646' for the UCS-4 (32-bit Unicode) character set. The name is not case-sensitive. Note that although the method of requesting UCS-4 characters is standardised, the compiler is not required to support them (in which case `-1` will be returned); the NAG Fortran Compiler supports UCS-4 in release 5.3 (as well as UCS-2 and JIS X 0213).

Assignment of a character value of one kind to a character value of a different kind is permitted if each kind is one of default character, ASCII character, or UCS-4 character. Assignment to and from a UCS-4 character variable preserves the original value.

Internal file input/output to variables of UCS-4 character kind is allowed (if the kind exists), including numeric conversions (e.g. the E edit descriptor), and conversions from/to default character and ASCII character. Similarly, writing default character, ASCII character and UCS-4 character values to a UTF-8 file and reading them back is permitted and preserves the value.

Finally, the intrinsic function `IACHAR` (for converting characters to the ASCII character set) accepts characters of any kind (in Fortran 95 it only accepted default kind).

71.24 Intrinsic function changes for 64-bit machines [5.2]

Especially to support machines with greater than 32-bit address spaces, but with 32-bit default integers, several intrinsic functions now all have an optional `KIND` argument at the end of the argument list, to specify the kind of integer they return. The functions are: `COUNT`, `INDEX`, `LBOUND`, `LEN`, `LEN_TRIM`, `SCAN`, `SHAPE`, `SIZE`, `UBOUND` and `VERIFY`.

71.25 Miscellaneous intrinsic procedure changes [5.2]

The intrinsic subroutine `DATE_AND_TIME` no longer requires the three character arguments (`DATE`, `TIME` and `ZONE`) to have a minimum length: if the actual argument is too small, it merely truncates the value assigned.

The intrinsic functions `IACHAR` and `ICHAR` now accept an optional `KIND` argument to specify the kind of integer to which to convert the character value. This serves no useful purpose since there are no character sets with characters bigger than 32 bits.

The intrinsic functions `MAX`, `MAXLOC`, `MAXVAL`, `MIN`, `MINLOC` and `MINVAL` all now accept character values; the comparison used is the native (`.LT.`) one, not the ASCII (`LLT`) one.

The intrinsic subroutine `SYSTEM_CLOCK` now accepts a `COUNT_RATE` argument of type real; this is to handle systems whose clock ticks are not an integral divisor of 1 second.

Fortran 2008 Extensions

72 Fortran 2008 Overview

This part of the manual describes those parts of the Fortran 2008 language which are not in Fortran 2003, and which are now supported by Release 5.3 of the NAG Fortran Compiler.

The new features of Fortran 2008 that are supported by the NAG Fortran Compiler can be grouped as follows:

- the **BLOCK** construct,
- additional intrinsic functions for bit manipulation,
- **ISO_FORTRAN_ENV** additions,
- miscellaneous and convenience features.

73 The BLOCK construct

This construct allows declarations of entities within executable code. For example,

```
Do i=1,n
  Block
    Real tmp
    tmp = a(i)**3
    If (tmp>b(i)) b(i) = tmp
  End Block
End Do
```

Here the variable **tmp** has its scope limited to the **BLOCK** construct, so will not affect anything outside it. This is particularly useful when including code by **INCLUDE** or by macro preprocessing.

All declarations are allowed within a **BLOCK** construct except for **COMMON**, **EQUIVALENCE**, **IMPLICIT**, **INTENT**, **NAMelist**, **OPTIONAL** and **VALUE**; also, statement function definitions are not permitted.

BLOCK constructs may be nested; like other constructs, branches into a **BLOCK** construct from outside are not permitted. A branch out of a **BLOCK** construct “completes” execution of the construct.

Entities within a **BLOCK** construct that do not have the **SAVE** attribute (including implicitly via initialisation), will cease to exist when execution of the construct is completed. For example, an allocated **ALLOCATABLE** variable will be automatically deallocated, and a variable with a **FINAL** procedure will be finalised.

74 Additional intrinsic functions for bit manipulation

- The elemental intrinsic functions **BGE**, **BGT**, **BLE** and **BLT** perform bitwise (i.e. unsigned) comparisons. They each have two arguments, **I** and **J**, which must be of type Integer but may be of different kind. The result is default Logical.

For example, **BGE(INT(Z'FF',INT8),128)** is true, while **INT(Z'FF',INT8)>=128** is false.

