



Intel® Fortran Compiler for Linux\* Systems  
User's Guide  
Volume II: Optimizing Applications

[Legal Information](#)

Copyright © 2003 Intel Corporation

Portions © Copyright 2001 Hewlett-Packard Development Company, L.P.

## Disclaimer and Legal Information

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

This User's Guide Volume II as well as the software described in it is furnished under license and may only be used or copied in accordance with the terms of the license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation. Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The software described in this User's Guide Volume II may contain software defects which may cause the product to deviate from published specifications. Current characterized software defects are available on request.

Intel SpeedStep, Intel Thread Checker, Celeron, Dialogic, i386, i486, iCOMP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetStructure, Intel Xeon, Intel XScale, Itanium, MMX, MMX logo, Pentium, Pentium II Xeon, Pentium III Xeon, Pentium M, and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\* Other names and brands may be claimed as the property of others.

Copyright © Intel Corporation 2003.

Portions © Copyright 2001 Hewlett-Packard Development Company, L.P.

# Table Of Contents

Disclaimer and Legal Information .....	ii
What's New in This Release .....	1
Improvements and New Optimization in This Release .....	1
Introduction to Volume II .....	2
The Subjects Covered.....	3
Notations and Conventions .....	4
Programming for High Performance.....	5
Programming for High Performance: Overview.....	5
Programming Guidelines.....	5
Analyzing and Timing Your Application .....	30
Compiler Optimizations .....	34
Compiler Optimizations Overview.....	34
Optimizing Compilation Process .....	34
Optimizing Different Application Types.....	57
Floating-point Arithmetic Optimizations .....	61
Optimizing for Specific Processors .....	69
Interprocedural Optimizations (IPO) .....	75
Profile-guided Optimizations .....	88
High-level Language Optimizations (HLO) .....	116
Parallel Programming with Intel® Fortran .....	120
Parallelism: an Overview.....	120

Auto-vectorization (IA-32 Only) .....	124
Auto-parallelization .....	137
Parallelization with OpenMP* .....	144
Debugging Multithreaded Programs .....	190
Optimization Support Features .....	199
Optimization Support Features Overview .....	199
Compiler Directives .....	199
Optimizations and Debugging .....	207
Optimizer Report Generation .....	209
Index .....	213

# What's New in This Release

This volume focuses on the coding techniques and compiler optimizations that make your application more efficient.

## Improvements and New Optimization in This Release

This document provides information about Intel® Fortran Compiler for IA-32-based applications and Itanium®-based applications. IA-32-based applications run on any processor of the Intel® Pentium® processor family generations, including the Intel® Xeon(TM) processor and Intel® Pentium® M processor. Itanium-based applications run on the Intel® Itanium® processor family.

The Intel Fortran Compiler has many options that provide high application performance. In this release, the Intel Fortran Compiler supports most Compaq\* Visual Fortran (CVF) options. Some of the CVF options are supported as synonyms for the Intel Fortran Compiler back-end optimizations. For a complete list of new options in this release, see New Compiler Options in the Intel Fortran Compiler Options Quick Reference Guide.

### Note

Please refer to the Release Notes for the most current information about features implemented in this release.

## New Processors Support

The new options `-xN`, `-xB`, and `-xP` support Intel® Pentium® 4 processors, Intel Pentium M processors Intel® Pentium® M processors and Intel processors code-named "Prescott," respectively. Correspondingly, the options `-axN`, `-axB`, and `-axP` optimize for Intel Pentium® 4 processors, Intel® Xeon™ processors, Intel® Pentium® M processors, and Intel processors code-named "Prescott" (new processor).

## Optimizing for Specific Processors at Run-time, IA-32 Systems

This release enhances processor-specific optimizations for IA-32 systems by performing the following run-time checks to determine the processor on which the program is running and:

- verify whether that processor supports the features of the new options `-xN`, `-xB`, and `-xP`
- set the appropriate flush-to-zero (FTZ) and denormals-are-zero (DAZ) flags

See Processor-specific Run-time Checks, IA-32 Systems for more details.

## Symbol Visibility Attribute Options

The Intel Fortran Compiler has the visibility attribute options that provide command-line control of the visibility attributes as well as a source syntax to set the complete range of these attributes. The options ensure immediate access to the feature without depending on header file modifications. The visibility options cause all global symbols to get the visibility specified by the option.

## IPO Functionality

Automatic generation and update of the intermediate language (.i1) and compiler files is part of the compilation process. You can create a library that retain versioned .i1 files and use them in IPO compilations. The compiler can extract the .i1 files from the library and use them to optimize the program.

## New Directive for Auto-vectorization

Added extended optimization directive `!DEC$ VECTOR NONTEMPORAL`.

## Miscellaneous

IA-32 option `-fpstkchk` checks whether a program makes a correct call to a function that should return a floating-point value. Marks the incorrect call and makes it easy to find the error.

The `-align keyword` option provides more data alignment control with additional keywords.

## Introduction to Volume II

This is the second volume in a two-volume Intel® Fortran Compiler User's Guide. It explains how you can use the Intel Fortran Compiler to enhance your application.

The variety of optimizations used by the Intel Fortran Compiler enables you to enhance the performance of your application. Each optimization is performed by a set of options discussed in the sections of this volume.

In addition to optimizations invoked by the compiler command line options, the compiler includes features which enhance your application performance such as directives, intrinsics, run-time library routines and various utilities. These features are discussed in the Optimization Support Features section.

 **Note**

This document explains how information and instructions apply differently to a targeted architecture, IA-32 or Itanium® architecture. If there is no specific reference to either architecture, the description applies to both architectures.

This documentation assumes that you are familiar with the Fortran Standard programming language and with the Intel® processor architecture. You should also be familiar with the host computer's operating system.

## The Subjects Covered

Programming for high performance by using the specifics of Intel Fortran:

- Setting Data Type and Alignment
- Using Arrays Efficiently
- Improving I/O Performance
- Improving Run-time Efficiency
- Coding for Intel Architectures

Implementing Intel Fortran Compiler optimizations

- Optimizing Compilation Process
- Options to Optimize Different Application Types
- Floating Point Arithmetic Optimizations
- Optimizing for Specific Processors
- Interprocedural Optimizations
- Profile-guided Optimizations
- High-level Language Optimizations (HLO)

Parallel Programming with Intel Fortran

- Auto-vectorization (IA-32 Only)
- Auto-parallelization
- Parallelization with OpenMP\*

Optimization Support Features

- Compiler Directives
- Optimizations and Debugging

## Notations and Conventions

This documentation uses the following conventions:

Intel® Fortran (later: Intel Fortran)	The name of the common compiler language supported by the Intel Fortran Compiler for Windows* and Intel Fortran Compiler for Linux* products.
Adobe Acrobat*	An asterisk at the end of a word or name indicates it is a third-party product trademark.
FORTTRAN 77 and later versions of Fortran	The references to the versions of the Fortran language. After FORTRAN 77, the references are Fortran 90 or Fortran 95. The default is "Fortran," which corresponds to all versions.
THIS TYPE STYLE	Statements, keywords, and directives are shown in all uppercase, in a normal font. For example, "add the USE statement..."
<b>This type style</b>	Bold normal font shows menu names, menu items, button names, dialog window names, and other user-interface items.
<b>File &gt; Open</b>	Menu names and menu items joined by a greater than (>) sign indicate a sequence of actions. For example, "Click <b>File &gt; Open</b> " indicates that in the <b>File</b> menu, click <b>Open</b> to perform this action.
<code>ifort</code>	The use of the compiler command in the examples for both IA-32 and Itanium processors is as follows: when there is no usage difference between the two architectures, only one command is given. Whenever there is a difference in usage, the commands for each architecture are given.
<code>This type style</code>	An element of syntax, a reserved word, a keyword, a file name, a variable, or a code example. The text appears in lowercase unless uppercase is required.
THIS TYPE STYLE	Fortran source text or syntax element.
<b>This type style</b>	Indicates what you type as command or input.
<i>This type style</i>	Command line arguments and option arguments you enter.
<i>This type style</i>	Indicates an argument on a command line or an option's argument in the text.
[options]	Indicates that the items enclosed in brackets are optional.
{value   value}	A value separated by a vertical bar ( ) indicates a version of an option.
...	Ellipses in the code examples indicate that part of the code is not shown.



# Programming for High Performance

## Programming for High Performance: Overview

This section consists of two sub-sections: Programming Guidelines and Analyzing and Timing Your Application. The first one discusses the programming guidelines geared to enhance application performance, including the specific coding practices to utilize the Intel® architecture features. The second discusses how to use the Intel performance analysis tools and how to time the program execution to collect information about the problem areas.

The correlation between the programming practices and related compiler options is explained and the related topics are linked.

## Programming Guidelines

### Setting Data Type and Alignment

Alignment of data concerns these kinds of variables:

- dynamically allocated
- members of a data structure
- global or local variables
- parameters passed on the stack.

For best performance, align data as follows:

- 8-bit data at any address
- 16-bit data to be contained within an aligned four byte word
- 32-bit data so that its base address is a multiple of four
- 64-bit data so that its base address is a multiple of eight
- 80-bit data so that its base address is a multiple of sixteen
- 128-bit data so that its base address is a multiple of sixteen.

### Causes of Unaligned Data and Ensuring Natural Alignment

For optimal performance, make sure your data is aligned naturally. A natural boundary is a memory address that is a multiple of the data item's size. For example, a `REAL (KIND=8)` data item aligned on natural boundaries has an address that is a multiple of 8. An array is aligned on natural boundaries if all of its elements are.

All data items whose starting address is on a natural boundary are **naturally aligned**. Data not aligned on a natural boundary is called **unaligned data**.

Although the Intel Fortran Compiler naturally aligns individual data items when it can, certain Fortran statements (such as `EQUIVALENCE`) can cause data items to become unaligned (see causes of unaligned data below).

You can use the command-line option `-align` to ensure naturally aligned data, but you should check and consider reordering data declarations of data items within common blocks, derived type and record structures as follows:

- carefully specify the order and sizes of data declarations to ensure naturally aligned data
- start with the largest size numeric items first, followed by smaller size numeric items, and then non-numeric (character) data.

Common blocks (`COMMON` statement), derived-type data, and FORTRAN 77 record structures (`RECORD` statement) usually contain multiple items within the context of the larger structure.

The following declaration statements can force data to be unaligned:

- Common blocks (`COMMON` statement)

The order of variables in the `COMMON` statement determines their storage order. Unless you are sure that the data items in the common block will be naturally aligned, specify either the `-align commons` or `-align dcommons` option, depending on the largest data size used. See Alignment Options.

- Derived-type (user-defined) data

Derived-type data members are declared after a `TYPE` statement.

If your data includes derived-type data structures, you should use the `-align records` option, unless you are sure that the data items in derived-type data structures will be naturally aligned.

If you omit the `SEQUENCE` statement, the `-align records` option (default) ensures all data items are naturally aligned.

If you specify the `SEQUENCE` statement, the `-align records` option is prevented from adding necessary padding to avoid unaligned data (data items are packed) unless you specify the `-align sequence` option. When you use `SEQUENCE`, you should specify data declaration order such that all data items are naturally aligned.

- Record structures (`RECORD` and `STRUCTURE` statements)

Intel Fortran record structures usually contain multiple data items. The order of variables in the `STRUCTURE` statement determines their storage order. The `RECORD` statement names the record structure.

If your data includes Intel Fortran record structures, you should use the `-align records` option, unless you are sure that the data items in derived-type data and Intel Fortran record structures will be naturally aligned.

- `EQUIVALENCE` statements

`EQUIVALENCE` statements can force unaligned data or cause data to span natural boundaries. For more information, see the *Intel® Fortran Language Reference Manual*.

To avoid unaligned data in a common block, derived-type data, or record structure (extension), use one or both of the following:

- For new programs or for programs where the source code declarations can be modified easily, plan the order of data declarations with care. For example, you should order variables in a `COMMON` statement such that numeric data is arranged from largest to smallest, followed by any character data (see the data declaration rules in *Ordering Data Declarations to Avoid Unaligned Data* below).
- For existing programs where source code changes are not easily done or for array elements containing derived-type or record structures, you can use command line options to request that the compiler align numeric data by adding padding spaces where needed.

**Other possible causes** of unaligned data include unaligned actual arguments and arrays that contain a derived-type structure or Intel Fortran record structure as detailed below.

- When actual arguments from outside the program unit are not naturally aligned, unaligned data access occurs. Intel Fortran assumes all passed arguments are naturally aligned and has no information at compile time about data that will be introduced by actual arguments during program execution.
- For arrays where each array element contains a derived-type structure or Intel Fortran record structure, the size of the array elements may cause some elements (but not the first) to start on an unaligned boundary.
- Even if the data items are naturally aligned within a derived-type structure without the `SEQUENCE` statement or a record structure, the size of an array element might require use of the Fortran `-align records` option to supply needed padding to avoid some array elements being unaligned.

- If you specify `-align norecords` or specify `-vms` without `-align records`, no padding bytes are added between array elements. If array elements each contain a derived-type structure with the `SEQUENCE` statement, array elements are packed without padding bytes regardless of the Fortran command options specified. In this case, some elements will be unaligned.
- When `-align records` option is in effect, the number of padding bytes added by the compiler for each array element is dependent on the size of the largest data item within the structure. The compiler determines the size of the array elements as an exact multiple of the largest data item in the derived-type structure without the `SEQUENCE` statement or a record structure. The compiler then adds the appropriate number of padding bytes. For instance, if a structure contains an 8-byte floating-point number followed by a 3-byte character variable, each element contains five bytes of padding (16 is an exact multiple of 8). However, if the structure contains one 4-byte floating-point number, one 4-byte integer, followed by a 3-byte character variable, each element would contain one byte of padding (12 is an exact multiple of 4).

### Checking for Inefficient Unaligned Data

During compilation, the Intel Fortran compiler naturally aligns as much data as possible. Exceptions that can result in unaligned data are described above.

Because unaligned data can slow run-time performance, it is worthwhile to:

- Double-check data declarations within common block, derived-type data, or record structures to ensure all data items are naturally aligned (see the data declaration rules in the subsection below). Using modules to contain data declarations can ensure consistent alignment and use of such data.
- Avoid the `EQUIVALENCE` statement or use it in a manner that cannot cause unaligned data or data spanning natural boundaries.
- Ensure that passed arguments from outside the program unit are naturally aligned.
- Check that the size of array elements containing at least one derived-type data or record structure (extension) cause array elements to start on aligned boundaries (see the previous subsection).
- There are two ways unaligned data might be reported:
- During compilation, warning messages are issued for any data items that are known to be unaligned (unless you specify the `-warn noalignments` (`-W0`) option that suppresses all warnings).
- During program execution, warning messages are issued for any data that is detected as unaligned. The message includes the address of the unaligned access. You can use the EDB debugger to locate unaligned data.

The following run-time message shows that:

- The statement accessing the unaligned data (program counter) is located at 3ff80805d60
- The unaligned data is located at address 140000154

```
Unaligned access pid=24821 <a.out> va=140000154,
pc=3ff80805d60,
ra=1200017bc
```

## Ordering Data Declarations to Avoid Unaligned Data

For new programs or when the source declarations of an existing program can be easily modified, plan the order of your data declarations carefully to ensure the data items in a common block, derived-type data, record structure, or data items made equivalent by an `EQUIVALENCE` statement will be naturally aligned.

Use the following rules to prevent unaligned data:

- Always define the largest size numeric data items first.
- If your data includes a mixture of character and numeric data, place the numeric data first.
- Add small data items of the correct size (or padding) before otherwise unaligned data to ensure natural alignment for the data that follows.

When declaring data, consider using explicit length declarations, such as specifying a `KIND` parameter. For example, specify `INTEGER(KIND=4)` (or `INTEGER(4)`) rather than `INTEGER`. If you do use a default size (such as `INTEGER`, `LOGICAL`, `COMPLEX`, and `REAL`), be aware that the compiler options `-i{2|4|8}` and `-r{8|16}` can change the size of an individual field's data declaration size and thus can alter the data alignment of a carefully planned order of data declarations.

Using the suggested data declaration guidelines minimizes the need to use the `-align keyword` options to add padding bytes to ensure naturally aligned data. In cases where the `-align keyword` options are still needed, using the suggested data declaration guidelines can minimize the number of padding bytes added by the compiler.

## Arranging Data Items in Common Blocks

The order of data items in a `common` statement determine the order in which the data items are stored. Consider the following declaration of a common block named `x`:


```

logical (kind=2) flag
integer          iarry_i(3)
character(len=5) name_ch
common /x/ flag, iarry_i(3),
name_ch

```

As shown in Figure 1-1, if you omit the appropriate Fortran command options, the common block will contain unaligned data items beginning at the first array element of `iarry_i`.


**Figure 1-1 Common Block with Unaligned Data**



ZK-6659A-GE

As shown in Figure 1-2, if you compile the program units that use the common block with the `-align commons` option, data items will be naturally aligned.

**Figure 1-2 Common Block with Naturally Aligned Data**



ZK-6660A-GE

Because the common block `x` contains data items whose size is 32 bits or smaller, specify `-align commons` option. If the common block contains data items whose size might be larger than 32 bits (such as `REAL (KIND=8)` data), use `-align commons` option.

If you can easily modify the source files that use the common block data, define the numeric variables in the `COMMON` statement in descending order of size and place the character variable last. This provides more portability, ensures natural alignment without padding, and does not require the Fortran command options `-align commons` or `-align commons` option:


```

logical (kind=2) flag
integer          iarry_i(3)
character(len=5) name_ch
common /x/ iarry_i(3), flag,
name_ch

```

As shown in Figure 1-3, if you arrange the order of variables from largest to smallest size and place character data last, the data items will be naturally aligned.

**Figure 1-3 Common Block with Naturally Aligned Reordered Data**



When modifying or creating all source files that use common block data, consider placing the common block data declarations in a module so the declarations are consistent. If the common block is not needed for compatibility (such as file storage or FORTRAN 77 use), you can place the data declarations in a module without using a common block.

### Arranging Data Items in Derived-Type Data

Like common blocks, derived-type data may contain multiple data items (members).

Data item components within derived-type data will be naturally aligned on up to 64-bit boundaries, with certain exceptions related to the use of the `SEQUENCE` statement and Fortran options. See Options Controlling Alignment for information about these exceptions.

Intel Fortran stores a derived data type as a linear sequence of values, as follows:

- If you specify the `SEQUENCE` statement, the first data item is in the first storage location and the last data item is in the last storage location. The data items appear in the order in which they are declared. The Fortran options have no effect on unaligned data, so data declarations must be carefully specified to naturally align data. The `-align sequence` option specifically aligns data items in a `SEQUENCE` derived-type on natural boundaries.
- If you omit the `SEQUENCE` statement, the Intel Fortran adds the padding bytes needed to naturally align data item components, unless you specify the `-align norecords` option.

Consider the following declaration of array `CATALOG_SPRING` of derived-type `PART_DT`:


```

module data_defs
type part_dt
integer          identifier
real            weight
character(len=15) description
end type part_dt
type(part_dt)
catalog_spring(30)
. . .
end module data_defs

```

As shown in Figure 1-4, the largest numeric data items are defined first and the character data type is defined last. There are no padding characters between data items and all items are naturally aligned. The trailing padding byte is needed because `CATALOG_SPRING` is an array; it is inserted by the compiler when the `-align records` option is in effect.

**Figure 1-4 Derived-Type Naturally Aligned Data (in `CATALOG_SPRING : ( , )`)**



### Arranging Data Items in Intel Fortran Record Structures

Intel Fortran supports record structures provided by Intel Fortran. Intel Fortran record structures use the `RECORD` statement and optionally the `STRUCTURE` statement, which are extensions to the FORTRAN 77 and Fortran standards. The order of data items in a `STRUCTURE` statement determine the order in which the data items are stored.

Intel Fortran stores a record in memory as a linear sequence of values, with the record's first element in the first storage location and its last element in the last storage location. Unless you specify `-align norecords`, padding bytes are added if needed to ensure data fields are naturally aligned.

The following example contains a structure declaration, a `RECORD` statement, and diagrams of the resulting records as they are stored in memory:

```


structure /stra/
character*1 chr
integer*4 int
end structure
. . .
record /stra/ rec

```



Figure 1-5 shows the memory diagram of record `REC` for naturally aligned records.

**Figure 1-5 Memory Diagram of `REC` for Naturally Aligned Records**



## Using Arrays Efficiently

This topic discusses how to efficiently access arrays and how to efficiently pass array arguments.