- The array reduction intrinsic functions **IALL**, **IANY** and **IPARITY** reduce arrays using bitwise operations. These are exactly the same as **SUM** and **PRODUCT**, except that instead of reducing the array by the **+** or ***** operation, they reduce it by the **IAND**, **IOR** and **IEOR** intrinsic functions respectively. That is, each element of the result is the bitwise-and, bitwise-or, or bitwise-exclusive-or of the reduced elements. If the number of reduced elements is zero, the result is zero for **IANY** and **IPARITY**, and **NOT(zero)** for **IALL**.

- The elemental intrinsic functions `LEADZ` and `TRAILZ` return the number of leading (most significant) and trailing (least significant) zero bits in the argument `I`, which must be of type Integer (of any kind). The result is default Integer.
- The elemental intrinsic functions `MASKL` and `MASKR` generate simple left-justified and right-justified bitmasks. The value of `MASKL(I,KIND)` is an integer with the specified kind that has its leftmost `I` bits set to one and the rest set to zero; `I` must be non-negative and less than or equal to the bitsize of the result. If `KIND` is omitted, the result is default integer. The value of `MASKR` is similar, but has its rightmost `I` bits set to one instead.
- The array reduction intrinsic function `PARITY` reduces Logical arrays. It is exactly the same as `ALL` and `ANY`, except that instead of reducing the array by the `.AND.` or `.OR.` operation, it reduces it by the `.NEQV.` operation. That is, each element of the result is `.TRUE.` if an odd number of reduced elements is `.TRUE.`.
- The elemental intrinsic function `POPCNT(I)` returns the number of bits in the Integer argument `I` that are set to 1. The elemental intrinsic function `POPPAR(I)` returns zero if the number of bits in `I` that are set to 1 are even, and one if it is odd. The result is default Integer.

75 ISO_FORTRAN_ENV additions

The intrinsic module `ISO_FORTRAN_ENV` contains additional named constants as follows.

- The additional scalar integer constants `INT8`, `INT16`, `INT32`, `INT64`, `REAL32`, `REAL64` and `REAL128` supply the kind type parameter values for integer and real kinds with the indicated bit sizes.
- The additional named array constants `CHARACTER_KINDS`, `INTEGER_KINDS`, `LOGICAL_KINDS` and `REAL_KINDS` list the available kind type parameter values for each type (in no particular order).

76 Miscellaneous and convenience features

- In a structure constructor, the value for an allocatable component may be omitted: this has the same effect as specifying `NULL()`.
- In a `STOP` statement, the *stop-code* may be any scalar constant expression of type integer or default character. (In the NAG Fortran Compiler this also applies to the `PAUSE` statement, but that statement is no longer standard Fortran.)
- `ENTRY` statements are regarded as obsolescent.
- An empty internal subprogram part, module subprogram part or type-bound procedure part is now permitted following a `CONTAINS` statement. In the case of the type-bound procedure part, an ineffectual `PRIVATE` statement may appear following the unnecessary `CONTAINS` statement.
- A type-bound procedure declaration statement may now declare multiple type-bound procedures. For example, instead of

```
PROCEDURE,NOPASS :: a
PROCEDURE,NOPASS :: b=>x
PROCEDURE,NOPASS :: c
```

the single statement

```
PROCEDURE,NOPASS :: a, b=>x, c
```

will suffice.

- The `NEWUNIT=` specifier has been added to the `OPEN` statement; this allocates a new unit number that cannot clash with any other logical unit (the value will be a special negative value). For example,

```
INTEGER unit
OPEN(FILE='output.log',FORM='FORMATTED',NEWUNIT=unit)
WRITE(unit,*) 'Logfile opened.'
```

The `NEWUNIT=` specifier can only be used if either the `FILE=` specifier is also used, or if the `STATUS=` specifier is used with the value `'SCRATCH'`.

- Fortran 2008 extends the rules that are used for generic resolution and for checking that procedures in a generic are unambiguous. The new rules are that
 - a dummy procedure is distinguishable from a dummy variable;
 - an `ALLOCATABLE` dummy variable is distinguishable from a `POINTER` dummy variable that does not have `INTENT(IN)`.

Appendices

77 Mixing Fortran and C

When mixing Fortran source code with C source code, without using the C interoperability features of Fortran 2003, the following points should be noted.