### Accessing Arrays Efficiently

Many of the array access efficiency techniques described in this section are applied automatically by the Intel Fortran [loop transformations](#) optimizations. Several aspects of array use can improve run-time performance:

- The fastest array access occurs when contiguous access to the whole array or most of an array occurs. Perform one or a few array operations that access all of the array or major parts of an array instead of numerous operations on scattered array elements. Rather than use explicit loops for array access, use elemental array operations, such as the following line that increments all elements of array variable `a`:

```
a = a + 1
```

When reading or writing an array, use the array name and not a `DO` loop or an implied `DO`-loop that specifies each element number. Fortran 95/90 array syntax allows you to reference a whole array by using its name in an expression. For example:

```
real :: a(100,100)
a = 0.0
a = a + 1      ! Increment all
elements      ! of a by 1
. . .
write (8) a    ! Fast whole array
use
```

Similarly, you can use derived-type array structure components, such as:

```

type x
INTEGER A(5)integer a(5)
end type x
. . .
type (x) z
write (8)z%a      ! Fast array
structure         ! component use

```

- Make sure multidimensional arrays are referenced using proper array syntax and are traversed in the **natural ascending storage order**, which is **column-major order** for Fortran. With column-major order, the leftmost subscript varies most rapidly with a stride of one. Whole array access uses column-major order. Avoid **row-major order**, as is done by C, where the rightmost subscript varies most rapidly. For example, consider the nested `do` loops that access a two-dimension array with the `j` loop as the innermost loop:

```

integer x(3,5), y(3,5), i, j
y = 0
do i=1,3                ! I outer loop varies slowest
do j=1,5                ! J inner loop varies fastest
x (i,j) = y(i,j) + 1   ! Inefficient row-major
storage order
end do                  ! (rightmost subscript varies
fastest)
end do
. . .
end program

```

Since `j` varies the fastest and is the second array subscript in the expression `x (i,j)`, the array is accessed in row-major order. To make the array accessed in natural column-major order, examine the array algorithm and data being modified. Using arrays `x` and `y`, the array can be accessed in natural column-major order by changing the nesting order of the `do` loops so the innermost loop variable corresponds to the leftmost array dimension:

```

integer x(3,5), y(3,5), i, j
y = 0
do j=1,5                ! J outer loop varies slowest
do i=1,3                ! I inner loop varies fastest
x (i,j) = y(i,j) + 1   ! Efficient column-major
storage order
end do                  ! (leftmost subscript varies
fastest)
end do
. . .
end program

```

The Intel Fortran whole array access ( $x = y + 1$ ) uses efficient column major order. However, if the application requires that  $J$  vary the fastest or if you cannot modify the loop order without changing the results, consider modifying the application program to use a rearranged order of array dimensions. Program modifications include rearranging the order of:

- Dimensions in the declaration of the arrays  $x(5,3)$  and  $y(5,3)$
- The assignment of  $x(j,i)$  and  $y(j,i)$  within the `do` loops
- All other references to arrays  $x$  and  $y$

In this case, the original `DO` loop nesting is used where  $J$  is the innermost loop:

```
integer x(3,5), y(3,5), i, j
y = 0
do i=1,3                ! I outer loop varies slowest
do j=1,5                ! J inner loop varies fastest
x (j,i) = y(j,i) + 1  ! Efficient column-major storage
order
end do                  ! (leftmost subscript varies
fastest)
end do
. . .
end program
```

Code written to access multidimensional arrays in row-major order (like C) or random order can often make use of the CPU memory cache less efficient. For more information on using natural storage order during record, see [Improving I/O Performance](#).

- Use the available Fortran 95/90 array intrinsic procedures rather than create your own.

Whenever possible, use Fortran 95/90 array intrinsic procedures instead of creating your own routines to accomplish the same task. Fortran 95/90 array intrinsic procedures are designed for efficient use with the various Intel Fortran run-time components.

Using the standard-conforming array intrinsics can also make your program more portable.

- With multidimensional arrays where access to array elements will be noncontiguous, avoid leftmost array dimensions that are a power of two (such as 256, 512).

Since the cache sizes are a power of 2, array dimensions that are also a power of 2 may make less efficient use of cache when array access is noncontiguous. If the cache size is an exact multiple of the leftmost

dimension, your program will probably make inefficient use of the cache. This does not apply to contiguous sequential access or whole array access.

One work-around is to increase the dimension to allow some unused elements, making the leftmost dimension larger than actually needed. For example, increasing the leftmost dimension of A from 512 to 520 would make better use of cache:

```
real a(512, 100)
do i= 2,511
do j = 2,99
a(i,j)=(a(i+1,j-1) + a(i-1, j+1)) * 0.5
end do
end do
```

In this code, array a has a leftmost dimension of 512, a power of two. The innermost loop accesses the rightmost dimension (row major), causing inefficient access. Increasing the leftmost dimension of a to 520 (`real a(520,100)`) allows the loop to provide better performance, but at the expense of some unused elements.

Because loop index variables I and J are used in the calculation, changing the nesting order of the `do` loops changes the results.

For more information on arrays and their data declaration statements, see the *Intel® Fortran Language Reference Manual*.

## Passing Array Arguments Efficiently

In Fortran, there are two general types of array arguments:

- Explicit-shape arrays used with FORTRAN 77.

These arrays have a fixed rank and extent that is known at compile time. Other dummy argument (receiving) arrays that are not deferred-shape (such as assumed-size arrays) can be grouped with explicit-shape array arguments.

- Deferred-shape arrays introduced with Fortran 95/90.

Types of deferred-shape arrays include array pointers and allocatable arrays. Assumed-shape array arguments generally follow the rules about passing deferred-shape array arguments.

When passing arrays as arguments, either the starting (base) address of the array or the address of an array descriptor is passed:

- When using explicit-shape (or assumed-size) arrays to receive an array, the starting address of the array is passed.
- When using deferred-shape or assumed-shape arrays to receive an array, the address of the array descriptor is passed (the compiler creates the array descriptor).

Passing an assumed-shape array or array pointer to an explicit-shape array can slow run-time performance. This is because the compiler needs to create an array temporary for the entire array. The array temporary is created because the passed array may not be contiguous and the receiving (explicit-shape) array requires a contiguous array. When an array temporary is created, the size of the passed array determines whether the impact on slowing run-time performance is slight or severe.

The following table summarizes what happens with the various combinations of array types. The amount of run-time performance inefficiency depends on the size of the array.

### Output Argument Array Types

Input Arguments Array Types	Explicit-Shape Arrays	Deferred-Shape and Assumed-Shape Arrays
Explicit-shape arrays	Very efficient. Does not use an array temporary. Does not pass an array descriptor. Interface block optional.	Efficient. Only allowed for assumed-shape arrays (not deferred-shape arrays). Does not use an array temporary. Passes an array descriptor. Requires an interface block.
Deferred-shape and assumed-shape arrays	When passing an allocatable array, very efficient. Does not use an array temporary. Does not pass an array descriptor. Interface block optional. When not passing an allocatable array, not efficient. Instead use allocatable arrays whenever possible. Uses an array temporary. Does not pass an array descriptor. Interface block optional.	Efficient. Requires an assumed-shape or array pointer as dummy argument. Does not use an array temporary. Passes an array descriptor. Requires an interface block.

## Improving I/O Performance

Improving overall I/O performance can minimize both device I/O and actual CPU time. The techniques listed in this topic can significantly improve performance in many applications.

An I/O flow problems limit the maximum speed of execution by being the slowest process in an executing program. In some programs, I/O is the bottleneck that prevents an improvement in run-time performance. The key to relieving I/O problems is to reduce the actual amount of CPU and I/O device time involved in I/O.

The problems can be caused by one or more of the following:

- A dramatic reduction in CPU time without a corresponding improvement in I/O time
- Such coding practices as:
  - Unnecessary formatting of data and other CPU-intensive processing
  - Unnecessary transfers of intermediate results
  - Inefficient transfers of small amounts of data
  - Application requirements

Improved coding practices can minimize actual device I/O, as well as the actual CPU time.

Intel offers software solutions to system-wide problems like minimizing device I/O delays.

### Use Unformatted Files Instead of Formatted Files

Use unformatted files whenever possible. Unformatted I/O of numeric data is more efficient and more precise than formatted I/O. Native unformatted data does not need to be modified when transferred and will take up less space on an external file.

Conversely, when writing data to formatted files, formatted data must be converted to character strings for output, less data can transfer in a single operation, and formatted data may lose precision if read back into binary form.

To write the array `A( 25 , 25 )` in the following statements, `S1` is more efficient than `S2`:

```
S1          WRITE ( 7 ) A
S2          WRITE ( 7,100 ) A
100        FORMAT ( 25( ' ',25F5.21) )
```

Although formatted data files are more easily ported to other systems, Intel Fortran can convert unformatted data in several formats; see Little-endian-to-Big-endian Conversion.

### **Write Whole Arrays or Strings**

To eliminate unnecessary overhead, write whole arrays or strings at one time rather than individual elements at multiple times. Each item in an I/O list generates its own calling sequence. This processing overhead becomes most significant in implied-DO loops. When accessing whole arrays, use the array name (Fortran array syntax) instead of using implied-DO loops.

### **Write Array Data in the Natural Storage Order**

Use the natural ascending storage order whenever possible. This is column-major order, with the leftmost subscript varying fastest and striding by 1. (See Accessing Arrays Efficiently.) If a program must read or write data in any other order, efficient block moves are inhibited.

If the whole array is not being written, natural storage order is the best order possible.

If you must use an unnatural storage order, in certain cases it might be more efficient to transfer the data to memory and reorder the data before performing the I/O operation.

### **Use Memory for Intermediate Results**

Performance can improve by storing intermediate results in memory rather than storing them in a file on a peripheral device. One situation that may not benefit from using intermediate storage is when there is a disproportionately large amount of data in relation to physical memory on your system. Excessive page faults can dramatically impede virtual memory performance.

If you are primarily concerned with the CPU performance of the system, consider using a memory file system (mfs) virtual disk to hold any files your code reads or writes.

### **Enable Implied-DO Loop Collapsing**

DO loop collapsing reduces a major overhead in I/O processing. Normally, each element in an I/O list generates a separate call to the Intel Fortran run-time

library (RTL). The processing overhead of these calls can be most significant in implied-DO loops.

Intel Fortran reduces the number of calls in implied-DO loops by replacing up to seven nested implied-DO loops with a single call to an optimized run-time library I/O routine. The routine can transmit many I/O elements at once.

Loop collapsing can occur in formatted and unformatted I/O, but only if certain conditions are met:

- The control variable must be an integer. The control variable cannot be a dummy argument or contained in an EQUIVALENCE or VOLATILE statement. Intel Fortran must be able to determine that the control variable does not change unexpectedly at run time.
- The format must not contain a variable format expression.

For information on VOLATILE attribute and statement, see the *Intel® Fortran Language Reference*.

For loop optimizations, see Loop Transformations, Loop Unrolling, and Optimization Levels.

### Use of Variable Format Expressions

Variable format expressions (an Intel Fortran extension) are almost as flexible as run-time formatting, but they are more efficient because the compiler can eliminate run-time parsing of the I/O format. Only a small amount of processing and the actual data transfer are required during run time.

On the other hand, run-time formatting can impair performance significantly. For example, in the following statements, S1 is more efficient than S2 because the formatting is done once at compile time, not at run time:

```
S1          WRITE (6,400) (A(I), I=1,N)
400  FORMAT (1X, <N> F5.2)
          .
          .
          .
S2          WRITE (CHFMT,500)
'(1X, ',N, 'F5.2)'
500  FORMAT (A,I3,A)
      WRITE (6,FMT=CHFMT) (A(I), I=1,N)
```

### Efficient Use of Record Buffers and Disk I/O

Records being read or written are transferred between the user's program buffers and one or more disk block I/O buffers, which are established when the file is



opened by the Intel Fortran RTL. Unless very large records are being read or written, multiple logical records can reside in the disk block I/O buffer when it is written to disk or read from disk, minimizing physical disk I/O.

You can specify the size of the disk block physical I/O buffer by using the `open` statement `BLOCKSIZE` specifier; the default size can be obtained from `fstat(2)`. If you omit the `BLOCKSIZE` specifier in the `open` statement, it is set for optimal I/O use with the type of device the file resides on (with the exception of network access).

The `open` statement `BUFFERCOUNT` specifier specifies the number of I/O buffers. The default for `BUFFERCOUNT` is 1. Any experiments to improve I/O performance should increase the `BUFFERCOUNT` value and not the `BLOCKSIZE` value, to increase the amount of data read by each disk I/O.

If the `open` statement has `BLOCKSIZE` and `BUFFERCOUNT` specifiers, then the internal buffer size in bytes is the product of these specifiers. If the `open` statement does not have these specifiers, then the default internal buffer size is 8192 bytes. This internal buffer will grow to hold the largest single record, but will never shrink.

The default for the Fortran run-time system is to use unbuffered disk writes. That is, by default, records are written to disk immediately as each record is written instead of accumulating in the buffer to be written to disk later.

To enable buffered writes (that is, to allow the disk device to fill the internal buffer before the buffer is written to disk), use one of the following:

- The `OPEN` statement `BUFFERED` specifier
- The `-assume buffered_io` command-line option
- The `FORT_BUFFERED` run-time environment variable

The `open` statement `BUFFERED` specifier takes precedence over the `-assume buffered_io` option. If neither one is set (which is the default), the `FORT_BUFFERED` environment variable is tested at run time.

The `open` statement `BUFFERED` specifier applies to a specific logical unit. In contrast, the `-assume nobuffered_io` option and the `FORT_BUFFERED` environment variable apply to all Fortran units.

Using buffered writes usually makes disk I/O more efficient by writing larger blocks of data to the disk less often. However, a system failure when using buffered writes can cause records to be lost, since they might not yet have been written to disk. (Such records would have been written to disk with the default unbuffered writes.)

When performing I/O across a network, be aware that the size of the block of network data sent across the network can impact application efficiency. When reading network data, follow the same advice for efficient disk reads, by increasing the `BUFFERCOUNT`. When writing data through the network, several items should be considered:

- Unless the application requires that records be written using unbuffered writes, enable buffered writes by a method described above.
- Especially with large files, increasing the `BLOCKSIZE` value increases the size of the block sent on the network and how often network data blocks get sent.
- Time the application when using different `BLOCKSIZE` values under similar conditions to find the optimal network block size.

When writing records, be aware that I/O records are written to unified buffer cache (UBC) system buffers. To request that I/O records be written from program buffers to the UBC system buffers, use the flush library routine (see `FLUSH` in *Intel® Fortran Library Reference*). Be aware that calling flush also discards read-ahead data in user buffer.

### Specify RECL

The sum of the record length (`RECL` specifier in an `open` statement) and its overhead is a multiple or divisor of the blocksize, which is device-specific. For example, if the `BLOCKSIZE` is 8192 then `RECL` might be 24576 (a multiple of 3) or 1024 (a divisor of 8).

The `RECL` value should fill blocks as close to capacity as possible (but not over capacity). Such values allow efficient moves, with each operation moving as much data as possible; the least amount of space in the block is wasted. Avoid using values larger than the block capacity, because they create very inefficient moves for the excess data only slightly filling a block (allocating extra memory for the buffer and writing partial blocks are inefficient).

The `RECL` value unit for formatted files is always 1-byte units. For unformatted files, the `RECL` unit is 4-byte units, unless you specify the `-assume byterecl` option to request 1-byte units (see `-assume byterecl`).

### Use the Optimal Record Type

Unless a certain record type is needed for portability reasons, choose the most efficient type, as follows:

- For sequential files of a consistent record size, the fixed-length record type gives the best performance.

- For sequential unformatted files when records are not fixed in size, the variable-length record type gives the best performance, particularly for `BACKSPACE` operations.
- For sequential formatted files when records are not fixed in size, the `Stream_LF` record type gives the best performance.

## Reading from a Redirected Standard Input File

Due to certain precautions that the Fortran run-time system takes to ensure the integrity of standard input, reads can be very slow when standard input is redirected from a file. For example, when you use a command such as `myprogram.exe < myinput.data>`, the data is read using the `READ(*)` or `READ(5)` statement, and performance is degraded. To avoid this problem, do one of the following:

- Explicitly open the file using the `open` statement. For example:

```
open(5, STATUS='OLD',  
FILE='myinput.dat')
```

Use an environment variable to specify the input file.

To take advantage of these methods, be sure your program does not rely on sharing the standard input file.

For More Information on Intel Fortran data files and I/O, see "Files, Devices, and I/O" in Volume I; on `open` statement specifiers and defaults, see "Open Statement" in the *Intel® Fortran Language Reference Manual*.

## Improving Run-time Efficiency

Source coding guidelines can be implemented to improve run-time performance. The amount of improvement in run-time performance is related to the number of times a statement is executed. For example, improving an arithmetic expression executed within a loop many times has the potential to improve performance, more than improving a similar expression executed once outside a loop.

### Avoid Small Integer and Small Logical Data Items

Avoid using integer or logical data less than 32 bits. Accessing a 16-bit (or 8-bit) data type can make data access less efficient, especially on Itanium-based systems.

To minimize data storage and memory cache misses with arrays, use 32-bit data rather than 64-bit data, unless you require the greater numeric range of 8-byte

integers or the greater range and precision of double precision floating-point numbers.

### Avoid Mixed Data Type Arithmetic Expressions

Avoid mixing integer and floating-point (`REAL`) data in the same computation. Expressing all numbers in a floating-point arithmetic expression (assignment statement) as floating-point values eliminates the need to convert data between fixed and floating-point formats. Expressing all numbers in an integer arithmetic expression as integer values also achieves this. This improves run-time performance.

For example, assuming that `I` and `J` are both `INTEGER` variables, expressing a constant number (2.) as an integer value (2) eliminates the need to convert the data:

Inefficient Code:

```
INTEGER I, J
I = J / 2.
```

Efficient Code:

```
INTEGER I, J
I = J / 2
```

You can use different sizes of the same general data type in an expression with minimal or no effect on run-time performance. For example, using `REAL`, `DOUBLE PRECISION`, and `COMPLEX` floating-point numbers in the same floating-point arithmetic expression has minimal or no effect on run-time performance.

### Use Efficient Data Types

In cases where more than one data type can be used for a variable, consider selecting the data types based on the following hierarchy, listed from most to least efficient:

- Integer (also see above example)
- Single-precision real, expressed explicitly as `REAL`, `REAL (KIND=4)`, or `REAL*4`
- Double-precision real, expressed explicitly as `DOUBLE PRECISION`, `REAL (KIND=8)`, or `REAL*8`
- Extended-precision real, expressed explicitly as `REAL (KIND=16)` or `REAL*16`

However, keep in mind that in an arithmetic expression, you should avoid mixing integer and floating-point (REAL) data (see example in the previous subsection).

### Avoid Using Slow Arithmetic Operators

Before you modify source code to avoid slow arithmetic operators, be aware that optimizations convert many slow arithmetic operators to faster arithmetic operators. For example, the compiler optimizes the expression  $H=J**2$  to be  $H=J*J$ .

Consider also whether replacing a slow arithmetic operator with a faster arithmetic operator will change the accuracy of the results or impact the maintainability (readability) of the source code.

Replacing slow arithmetic operators with faster ones should be reserved for critical code areas. The following hierarchy lists the Intel Fortran arithmetic operators, from fastest to slowest:

- Addition (+), subtraction (-), and floating-point multiplication (\*)
- Integer multiplication (\*)
- Division (/)
- Exponentiation (\*\*)

### Avoid Using EQUIVALENCE Statements

Avoid using EQUIVALENCE statements. EQUIVALENCE statements can:

- Force unaligned data or cause data to span natural boundaries.
- Prevent certain optimizations, including:
  - Global data analysis under certain conditions (see -O2 in Setting Optimization with -On options).
  - Implied-DO loop collapsing when the control variable is contained in an EQUIVALENCE statement

### Use Statement Functions and Internal Subprograms

Whenever the Intel Fortran compiler has access to the use and definition of a subprogram during compilation, it may choose to inline the subprogram. Using statement functions and internal subprograms maximizes the number of subprogram references that will be inlined, especially when multiple source files are compiled together at optimization level -O3.

For more information, see Efficient Compilation.

## Code DO Loops for Efficiency

Minimize the arithmetic operations and other operations in a DO loop whenever possible. Moving unnecessary operations outside the loop will improve performance (for example, when the intermediate nonvarying values within the loop are not needed).

For More Information on loop optimizations, see *Pipelining for Itanium®-based Applications and Loop Unrolling*; on coding Intel Fortran statements, see the *Intel® Fortran Language Reference Manual*.