77.1 Naming Conventions

External procedure names and common block names are in lower case with an underscore appended (this is the same as the standard UNIX f77). The main program-unit is called `'main'`.

Module variables and module procedure names are formed by taking the module name (in lower case), appending `'_MP_'` and then appending the entity's name (in lower case).

These conventions differ when the `—compatible` option is used. On Windows this causes external procedures and common blocks to be in upper case with no trailing underscore, while module entities have the module name in upper case instead of lower case (64-bit Windows is always `—compatible`). On most other platforms this option adds an extra underscore (after the one that is always added) for names that already had an underscore.

77.2 Initialisation and Termination

If the main program is not written in Fortran then either the initialisation routine

```
f90_init(int argc,char*argv[])
```

or the i/o initialisation routine

```
f90_io_init(void)
```

should be called from C. The `f90_init` routine initialises the Fortran floating-point environment as well as allowing command-line arguments to be accessed via the Fortran 2003 intrinsic functions and also via the `F90_UNIX` module.

Additionally, the program should be terminated with `f90_finish(int status)` or alternatively `f90_io_finish(void)` may be called before the usual C termination routine to close all Fortran files and flush the Fortran i/o buffers (failing to do this might corrupt an output file that is still open).

77.3 Calling Conventions

The following sections describe in detail the calling conventions used by Fortran programs compiled with the NAG Fortran Compiler in C terms. This information is thus mostly useful to those wishing to mix Fortran and C code.

The conventions used for the Fortran 77 subset of Modern Fortran are compatible with the de facto UNIX f77 conventions (except for `COMPLEX` functions compiled without the `—compatible` option).

77.4 Data Types

Definitions of data types useful in communicating from C to Fortran are in the files `dope.h` and `nagfortran.h`, which are located in the compiler library directory (usually `/usr/local/lib/NAG_Fortran` on UNIX-like systems).

Fortran Data Type	Fortran Precision	C typedef name
INTEGER	8 bits	Integer1
	16 bits	Integer2
	32 bits	Integer <i>or</i> Integer3
	64 bits	Integer4
LOGICAL	8 bits	Logical1
	16 bits	Logical2
	32 bits	Logical
	64 bits	Logical4
REAL	single	Real
	double	Double
	double-double	DDReal
	quadruple	QReal
COMPLEX	single	Complex
	double	DComplex
	double-double	DDComplex
	quadruple	QComplex

Note that DDReal and QReal are the same on most systems; on Sun Solaris these are different (the latter being an IEEE-conformant 128-bit floating-point type).

77.4.1 Pointers

Scalar non-polymorphic non-CHARACTER POINTER types are simply C pointers to the object.

An array POINTER is a dope vector describing the array pointed to. Unlike a simple address, these dope vectors are capable of directly describing non-contiguous array sections of contiguous arrays. See below (DopeN and ChDopeN) for further details. Polymorphic dope vectors are NPDopeN except for CLASS(*) pointers which are CSDopeN.

77.4.2 Derived types

Fortran derived types are translated into C structs. If BIND(C) or SEQUENCE is used, the order of the items within the struct is the same as the order within the derived type definition. Otherwise, the order of items is permuted to put larger types at the front of the struct so as to improve alignment and reduce storage; the C output code can be inspected to determine the ordering used.

77.4.3 Supporting types

Char unsigned char
Data type for default (single-byte) character storage.

Char2 unsigned short
Data type for 16-bit (JIS and UCS-2) character storage.

Char4 unsigned int
Data type for 32-bit (UCS-4) character storage.

Substring

```
struct { Char *addr; ChrLen len;}
```

Describes a single-byte (default) CHARACTER string; used for deferred-length default CHARACTER variables and as the return type of variable-length scalar non-POINTER default CHARACTER functions and all POINTER default CHARACTER functions.

Substring2

```
struct { Char2 *addr; ChrLen len;}
```

Describes a 16-bit (JIS or UCS-2) CHARACTER string; used for deferred-length 16-bit CHARACTER variables

and as the return type of variable-length scalar non-POINTER 16-bit CHARACTER functions and all POINTER 16-bit CHARACTER functions.

Substring4

```
struct { Char4 *addr; ChrLen len;}
```

Describes a 32-bit ISO_10646 (UCS-4) CHARACTER string; used for deferred-length 32-bit CHARACTER variables and as the return type of variable-length scalar non-POINTER UCS-4 CHARACTER functions and all POINTER UCS-4 CHARACTER functions.