## Using Ininsics for Itanium®-based Systems

Intel® Fortran supports all standard Fortran intrinsic procedures and in addition, provides Intel-specific intrinsic procedures to extend the functionality of the language. Intel Fortran intrinsic procedures are provided in the library `libintrins.a`. See the *Intel® Fortran Language Reference*.

This topic provides examples of the Intel-extended intrinsics that are helpful in developing efficient applications.

### Cache Size Intrinsic (Itanium® Compiler)

Intrinsic `cachesize(n)` is used only with Intel® Itanium® Compiler. `cachesize(n)` returns the size in kilobytes of the cache at level `n`; 1 represents the first level cache. Zero is returned for a nonexistent cache level.

This intrinsic can be used in many scenarios where application programmer would like to tailor their algorithms for target processor's cache hierarchy. For example, an application may query the cache size and use it to select block sizes in algorithms that operate on matrices.

```
subroutine foo (level)
integer level
if (cachesize(level) >
threshold)
call big_bar()
else
call small_bar()
end if
end subroutine
```

## Coding Guidelines for Intel® Architectures

This section provides general guidelines for coding practices and techniques that insure most benefits of using:

- IA-32 architecture supporting MMX(TM) technology and Streaming SIMD Extensions (SSE) and Streaming SIMD Extensions 2 (SSE2)
- Itanium® architecture

This section describes practices, tools, coding rules and recommendations associated with the architecture features that can improve the performance on IA-32 and Itanium processors families. For all details about optimization for IA-32 processors, see Intel® Architecture Optimization Reference Manual. For all details about optimization for Itanium processor family, see the Intel Itanium 2 Processor Reference Manual for Software Development and Optimization.

### Note

If a guideline refers to a particular architecture only, this architecture is explicitly named. The default is for both IA-32 and Itanium architectures.

Performance of compiler-generated code may vary from one compiler to another. Intel® Fortran Compiler generates code that is highly optimized for Intel architectures. You can significantly improve performance by using various compiler optimization options. In addition, you can help the compiler to optimize your Fortran program by following the guidelines described in this section.

When coding in Fortran, the most important factors to consider in achieving optimum processor performance are:

- avoiding memory access stalls
- ensuring good floating-point performance
- ensuring good SIMD integer performance
- using vectorization.

The following sections summarize and describe coding practices, rules and recommendations associated with the features that will contribute to optimizing the performance on Intel architecture-based processors.

## Memory Access

The Intel compiler lays out Fortran arrays in column-major order. For example, in a two-dimensional array, elements  $A(22, 34)$  and  $A(23, 34)$  are contiguous in memory. For best performance, code arrays so that inner loops access them in a contiguous manner. Consider the following examples.

The code in example 1 will likely have higher performance than the code in example 2.

### Example 1

```
DO J = 1, N
DO I = 1, N
B(I,J) = A(I, J) + 1
END DO
END DO
```

The code above illustrates access to arrays A and B in the inner loop I in a contiguous manner which results in good performance.

### Example 2

```
DO I = 1, N
DO J = 1, N
B(I,J) = A(I, J) + 1
END DO
END DO
```

The code above illustrates access to arrays A and B in inner loop J in a non-contiguous manner which results in poor performance.

The compiler itself can transform the code so that inner loops access memory in a contiguous manner. To do that, you need to use advanced optimization options, such as `-O3` for both IA-32 and Itanium architectures, and `-O3` and `-axW|N|B|P` for IA-32 only.

### Memory Layout

Alignment is a very important factor in ensuring good performance. Aligned memory accesses are faster than unaligned accesses. If you use the interprocedural optimization on multiple files (the `-ipo` option), the compiler analyzes the code and decides whether it is beneficial to pad arrays so that they start from an aligned boundary. Multiple arrays specified in a single common block can impose extra constraints on the compiler. For example, consider the following `COMMON` statement:

```
COMMON /AREA1/ A(200), X, B(200)
```

If the compiler added padding to align `A(1)` at a 16-byte aligned address, the element `B(1)` would not be at a 16-byte aligned address. So it is better to split `AREA1` as follows.

```
COMMON /AREA1/
A(200)
COMMON /AREA2/ X
COMMON /AREA3/
B(200)
```



The above code provides the compiler maximum flexibility in determining the padding required for both A and B.

## Optimizing for Floating-point Applications

To improve floating-point performance, generally follow these rules:

- Avoid exceeding representable ranges during computation since handling these cases can have a performance impact. Use `REAL` variables in single-precision format unless the extra precision obtained through `DOUBLE` or `REAL*10` variables is required. Using variables with a larger precision formation will also increase memory size and bandwidth requirements.
- **For IA-32 only:** Avoid repeatedly changing rounding modes between more than two values, which can lead to poor performance when the computation is done using non-SSE instructions. Hence avoid using `FLOOR` and `TRUNC` instructions together when generating non-SSE code. The same applies for using `CEIL` and `TRUNC`.

Another way to avoid the problem is to use the `-x{K|W|N|B|P}` options to do the computation using SSE instructions.

- Reduce the impact of denormal exceptions for both architectures as described below.

### Denormal Exceptions

Floating point computations with underflow can result in denormal values that have an adverse impact on performance.

**For IA-32:** take advantage of the SIMD capabilities of Streaming SIMD Extensions (SSE), and Streaming SIMD Extensions 2 (SSE2) instructions. The `-x{K|W|N|B|P}` options enable the flush-to-zero (FTZ) mode in SSE and SSE2 instructions, whereby underflow results are automatically converted to zero, which improves application performance. In addition, the `-xP` option also enables the denormals-as-zero (DAZ) mode, whereby denormals are converted to zero on input, further improving performance. An application developer willing to trade pure IEEE-754 compliance for speed would benefit from these options. For more information on FTZ and DAZ, see Setting FTZ and DAZ Flags and "Floating-point Exceptions" in the Intel® Architecture Optimization Reference Manual.

**For Itanium architecture:** enable flush-to-zero (FTZ) mode with the `-ftz` option set by `-O3` option.

## Auto-vectorization

Many applications significantly increase their performance if they can implement vectorization, which uses streaming SIMD SSE2 instructions for the main computational loops. The Intel Compiler turns vectorization on (auto-vectorization) or you can implement it with compiler directives.

See Auto-vectorization (IA-32 Only) section for complete details.

## Creating Multithreaded Applications

The Intel Fortran Compiler and the Intel® Threading Toolset have the capabilities that make developing multithreaded application easy. See Parallel Programming with Intel Fortran. Multithreaded applications can show significant benefit on multiprocessor Intel symmetric multiprocessing (SMP) systems or on Intel processors with Hyper-Threading technology.

# Analyzing and Timing Your Application

## Using Intel Performance Analysis Tools

Intel offers an array of application performance tools that are optimized to take the best advantage of the Intel architecture-based processors. You can employ these tools for developing the most efficient programs without having to write assembly code.

The performance tools to help you analyze your application and find and resolve the problem areas are as follows:

- Intel® Enhanced Debugger for IA-32 systems and Intel® Debugger (IDB) for Itanium®-based systems.

The Enhanced Debugger (EDB) enables you to debug your programs and view the XMM registers in a variety of formats corresponding to the data types supported by Streaming SIMD Extensions and Streaming SIMD Extensions 2.

The IDB debugger provides extensive support for debugging programs by using a command-line or graphical user interface.

- Intel® VTune(TM) Performance Analyzer

The VTune analyzer collects, analyzes, and provides Intel architecture-specific software performance data from the system-wide view down to a specific module, function, and instruction in your code. For information, see <http://www.intel.com/software/products/vtune/>.

- Intel® Threading Tools. The Intel Threading Tools consist of the following:
  - Intel® Thread Checker
  - Intel® Thread Profiler

For general information, see <http://www.intel.com/software/products/threadtool.htm>.

## Timing Your Application

One of the performance indicators is your application timing. Use the `time` command to provide information about program performance. The following considerations apply to timing your application:

- Run program timings when other users are not active. Your timing results can be affected by one or more CPU-intensive processes also running while doing your timings.
- Try to run the program under the same conditions each time to provide the most accurate results, especially when comparing execution times of a previous version of the same program. Use the same CPU system (model, amount of memory, version of the operating system, and so on) if possible.
- If you do need to change systems, you should measure the time using the same version of the program on both systems, so you know each system's effect on your timings.
- For programs that run for less than a few seconds, run several timings to ensure that the results are not misleading. Overhead functions like loading shared libraries might influence short timings considerably.

Using the form of the `time` command that specifies the name of the executable program provides the following:

- The elapsed, real, or "wall clock" time, which will be greater than the total charged actual CPU time.
- Charged actual CPU time, shown for both system and user execution. The total actual CPU time is the sum of the actual user CPU time and actual system CPU time.

### Example

In the following example timings, the sample program being timed displays the following line:

```
Average of all the numbers is:      4368488960.000000
```

Using the Bourne shell, the following program timing reports that the program uses 1.19 seconds of total actual CPU time (0.61 seconds in actual CPU time for

user program use and 0.58 seconds of actual CPU time for system use) and 2.46 seconds of elapsed time:

```
$ time a.out

Average of all the numbers is:
 4368488960.000000

real    0m2.46s
user    0m0.61s
sys     0m0.58s
```

Using the C shell, the following program timing reports 1.19 seconds of total actual CPU time (0.61 seconds in actual CPU time for user program use and 0.58 seconds of actual CPU time for system use), about 4 seconds (0:04) of elapsed time, the use of 28% of available CPU time, and other information:

```
% time a.out

Average of all the numbers is:
 4368488960.000000

0.61u 0.58s 0:04 28% 78+424k 9+5io 0pf+0w
```

Using the bash shell, the following program timing reports that the program uses 1.19 seconds of total actual CPU time (0.61 seconds in actual CPU time for user program use and 0.58 seconds of actual CPU time for system use) and 2.46 seconds of elapsed time:

```
[user@system user]$ time ./a.out

Average of all the numbers is:
 4368488960.000000

elapsed 0m2.46s
user    0m0.61s
sys     0m0.58s
```


Timings that show a large amount of system time may indicate a lot of time spent doing I/O, which might be worth investigating.

If your program displays a lot of text, you can redirect the output from the program on the time command line. Redirecting output from the program will change the times reported because of reduced screen I/O.

For more information, see `time(1)`.

In addition to the `time` command, you might consider modifying the program to call routines within the program to measure execution time. For example, use the

Intel Fortran intrinsic procedures, such as SECNDS, DCLOCK, CPU\_TIME, SYSTEM\_CLOCK, TIME, and DATE\_AND\_TIME. See "Intrinsic Procedures" in the *Intel® Fortran Language Reference*.