Offset `int, long or long long`

An integer type for addressing and subscript calculations; this is `int` (32-bit) on 32-bit systems and small model 64-bit systems, and a 64-bit integer type on large model 64-bit systems.

ChrLen usually `int`, or `long long` on 64-bit Windows.

An integer type for representing character length.

Pointer usually `char *`.

A byte pointer used to refer to any type and for pointer arithmetic.

Triplet `struct { Offset extent, mult, lower; }`

Contains the parameters of an array dimension. `extent` is the size of that dimension, `mult` is the stride (i.e., the distance between successive elements in bytes) and `lower` is the lower bound. It is a component of the `DopeN` and `ChDopeN` structs.

DopeN `struct { Pointer addr; Offset offset; Triplet dim[N]; }`

Dope vectors for all non-polymorphic non-CHARACTER arrays (including arrays of derived type). *N* is the rank and is from 1 to 7. `addr` is the address of the first element, `dim` describes each dimension and `offset` is the offset to add to the subscript calculation, i.e., `SUM(mult*lower)`. This is used as the return type for POINTER array functions; a pointer to it is used as the argument type for assumed-shape and POINTER array arguments.

An array pointer which has been nullified has an `addr` field which is a null pointer; note that zero-sized arrays have an `addr` field which is not a null pointer.

ChDopeN `struct { Pointer addr; ChrLen len; Offset offset; Triplet dim[N]; }`

These are exactly the same as the `DopeN` structs with the addition of the `len` component which specifies the CHARACTER length.

ArrayTemp_type

```
struct { type *addr; Offset extent[7];}
```

Describes a contiguous array of *type*, which is one of: `Integer1`, `Integer2`, `Integer`, `Integer4`, `Logical1`, `Logical2`, `Logical`, `Logical4`, `Real`, `Double`, `QReal`, `Complex`, `DComplex` or `QComplex`. It is used as the return type for non-POINTER array functions. Note that `extent` values after the rank of the array being described are undefined.

ArrayTemp_Character

```
struct { Char *addr; ChrLen len; Offset extent[7];}
```

Describes a contiguous CHARACTER array; it is the same as the other array types with the addition of the `len` component for the CHARACTER length.

ArrayTemp_Derived

synonym for `ArrayTemp_Character`.

Describes a contiguous array of any derived type. The `len` field in this case is the size of the derived type array element.

77.5 SUBROUTINE return types

77.5.1 SUBROUTINEs with label arguments

The return type is `int`; its value is the index of the label to which control is to be transferred (1 for the first label, etc.). Zero or an out-of-range value indicates no control transfer is to take place.

77.5.2 SUBROUTINEs with no label arguments

Return type is void.

77.6 FUNCTION return types

77.6.1 Scalar

1. INTEGER, LOGICAL and REAL.
The intrinsic type as listed above.
2. COMPLEX, *-compatible* option not used.
Complex or DComplex according to the precision.
3. COMPLEX, *-compatible* option used.
Return type is void. The address of a temporary of type Complex or DComplex is passed as the initial argument of the function (the result is written to this location).
4. CHARACTER with fixed or assumed length.
Return type is void. Two additional initial arguments are passed to the function, the first (Char*, Char2* or Char4*) being the address to which the result is to be written and the second (Chrlen) being the length of the result (in case the called function is of assumed length).
5. CHARACTER with variable length.
Return type is Substring, Substring2 or Substring4 (described above). The called function allocates the storage for the string being returned; it is the caller's responsibility to deallocate it when it is no longer required.
6. Derived type.
The derived-type struct.

77.6.2 Scalar POINTER functions

Note that with all POINTER-valued functions the storage to which the result points may have been allocated within the called function or the result may point to pre-existing storage.

1. INTEGER, LOGICAL, REAL and COMPLEX.
A pointer to the appropriate intrinsic type (e.g., Complex*).
2. CHARACTER
Return type is Substring, Substring2 or Substring4.
3. Derived type.
A pointer to the derived-type struct.

77.6.3 Array non-POINTER functions

1. Intrinsic types.
The appropriate ArrayTemp_ struct for the intrinsic type, as described above.
2. Derived types.
ArrayTemp_Derived is returned with the len component set to the size of the derived-type struct.

77.6.4 Array POINTER functions

1. CLASS(*)
CSDope1, CSDope2, ... or CSDope7, depending on the rank of the array.
2. CLASS(*derived-type-name*) NPDope1, NPDope2, ... or NPDope7, depending on the rank of the array.

3. Non-polymorphic non-CHARACTER type.
Dope1, Dope2, ... or Dope7, depending on the rank of the array.
4. CHARACTER.
ChDope1, ChDope2, ... or ChDope7, depending on the rank of the array.

Note that non-polymorphic derived-type arrays are returned as Dope N structs.

77.7 Argument types

77.7.1 CHARACTER type

All normal arguments of CHARACTER type, whether default CHARACTER or multi-byte CHARACTER, have an additional `ChrLen` argument being the length of the CHARACTER entity; this additional argument is passed at the end of the argument list after all the normal arguments. When there are several CHARACTER arguments their lengths are passed in order from left to right.

This is except on 32-bit Windows when the *—compatible* option is specified, the additional argument immediately follows the CHARACTER argument.

The other exception to this rule is for assumed-shape CHARACTER arrays; in this case the length of the dummy argument is taken directly from the field in the dope vector and is not passed separately.

77.7.2 non-POINTER non-ALLOCATABLE Scalar

1. non-CHARACTER type.
The address of the argument is passed (e.g., `Integer*` for an INTEGER argument).
2. CHARACTER.
The address of the argument is passed, and additionally the length of the argument is passed as a separate `ChrLen` argument at the end of the argument list.

77.7.3 POINTER and ALLOCATABLE Scalar

These are passed exactly the same as the normal case except that the address of the pointer (or allocatable descriptor) is passed (e.g., `Integer**` for an INTEGER POINTER).

77.7.4 non-POINTER Array

1. Assumed shape.
The address of an appropriate dope vector describing the array is passed (i.e., a `CSDopeN*`, `NPDopeN*`, `DopeN*` or a `ChDopeN*`, depending on the polymorphism and type of the dummy argument). There is no need for the array to be contiguous as long as it can be described by an array section.
2. Other.
The address of the first element of the array is passed. For CHARACTER arrays the length of each array element is passed as an `ChrLen` at the end of the argument list, the same as for scalars. The array must be contiguous.

77.7.5 POINTER Array

The address of an appropriate dope vector is passed, the same as for assumed-shape arrays. For CHARACTER arrays the length is passed at the end as a separate argument.

77.7.6 Procedures

The address of the procedure is passed. CHARACTER functions have the length passed as a separate argument at the end of the list; if the function is of variable length this length will be negative.

77.7.7 OPTIONAL arguments

If an `OPTIONAL` argument is `.NOT.PRESENT()` a null pointer of the appropriate type is passed (e.g., for an `INTEGER` scalar, an `(Integer*)0` is passed).

78 ASCII Collating Sequence

78.1 Printing Characters

Decimal value and character. The value 32 is a space.

	30	40	50	60	70	80	90	100	110	120
0		(2	<	F	P	Z	d	n	x
1)	3	=	G	Q	[e	o	y
2		*	4	>	H	R	\	f	p	z
3	!	+	5	?	I	S]	g	q	{
4	"	,	6	@	J	T	^	h	r	
5	#	-	7	A	K	U	_	i	s	}
6	\$.	8	B	L	V	'	j	t	~
7	%	/	9	C	M	W	a	k	u	
8	&	0	:	D	N	X	b	l	v	
9	'	1	;	E	O	Y	c	m	w	

Octal value and character. The value 40 is a space.

	40	50	60	70	100	110	120	130	140	150	160	170
0		(0	8	@	H	P	X	'	h	p	x
1	!)	1	9	A	I	Q	Y	a	i	q	y
2	"	*	2	:	B	J	R	Z	b	j	r	z
3	#	+	3	;	C	K	S	[c	k	s	{
4	\$,	4	<	D	L	T	\	d	l	t	
5	%	-	5	=	E	M	U]	e	m	u	}
6	&	.	6	>	F	N	V	^	f	n	v	~
7	'	/	7	?	G	O	W	_	g	o	w	