# Compiler Optimizations

## Compiler Optimizations Overview

The variety of optimizations used by the Intel® Fortran Compiler enable you to enhance the performance of your application. Each optimization is performed by a set of options discussed in each section dedicated to the following optimizations:

- Optimizing compilation process
- Optimizing different types of applications
- Floating-point arithmetic operations
- Optimizing applications for specific processors
- Interprocedural optimizations (IPO)
- Profile-guided optimizations
- High-level Language optimizations

In addition to optimizations invoked by the compiler command-line options, the compiler includes features which enhance your application performance such as directives, intrinsics, run-time library routines and various utilities. These features are discussed in the Optimization Support Features section.

## Optimizing Compilation Process

### Optimizing Compilation Process Overview

This section describes the Intel® Fortran Compiler options that optimize the compilation process. By default, the compiler converts source code directly to an executable file. Appropriate options enable you not only to control the process and obtain desired output file produced by the compiler, but also make the compilation itself more efficient.

A group of options monitors the outcome of Intel compiler-generated code without interfering with the way your program runs. These options control some computation aspects, such as allocating the stack memory, setting or modifying variable settings, and defining the use of some registers.

The options in this section provide you with the following capabilities of efficient compilation:

- automatic allocation of variables and stacks
- aligning data
- symbol visibility attribute options

## Efficient Compilation

Understandably, efficient compilation contributes to performance improvement. Before you analyze your program for performance improvement, and improve program performance, you should think of efficient compilation itself. Based on the analysis of your application, you can decide which Intel Fortran Compiler optimizations and command-line options can improve the run-time performance of your application.

### Efficient Compilation Techniques

The efficient compilation techniques can be used during the earlier stages and later stages of program development.

During the earlier stages of program development, you can use incremental compilation with minimal optimization. For example:

```
ifort -c -O1 sub2.f90 (generates object file of sub2)
```

```
ifort -c -O1 sub3.f90 (generates object file of sub3)
```

```
ifort -omain.exe -g -O0 main.f90 sub2.obj sub3.obj
```

The above command turns off all compiler default optimizations (for example, -O2) with -O0. You can use the -g option to generate symbolic debugging information and line numbers in the object code for use by a source-level debugger. The main.exe file created in the above command contains symbolic debugging information.

During the later stages of program development, you should specify multiple source files together and use an optimization level of at least -O2 (default) to allow more optimizations to occur. For instance, the following command compiles all three source files together using the default level of optimization, -O2:

```
ifort -omain.exe main.f90 sub2.f90 sub3.f90
```

Compiling multiple source files lets the compiler examine more code for possible optimizations, which results in:

- Inlining more procedures
- More complete data flow analysis
- Reducing the number of external references to be resolved during linking

For very large programs, compiling all source files together may not be practical. In such instances, consider compiling source files containing related routines

together using multiple `ifort` commands, rather than compiling source files individually.

## Options That Improve Run-Time Performance

The table below lists the options that can directly improve run-time performance. Most of these options do not affect the accuracy of the results, while others improve run-time performance but can change some numeric results. The Intel Fortran Compiler performs [some optimizations by default](#) unless you turn them off by corresponding command-line options. Additional optimizations can be enabled or disabled using command options.

Option	Description
<code>-align</code> <i>keyword</i>	Analyzes and reorders memory layout for variables and arrays.  Controls whether padding bytes are added between data items within common blocks, derived-type data, and record structures to make the data items naturally aligned.
<code>-ax{K W N B P}</code> <b>IA-32 only</b>	Optimizes your application's performance for specific processors. Regardless of which <code>-ax</code> suboption you choose, your application is optimized to use all the benefits of that processor with the resulting binary file capable of being run on any Intel IA-32 processor.
<code>-O2 (-fast)</code>	Sets the following performance-related options: <code>-align</code> <i>dcommons</i> , <code>-align</code> <i>sequence</i>
<code>-O1</code> , <code>-inline</code> <i>all</i>	Inlines every call that can possibly be inlined while generating correct code. Certain recursive routines are not inlined to prevent infinite loops
<code>-parallel</code>	Enables parallel processing using directed decomposition (directives inserted in source code). This can improve the performance of certain programs running on shared memory multiprocessor systems
<code>-On</code>	Controls the types of optimization performed. The default optimizations set is <code>-O2</code> , unless you specify <code>-O0</code> (no optimizations). Use <code>-O3</code> to activate loop transformation optimizations.
<code>-openmp</code>	Enables parallel processing using directed decomposition (directives inserted in source code). This can improve the performance of certain programs running on shared memory multiprocessor systems.




<code>-qp</code>	Requests profiling information, which you can use to identify those parts of your program where improving source code efficiency would most likely improve run-time performance. After you modify the appropriate source code, recompile the program and test the run-time performance.
<code>-tpp{n}</code>	Specifies the target processor generation architecture on which the program will be run, allowing the optimizer to make decisions about instruction tuning optimizations needed to create the most efficient code. Keywords allow specifying one particular processor generation type, multiple processor generation types, or the processor generation type currently in use during compilation. Regardless of the setting of <code>-tpp{n}</code> , the generated code will run correctly on all implementations of the Intel® IA-32 or Itanium® architectures.
<code>-unrolln</code>	Specifies the number of times a loop is unrolled ( <i>n</i> ) when specified with optimization level <code>-O3</code> . If you omit <i>n</i> in <code>-unroll</code> , the optimizer determines how many times loops can be unrolled.

### Options That Slow Down the Run-time Performance

The table below lists options that can slow down the run-time performance. Some applications that require floating-point exception handling or rounding might need to use the `-fpen` dynamic option. Other applications might need to use the `-assume dummy_aliases` or `-vms` options for compatibility reasons. Other options that can slow down the run-time performance are primarily for troubleshooting or debugging purposes.

Table below lists the options that can slow down run-time performance.

Option	Description
<code>-assume dummy_aliases</code>	<p>Forces the compiler to assume that dummy (formal) arguments to procedures share memory locations with other dummy arguments or with variables shared through use association, host association, or common block use. These program semantics slow performance, so you should specify <code>-assume dummy_aliases</code> only for the called subprograms that depend on such aliases.</p> <p>The use of dummy aliases violates the FORTRAN 77 and Fortran 95/90 standards but occurs in some older programs.</p>

-c	If you use <code>-c</code> when compiling multiple source files, also specify <code>-ooutputfile</code> to compile many source files together into one object file. Separate compilations prevent certain interprocedural optimizations, such as when using multiple <code>f90</code> commands or using <code>-c</code> without the <code>-ooutputfile</code> option.
<code>-check bounds</code>	Generates extra code for array bounds checking at run time.
<code>-check overflow</code>	Generates extra code to check integer calculations for arithmetic overflow at run time. Once the program is debugged, omit this option to reduce executable program size and slightly improve run-time performance.
<code>-fpe 3</code>	Using this option slows program execution. It enables certain types of floating-point exception handling, which can be expensive.
<code>-g</code>	<p>Generate extra symbol table information in the object file. Specifying this option also reduces the default level of optimization to <code>-O0</code> or <code>-O0</code> (no optimization).</p> <p> <b>Note</b></p> <p>The <code>-g</code> option only slows your program down when no optimization level is specified, in which case <code>-g</code> turns on <code>-O0</code>, which slows the compilation down. If <code>-g</code>, <code>-O2</code> are specified, the code runs very much the same speed as if <code>-g</code> were not specified.</p>
<code>-inline none</code> <code>-inline manual</code>	Prevents the inlining of all procedures except statement functions.
<code>-save</code>	Forces the local variables to retain their values from the last invocation terminated. This may change the output of your program for floating-point values as it forces operations to be carried out in memory rather than in registers, which in turn causes more frequent rounding of your results.
<code>-O0</code>	Turns off optimizations. Can be used during the early stages of program development or when you use the debugger.
<code>-vms</code>	Controls certain VMS-related run-time defaults, including alignment. If you specify the <code>-vms</code> option, you may need to also specify the <code>-align records</code> option to obtain optimal run-time performance.

## Little-endian-to-Big-endian Conversion (IA-32)

The Intel Fortran Compiler writes unformatted sequential files in big-endian format and reads files produced in big-endian format.

The little-endian-to-big-endian conversion feature is intended for Fortran unformatted input/output operations in unformatted sequential files. It enables the development and processing of files with big-endian data organization on the IA-32-based processors, which usually process the data in the little endian format.

The feature also enables processing of the files developed on processors that accept big-endian data format and producing the files for such processors on IA-32-based little-endian systems.

The little-endian-to-big-endian conversion is accomplished by the following operations:

- The `WRITE` operation converts little endian format to big endian format.
- The `READ` operation converts big endian format to little endian format.

The feature enables the conversion of variables and arrays (or array subscripts) of basic data types. Derived data types are not supported.

### Little-to-Big Endian Conversion Environment Variable

In order to use the little-endian-to-big-endian conversion feature, specify the numbers of the units to be used for conversion purposes by setting the `F_UFMTENDIAN` environment variable. Then, the `READ/WRITE` statements that use these unit numbers, will perform relevant conversions. Other `READ/WRITE` statements will work in the usual way.

In the general case, the variable consists of two parts divided by a semicolon. No spaces are allowed inside the `F_UFMTENDIAN` value. The variable has the following syntax:

```
F_UFMTENDIAN=MODE | [MODE;] EXCEPTION
```

where:

```
MODE = big | little
EXCEPTION = big:ULIST | little:ULIST | ULIST
ULIST = U | ULIST,U
U = decimal | decimal -decimal
```

- `MODE` defines current format of data, represented in the files; it can be omitted.  
The keyword `little` means that the data have little endian format and will

not be converted. For IA-32 systems, this keyword is a default.

The keyword `big` means that the data have big endian format and will be converted. This keyword may be omitted together with the colon.

- `EXCEPTION` is intended to define the list of exclusions for `MODE`; it can be omitted. `EXCEPTION` keyword (`little` or `big`) defines data format in the files that are connected to the units from the `EXCEPTION` list. This value overrides `MODE` value for the units listed.
- Each list member `U` is a simple unit number or a number of units. The number of list members is limited to 64.  
`decimal` is a non-negative decimal number less than  $2^{32}$ .

Converted data should have basic data types, or arrays of basic data types. Derived data types are disabled.