78.2 Non-printing Characters

Decimal	Octal	Mnemonic	Control	Description
0	000	NUL	ctrl/@	Null character
1	001	SOH	ctrl/A	Start of heading
2	002	STX	ctrl/B	Start of text
3	003	ETX	ctrl/C	End of text
4	004	EOT	ctrl/D	End of transmission
5	005	ENQ	ctrl/E	Enquire
6	006	ACK	ctrl/F	Acknowledge
7	007	BEL	ctrl/G	Bell
8	010	BS	ctrl/H	Backspace
9	011	HT	ctrl/I	Horizontal tab
10	012	LF	ctrl/J	Line feed
11	013	VT	ctrl/K	Vertical tab
12	014	FF	ctrl/L	Form feed
13	015	CR	ctrl/M	Carriage return
14	016	SO	ctrl/N	Shift out
15	017	SI	ctrl/O	Shift in
16	020	DLE	ctrl/P	Data link escape
17	021	DC1	ctrl/Q	Device control 1 (XON)
18	022	DC2	ctrl/R	Device control 2
19	023	DC3	ctrl/S	Device control 3 (XOFF)
20	024	DC4	ctrl/T	Device control 4
21	025	NAK	ctrl/U	Negative acknowledge
22	026	SYN	ctrl/V	Synchronise
23	027	ETB	ctrl/W	End transmission of block
24	030	CAN	ctrl/X	Cancel
25	031	EM	ctrl/Y	End of medium
26	032	EOF	ctrl/Z	End-of-file
27	033	ESC	ctrl/[Escape
28	034	FS	ctrl/\	File separator
29	035	GS	ctrl/]	Group separator
30	036	RS	ctrl/^	Record separator
31	037	US	ctrl/_	Unit separator
127	177	DEL		Delete

Detailed Contents

Contents

1	Introduction to the Compiler	1
1.1	Other Fortran-related Activities at NAG	1
1.2	This Manual	1
2	Usage	2
3	Description	2
4	File Types	2
5	Compiler Options	2
6	Files	10
7	Compilation Messages	11
8	Compiler Limits	12
9	Input/Output Information	12
10	OpenMP Support	12
11	Automatic File Preconnection	14
12	IEEE 754 Arithmetic Support	14
13	Random Number Algorithm	14
14	Automatic Garbage Collection	14
15	Memory Tracing	15
16	Undefined Variable Detection	15
17	Data Types	16
18	Modules	16
19	Runtime Environment Variables	17
20	Debugging	17
21	Producing a Call Graph	17

22	Dependency Analysis	18
23	Source File Polishing	18
24	dbx90 command line	22
25	Description of dbx90	22
26	dbx90 options	22
27	dbx90 Commands	22
28	dbx90 Expressions	24
28.1	Scalar expressions	24
28.2	Array sections	24
28.3	Derived type component specification	24
29	dbx90 Command aliases	25
30	dbx90 limitations	25
31	Example of dbx90	25
32	Troubleshooting dbx90	27
33	Overview of fpp	28
34	Usage	28
35	Description	28
36	Options	28
37	Using fpp	29
37.1	Source files	29
37.2	Output	29
37.3	Directives	30
37.4	Macro definition	30
37.5	Including external files	30
37.6	Line number control	30
37.7	Conditional selection of source text	31
38	Details	31
38.1	Scope of macro or variable definitions	31
38.2	End of macro definition	32

38.3	Function-like macro definition	32
38.4	Cancelling macro definitions	32
38.5	Conditional source code selection	32
38.6	Including external files	33
38.7	Comments	33
38.8	Macro functions	34
38.9	Macro expression	34
39	Diagnostics	34
40	Non-standard Extensions	36
40.1	BOZ literal constants outside DATA statements	36
40.2	High Performance Fortran	36
40.2.1	HPF Directives	36
40.2.2	HPF intrinsic functions	36
40.2.3	EXTRINSIC procedures	36
40.3	Longer Names	37
40.4	Dollar Sign in Names	37
40.5	Input/output endian/format conversion	37
40.6	Elemental BIND(C) procedures	37
41	Obsolete Extensions	37
41.1	Byte Sizes	38
41.2	TAB Format	38
41.3	Hollerith Constants	38
41.4	Increased Line Length in Fixed Source Form	38
41.5	Increased Maximum Number of Continuation Lines	39
41.6	Legacy Application Support	39
41.7	Mismatched Argument Lists	39
41.8	Double Precision Complex Extensions	39
42	Intrinsic Module Overview	40
43	f90_gc	40
44	f90_iostat	42
45	f90_kind	43
46	f90_preconn_io	44
46.1	Procedures	44

46.2 Example	45
47 f90_stat	45
47.1 Parameters	45
47.2 Example	46
48 f90_unix_*	46
49 ieee_*, iso_c_binding, iso_fortran_env	46
50 hpf_library	46
51 Posix Module Overview	47
52 f90_unix_dir	47
52.1 Parameters	47
52.2 Procedures	47
53 f90_unix_dirent	49
53.1 Procedures	49
54 f90_unix_env	50
54.1 Parameters	50
54.2 Types	51
54.3 Procedures	52
55 f90_unix_errno	55
55.1 Error Handling	55
55.2 Parameters	55
56 f90_unix_file	56
56.1 Parameters	56
56.2 Types	58
56.3 Procedures	59
57 f90_unix_io	61
57.1 Procedures	61
58 f90_unix_proc	61
58.1 Parameters	61
58.2 Procedures	62
59 HPF_LIBRARY Module Summary	67

59.1	Elemental Functions	67
59.2	Reduction Functions	67
59.3	Scan Functions	67
59.4	Scatter Functions	68
59.5	Sorting Functions	68
59.6	Enquiry Functions	69
60	Fortran 95 Program Structure	70
61	Fortran 95 Expressions	73
62	Fortran 95 Statements	77
63	Fortran 95 Intrinsic Procedures	93
64	Fortran 2003 Overview	97
65	Object-oriented programming	98
65.1	Type Extension	98
65.1.1	Extending Types [5.0]	98
65.1.2	Polymorphic Variables [5.0]	98
65.1.3	Type Selection [5.0]	99
65.1.4	Unlimited polymorphism [5.2]	99
65.1.5	Ad hoc type comparison [5.3]	100
65.2	Typed allocation [5.1]	100
65.3	Sourced allocation (cloning) [5.1]	100
65.4	Type-bound procedures [5.1]	101
65.4.1	The type-bound procedure part	101
65.4.2	Specific type-bound procedures	101
65.4.3	Generic type-bound procedures	102
65.5	Abstract derived types [5.1]	102
65.6	Object-bound procedures [5.2]	103
66	ALLOCATABLE extensions	103
66.1	Allocatable Dummy Arrays [4.x]	103
66.2	Allocatable Function Results [4.x]	104
66.3	Allocatable Structure Components [4.x]	105
66.4	Allocatable Component Example	107
66.5	The MOVE_ALLOC intrinsic subroutine [5.2]	108
66.6	Allocatable scalars [5.2]	109
66.7	Automatic reallocation [5.2]	109

67 Other data-oriented enhancements	109
67.1 Parameterised derived types [n/a]	109
67.1.1 Basic Syntax and Semantics	110
67.1.2 More Semantics	111
67.1.3 Assumed type parameters	111
67.1.4 Deferred type parameters	111
67.2 Finalisation [5.3]	112
67.3 The PROTECTED attribute [5.0]	113
67.3.1 Syntax	113
67.3.2 Semantics	113
67.3.3 Example	113
67.4 Pointer enhancements	114
67.4.1 INTENT for pointers [5.1]	114
67.4.2 Pointer bounds specification [5.2]	114
67.4.3 Rank-remapping Pointer Assignment [5.0]	114
67.5 Individual component accessibility [5.1]	114
67.6 Public entities of private type [5.1]	115
68 C interoperability [mostly 5.1]	115
68.1 The ISO_C_BINDING module	115
68.1.1 The kind parameters	115
68.1.2 Using C_PTR and C_FUNPTR	116
68.2 BIND(C) types	117
68.3 BIND(C) variables	117
68.4 BIND(C) procedures	117
68.5 Enumerations	118
69 IEEE arithmetic support [4.x except as otherwise noted]	118
69.1 Introduction	118
69.2 Exception flags, modes and information flow	119
69.3 Procedures in the modules	119
69.4 The IEEE_FEATURES module	119
69.5 IEEE_EXCEPTIONS	120
69.5.1 Types and constants	120
69.5.2 Procedures	120
69.6 IEEE_ARITHMETIC module	121
69.6.1 IEEE datatype selection	121
69.6.2 Enquiry functions	121