**Command lines** for variable setting with different shells:

```
Sh: export F_UFMTENDIAN=MODE;EXCEPTION
```

```
Csh: setenv F_UFMTENDIAN MODE;EXCEPTION
```

### Note

Environment variable value should be enclosed in quotes if semicolon is present.

### Another Possible Environment Variable Setting

The environment variable can also have the following syntax:

```
F_UFMTENDIAN=u[,u] . . .
```

**Command lines** for the variable setting with different shells:

- Sh: `export F_UFMTENDIAN=u[,u] . . .`
- Csh: `setenv F_UFMTENDIAN u[,u] . . .`

See error messages that may be issued during the little endian – big endian conversion. They are all fatal. You should contact Intel if such errors occur.

### Usage Examples

1. `F_UFMTENDIAN=big`

All input/output operations perform conversion from big-endian to little-endian on `READ` and from little-endian to big-endian on `WRITE`.

2. `F_UFMTENDIAN="little;big:10,20"`  
or `F_UFMTENDIAN=big:10,20`  
or `F_UFMTENDIAN=10,20`

In this case, only on unit numbers 10 and 20 the input/output operations perform big-little endian conversion.

3. `F_UFMTENDIAN="big;little:8"`

In this case, on unit number 8 no conversion operation occurs. On all other units, the input/output operations perform big-little endian conversion.

4. `F_UFMTENDIAN=10-20`

Define 10, 11, 12 ... 19, 20 units for conversion purposes; on these units, the input/output operations perform big-little endian conversion.

5. Assume you set `F_UFMTENDIAN=10,100` and run the following program.

```
integer*4   cc4
integer*8   cc8
integer*4   c4
integer*8   c8
c4 = 456
c8 = 789

C prepare a little endian representation
of data

open(11,file='lit.tmp',form='unformatted')
write(11) c8
write(11) c4
close(11)

C prepare a big endian representation of
data

open(10,file='big.tmp',form='unformatted')
write(10) c8
write(10) c4
close(10)

C read big endian data and operate with
them on
C little endian machine.

open(100,file='big.tmp',form='unformatted')
read(100) cc8
read(100) cc4
```

```

C Any operation with data, which have been
read

C
close(100)
stop
end

```

Now compare `lit.tmp` and `big.tmp` files with the help of `od` utility.

```
> od -t x4 lit.tmp
```

```

0000000 00000008 00000315 00000000 00000008
0000020 00000004 000001c8 00000004
0000034

```

```
> od -t x4 big.tmp
```

```

0000000 08000000 00000000 15030000 08000000
0000020 04000000 c8010000 04000000
0000034

```

You can see that the byte order is different in these files.

## Default Compiler Optimizations

If you invoke the Intel® Fortran Compiler without specifying any compiler options, the default state of each option takes effect. The following tables summarize the options whose default status is ON as they are required for Intel Fortran Compiler default operation. The tables group the options by their functionality.

For the default states and values of all options, see Compiler Options Quick Reference Alphabetical table in the *Intel® Fortran Compiler Options Quick Reference*. The table provides links to the sections describing the functionality of the options. If an option has a default value, such value is indicated.

Per your application requirement, you can disable one or more options. For general methods of disabling optimizations, see Volume I.

The following tables list all options that compiler uses for its default optimizations.

### Data Setting and Fortran Language Conformance

Default Option	Description
<code>-align</code> <code>-align <i>records</i></code>	Analyzes and reorders memory layout for variables and arrays.

<code>-align rec8bytes</code>	Specifies 8-byte boundary for alignment constraint.
<code>-altparam</code>	Specifies if alterate form of parameter constant declarations is recognized or not.
<code>-ansi_alias</code>	Enables assumption of the program's ANSI conformance.
<code>-assume cc_omp</code>	Enables OpenMP conditional compilation directives.
<code>-ccdefault default</code>	Specifies default carriage control for units 6 and *.
<code>-double_size 64</code>	Defines the default KIND for double-precision variables to be 64.  <code>-double_size 64 n</code> is 64 (KIND=8)
<code>-dps</code>	Enables DEC* parameter statement recognition.
<code>-error_limit 30</code>	Specifies the maximum number of error-level or fatal-level compiler errors permissible.
<code>-fpe 3</code>	Specifies floating-point exception handling at run time for the main program.
<code>-iface nomixed_str_len_arg</code>	Specifies the type of argument-passing conventions used for general arguments and for hidden-length character arguments.
<code>-integer_size 32</code>	Specifies the default size of integer and logical variables.
<code>-libdir all</code>	Controls the library names that should be emitted into the object file.
<code>-pad</code>	Enables changing variable and array memory layout.
<code>-pc80</code> IA-32 only	<code>-pc{32 64 80}</code> enables floating-point significand precision control as follows: <code>-pc32</code> to 24-bit significand, <code>-pc64</code> to 53-bit significand, and <code>-pc80</code> to 64-bit significand.
<code>-real_size 64</code>	Specifies the size of REAL and COMPLEX declarations, constants, functions, and intrinsics.
<code>-save</code>	Saves all variables in static allocation. Disables <code>-auto</code> , that is, disables setting all

	variables <code>AUTOMATIC</code> .
<code>-Zp8</code>	<code>-Zp{n}</code> specifies alignment constraint for structures on 1-, 2-, 4-, 8-, or 16-byte boundary. To disable, use <code>-align-</code> .

## Optimizations

Default Option	Description
<code>-assume cc_omp</code>	Enables OpenMP conditional compilation directives.
<code>-fp</code> IA-32 only	Disables the use of the <code>ebp</code> register in optimizations. Directs to use the <code>ebp</code> -based stack frame for all functions.
<code>-fpe 3</code>	Specifies floating-point exception handling at run time for the main program. <code>-fpe 0</code> disables the option.
<code>-ip_no_inlining</code>	Disables full or partial inlining that would result from the <code>-ip</code> interprocedural optimizations. Requires <code>-ip</code> or <code>-ipo</code> .
<code>-IPFfltacc-</code> Itanium® compiler	Enables the compiler to apply optimizations that affect floating-point accuracy.
<code>-IPFfma</code> Itanium compiler	Enables the contraction of floating-point multiply and add/subtract operations into a single operation.
<code>-IPFfp_speculation</code> <i>fast</i> Itanium compiler	Sets the compiler to speculate on floating-point operations. <code>-IPFfp_speculationoff</code> disables this optimization.
<code>-ipo_obj</code> Itanium compiler	Forces the generation of real object files. Requires <code>-ipo</code> . IA-32 systems: OFF
<code>-O, -O1, -O2</code>	Optimize for maximum speed.
<code>-Ob1</code>	Disables inlining unless <code>-ip</code> or <code>-Ob2</code> is specified.
<code>-openmp_report1</code>	Indicates loops, regions, and sections parallelized.



<code>-opt_report_levelmin</code>	Specifies the minimal level of the optimizations report.
<code>-par_report1</code>	Indicates loops successfully auto-parallelized.
<code>-tpp2</code> Itanium compiler	Optimizes code for the Intel® Itanium® 2 processor for Itanium-based applications. Generated code is compatible with the Itanium processor.
<code>-tpp7</code> IA-32 only	Optimizes code for the Intel® Pentium® 4 and Intel® Xeon(TM) processor for IA-32 applications.
<code>-unroll</code>	<code>-unroll[n]</code> : omit <i>n</i> to let the compiler decide whether to perform unrolling or not (default). Specify <i>n</i> to set maximum number of times to unroll a loop. The Itanium compiler currently uses only <i>n</i> = 0, <code>-unroll0</code> (disabled option) for compatibility.
<code>-vec_report1</code>	Indicates loops successfully vectorized.

## Disabling Default Options

To disable an option, use one of the following as applies:

- Generally, to disable one or a group of optimization options, use `-O0 option`. For example:

```
ifort -O2 -O0 input_file(s)
```

### Note

The `-O0` option is part of a mutually-exclusive group of options that includes `-O0`, `-O`, `-O1`, `-O2`, and `-O3`. The last of any of these options specified on the command line will override the previous options from this group.

- To disable options that include optional "-" shown as `[-]`, use that version of the option in the command line, for example: `-align-`.
- To disable options that have `{n}` parameter, use `n=0` version, for example: `-unroll0`.

 **Note**

If there are enabling and disabling versions of switches on the line, the last one takes precedence.

## Using Compilation Options

### Stacks: Automatic Allocation and Checking

The options in this group enable you to control the computation of stacks and variables in the compiler generated code.

#### Automatic Allocation of Variables

##### -auto

The `-auto` option specifies that locally declared variables are allocated to the run-time stack rather than static storage. If variables defined in a procedure appear in an `AUTOMATIC` statement, or if the procedure is recursive and the variables do not have the `SAVE` or `ALLOCATABLE` attribute, they are allocated to the stack. It does not affect variables that appear in an `EQUIVALENCE` or `SAVE` statement, or those that are in `COMMON`.

`-auto` is the same as `-automatic` and `-nosave`.

`-auto` may provide a performance gain for your program, but if your program depends on variables having the same value as the last time the routine was invoked, your program may not function properly. Variables that need to retain their values across routine calls should appear in a `SAVE` statement.

If you specify `-recursive` or `-openmp`, the default is `-auto`.

##### -auto\_scalar

The `-auto_scalar` option causes allocation of local scalar variables of intrinsic type `INTEGER`, `REAL`, `COMPLEX`, or `LOGICAL` to the stack. This option does not affect variables that appear in an `EQUIVALENCE` or `SAVE` statement, or those that are in `COMMON`.

`-auto_scalar` may provide a performance gain for your program, but if your program depends on variables having the same value as the last time the routine was invoked, your program may not function properly. Variables that need to retain their values across subroutine calls should appear in a `SAVE` statement. This option is similar to `-auto`, which causes all local variables to be allocated on the stack. The difference is that `-auto_scalar` allocates only scalar variables of the stated above intrinsic types to the stack.

`-auto_scalar` enables the compiler to make better choices about which variables should be kept in registers during program execution.

### **-save, -zero**

The `-save` option is opposite of `-auto`: the `-save` option saves all variables in static allocation except local variables within a recursive routine. If a routine is invoked more than once, this option forces the local variables to retain their values from the last invocation terminated. This may cause a performance degradation and may change the output of your program for floating-point values as it forces operations to be carried out in memory rather than in registers, which in turn causes more frequent rounding of your results. `-save` is the same as `-noauto`.

The `-[no]zero` option initializes to zero all local scalar variables of intrinsic type `INTEGER`, `REAL`, `COMPLEX`, or `LOGICAL`, which are saved and not initialized yet. Used in conjunction with `-save`. The default is `-nozero`.