69.6.3 Rounding mode	123
69.6.4 Underflow mode	123
69.6.5 Number Classification	124
69.6.6 Test functions	124
69.6.7 Arithmetic functions	125
70 Input/output Features	126
70.1 Stream input/output [5.1]	126
70.2 The BLANK= and PAD= specifiers [5.1]	126
70.3 Decimal Comma [5.1]	126
70.4 The DELIM= specifier [5.1]	127
70.5 The ENCODING= specifier [5.1]	127
70.6 The IOMSG= specifier [5.1]	127
70.7 The IOSTAT= specifier [5.1]	127
70.8 The SIGN= specifier [5.1]	127
70.9 Intrinsic functions for testing IOSTAT= values [5.1]	127
70.10 Input/output of IEEE infinities and NaNs [5.1]	127
70.11 Output of floating-point zero [5.1]	128
70.12 NAMELIST and internal files [5.1]	128
70.13 Recursive input/output [5.2]	128
70.14 Asynchronous input/output	128
70.14.1 Basic syntax [5.1]	128
70.14.2 Basic Example	129
70.14.3 The ASYNCHRONOUS attribute [5.2]	129
70.14.4 The WAIT statement [5.2]	129
70.14.5 Execution Semantics	130
70.15 Scale factor followed by repeat count [5.1]	130
70.16 FLUSH statement [5.2]	130
70.17 Defined input/output [n/a]	131
71 Miscellaneous Fortran 2003 Features	132
71.1 Abstract interfaces and the PROCEDURE statement [5.1]	132
71.2 Named procedure pointers [5.2]	133
71.3 Intrinsic modules [4.x]	133
71.4 Renaming user-defined operators on the USE statement [5.2]	134
71.5 The ISO_FORTRAN_ENV module [5.1]	134
71.6 The IMPORT statement [5.1]	135
71.7 Length of names and statements	135

71.8 Array constructor syntax enhancements	135
71.9 Structure constructor syntax enhancements [5.3]	135
71.10 Deferred character length [5.2]	136
71.11 The ERRMSG= specifier [5.1]	136
71.12 Intrinsic functions in constant expressions [5.2 partial; 5.3 complete]	137
71.13 Access to the command line [5.1]	137
71.14 Access to environment variables [5.1]	137
71.15 Character kind selection [5.1]	138
71.16 Argument passing relaxation [5.1]	138
71.17 The MAXLOC and MINLOC intrinsic functions [5.1]	138
71.18 The VALUE attribute [4.x]	138
71.18.1 Syntax	138
71.18.2 Semantics	138
71.18.3 Example	139
71.19 The VOLATILE attribute [5.0]	139
71.20 Enhanced complex constants [5.2]	139
71.21 The ASSOCIATE construct [5.2]	139
71.22 Binary, octal and hexadecimal constants [5.2]	140
71.23 Character sets [5.1; 5.3]	140
71.24 Intrinsic function changes for 64-bit machines [5.2]	140
71.25 Miscellaneous intrinsic procedure changes [5.2]	140
72 Fortran 2008 Overview	142
73 The BLOCK construct	142
74 Additional intrinsic functions for bit manipulation	142
75 ISO_FORTRAN_ENV additions	143
76 Miscellaneous and convenience features	143
77 Mixing Fortran and C	145
77.1 Naming Conventions	145
77.2 Initialisation and Termination	145
77.3 Calling Conventions	145
77.4 Data Types	145
77.4.1 Pointers	146
77.4.2 Derived types	146
77.4.3 Supporting types	146

77.5 SUBROUTINE return types	147
77.5.1 SUBROUTINEs with label arguments	147
77.5.2 SUBROUTINEs with no label arguments	147
77.6 FUNCTION return types	147
77.6.1 Scalar	147
77.6.2 Scalar POINTER functions	148
77.6.3 Array non-POINTER functions	148
77.6.4 Array POINTER functions	148
77.7 Argument types	148
77.7.1 CHARACTER type	148
77.7.2 non-POINTER Scalar	148
77.7.3 POINTER Scalar	149
77.7.4 non-POINTER Array	149
77.7.5 POINTER Array	149
77.7.6 Procedures	149
77.7.7 OPTIONAL arguments	149
78 ASCII Collating Sequence	150
78.1 Printing Characters	150
78.2 Non-printing Characters	151