### **Summary**

There are three choices for allocating variables: `-save`, `-auto`, and `-auto_scalar`. Only one of these three can be specified. The correlation among them is as follows:

- `-save` disables `-auto`, sets `-noautomatic`, and allocates all variables not marked `AUTOMATIC` to static memory.
- `-auto` disables `-save`, sets `-automatic`, and allocates all variables—scalars and arrays of all types—not marked `SAVE` to the stack.
- `-auto_scalar`:
  - It makes local scalars of intrinsic types `INTEGER`, `REAL`, `COMPLEX`, and `LOGICAL` automatic.
  - This is the default; there is no `-noauto_scalar`; however, `-recursive` or `-openmp` disables `-auto_scalar` and makes `-auto` the default.

### **Checking the Floating-point Stack State (IA-32 only), `-fpstkchk`**

The `-fpstkchk` option (IA-32 only) checks whether a program makes a correct call to a function that should return a floating-point value. If an incorrect call is detected, the option places a code that marks the incorrect call in the program.

When an application calls a function that returns a floating-point value, the returned floating-point value is supposed to be on the top of the floating-point stack. If return value is not used, the compiler must pop the value off of the floating-point stack in order to keep the floating-point stack in correct state.

If the application calls a function, either without defining or incorrectly defining the function's prototype, the compiler does not know whether the function must return a floating-point value, and the return value is not popped off of the floating-point stack if it is not used. This can cause the floating-point stack overflow.

The overflow of the stack results in two undesirable situations:

- a NAN value gets involved in the floating-point calculations
- the program results become unpredictable; the point where the program starts making errors can be arbitrarily far away from the point of the actual error.

The `-fpstkchk` option marks the incorrect call and makes it easy to find the error.

#### Note

This option causes significant code generation after every function/subroutine call to insure a proper state of a floating-point stack and slows down compilation. It is meant only as a debugging aid for finding floating point stack underflow/overflow problems, which can be otherwise hard to find.

## Aliases

`-common_args`

The `-common_args` option assumes that the "by-reference" subprogram arguments may have aliases of one another.

### Preventing CRAY\* Pointer Aliasing

Option `-safe_cray_ptr` specifies that the CRAY\* pointers do not alias with other variables. The default is OFF.

Consider the following example.

```
pointer (pb, b)
pb =
getstorage()
do i = 1, n
b(i) = a(i) + 1
enddo
```

When `-safe_cray_ptr` is not specified (default), the compiler assumes that `b` and `a` are aliased. To prevent such an assumption, specify this option, and the compiler will treat `b(i)` and `a(i)` as independent of each other.

However, if the variables are intended to be aliased with CRAY pointers, using the `-safe_cray_ptr` option produces incorrect result. For the code example below, `-safe_cray_ptr` should not be used.

```
pb = loc(a(2))
do i=1, n
  b(i) = a(i) +1
enddo
```

### **Cross-platform, -ansi alias**

The `-ansi_alias[-]` enables (default) or disables the compiler to assume that the program adheres to the ANSI Fortran type aliasability rules. For example, an object of type `real` cannot be accessed as an integer. You should see the ANSI standard for the complete set of rules.

The option directs the compiler to assume the following:

- Arrays are not accessed out of arrays' bounds.
- Pointers are not cast to non-pointer types and vice-versa.
- References to objects of two different scalar types cannot alias. For example, an object of type `integer` cannot alias with an object of type `real` or an object of type `real` cannot alias with an object of type `double precision`.

If your program satisfies the above conditions, setting the `-ansi_alias` option will help the compiler better optimize the program. However, if your program may not satisfy one of the above conditions, the option must be disabled, as it can lead the compiler to generate incorrect code.

The synonym of `-ansi_alias` is `-assume [no]dummy_aliases`.

### **Alignment Options**

#### **`-align recnbyte` or `-Zp[n]`**

Use the `-align recnbyte` (or `-Zp[n]`) option to specify the alignment constraint for structures on `n`-byte boundaries (where `n = 1, 2, 4, 8, or 16` with `-Zp[n]`).

When you specify this option, each structure member after the first is stored on either the size of the member type or `n`-byte boundaries (where `n = 1, 2, 4, 8, or 16`), whichever is smaller.

For example, to specify 2 bytes as the packing boundary (or alignment constraint) for all structures and unions in the file `prog1.f`, use the following command:

**ifort -Zp2 prog1.f**

The default for IA-32 and Itanium-based systems is `-align rec8byte` or `-Zp8`. The `-Zp16` option enables you to align Fortran structures such as common blocks. For Fortran structures, see `STRUCTURE` statement in *Intel® Fortran Language Reference Manual*.

If you specify `-Zp` (omit `n`), structures are packed at 8-byte boundary.

### **-align and -pad**

The `-align` option is a front-end option that changes alignment of variables in a common block.

Example:

```
common
/block1/ch,doub,ch1,int
integer int
character(len=1) ch,
ch1
double precision doub
end
```

The `-align` option enables padding inserted to assure alignment of `doub` and `int` on natural alignment boundaries. The `-noalign` option disables padding.

The `-align` option applies mainly to structures. It analyzes and reorders memory layout for variables and arrays and basically functions as `-Zp{n}`. You can disable either option with `-noalign`.

For `-align keyword` options, see *Command-line Options*.

The `-pad` option is effectively not different from `-align` when applied to structures and derived types. However, the scope of `-pad` is greater because it applies also to common blocks, derived types, sequence types, and VAX\* structures.

### **Recommendations on Controlling Alignment with Options**

The following options control whether the Intel Fortran compiler adds padding (when needed) to naturally align multiple data items in common blocks, derived-type data, and Intel Fortran record structures:

- By default (with `-O2`), the `-align commons` option requests that data in common blocks be aligned on up to 4-byte boundaries, by adding padding bytes as needed.

The `-align nocommons` arbitrarily aligns the bytes of common block data. In this case, unaligned data can occur unless the order of data items specified in the `COMMON` statement places the largest numeric data item first, followed by the next largest numeric data (and so on), followed by any character data.

- By default (with `-O2`), the `-align dcommons` option requests that data in common blocks be aligned on up to 8-byte boundaries, by adding padding bytes as needed.  
The `-align nodcommons` arbitrarily aligns the bytes of data items in a common data.

Specify the `-align dcommons` option for applications that use common blocks, unless your application has no unaligned data or, if the application might have unaligned data, all data items are four bytes or smaller. For applications that use common blocks where all data items are four bytes or smaller, you can specify `-align commons` instead of `-align dcommons`.

- The `-align norecords` option requests that multiple data items in derived-type data and record structures (an Intel Fortran extension) be aligned arbitrarily on byte boundaries instead of being naturally aligned. The default is `-align records`.
- The `-align records` option requests that multiple data items in record structures (extension) and derived-type data without the `SEQUENCE` statement be naturally aligned, by adding padding bytes as needed.
- The `-align recnbyte` option requests that fields of records and components of derived types be aligned on either the size byte boundary specified or the boundary that will naturally align them, whichever is smaller. This option does not affect whether common blocks are naturally aligned or packed.
- The `-align sequence` option controls alignment of derived types with the `SEQUENCE` attribute.

The `-align nosequence` option means that derived types with the `SEQUENCE` attribute are packed regardless of any other alignment rules. Note that `-align none` implies `-align nosequence`.

The `-align sequence` option means that derived types with the `SEQUENCE` attribute obey whatever alignment rules are currently in use. Consequently, since `-align record` is a default value, then `-align sequence` alone on the command line will cause the fields in these derived types to be naturally aligned.

The default behavior is that multiple data items in derived-type data and record structures *will* be naturally aligned; data items in common blocks will not ( `-align records` with `-align nocommons` ). In derived-type data, using the

SEQUENCE statement prevents `-align records` from adding needed padding bytes to naturally align data items.

## Symbol Visibility Attribute Options

Applications that do not require symbol preemption or position-independent code can obtain a performance benefit by taking advantage of the generic ABI visibility attributes.

### Note

The visibility options are supported by both IA-32 and Itanium compilers, but currently the optimization benefits are for Itanium-based systems only.

## Global Symbols and Visibility Attributes

A global symbol is a symbol that is visible outside the compilation unit in which it is declared (compilation unit is a single-source file with its include files). Each global symbol definition or reference in a compilation unit has a visibility attribute that controls how it may be referenced from outside the component in which it is defined. The values for visibility are defined in the table that follows.

EXTERN	The compiler must treat the symbol as though it is defined in another component. This means that the compiler must assume that the symbol will be overridden (preempted) by a definition of the same name in another component. (See Symbol Preemption.) If a function symbol has external visibility, the compiler knows that it must be called indirectly and can inline the indirect call stub.
DEFAULT	Other components can reference the symbol. Furthermore, the symbol definition may be overridden (preempted) by a definition of the same name in another component.
PROTECTED	Other components can reference the symbol, but it cannot be preempted by a definition of the same name in another component.
HIDDEN	Other components cannot directly reference the symbol. However, its address might be passed to other components indirectly; for example, as an argument to a call to a function in another component, or by having its address stored in a data item referenced by a function in another component.
INTERNAL	The symbol cannot be referenced outside the component where it is defined, either directly or indirectly.



**Note**

Visibility applies to both references and definitions. A symbol reference's visibility attribute is an assertion that the corresponding definition will have that visibility.

**Symbol Preemption and Optimization**

Sometimes programmers need to use some of the functions or data items from a shareable object, but at the same time, they need to replace other items with definitions of their own. For example, an application may need to use the standard run-time library shareable object, `libc.so`, but to use its own definitions of the heap management routines `malloc` and `free`. In this case it is important that calls to `malloc` and `free` within `libc.so` use the user's definition of the routines and not the definitions in `libc.so`. The user's definition should then override, or *preempt*, the definition within the shareable object.

This functionality of redefining the items in shareable objects is called symbol preemption. When the run-time loader loads a component, all symbols within the component that have default visibility are subject to preemption by symbols of the same name in components that are already loaded. Note that since the main program image is always loaded first, none of the symbols it defines will be preempted (redefined).

The possibility of symbol preemption inhibits many valuable compiler optimizations because symbols with default visibility are not bound to a memory address until run-time. For example, calls to a routine with default visibility cannot be inlined because the routine might be preempted if the compilation unit is linked into a shareable object. A preemptable data symbol cannot be accessed using GP-relative addressing because the name may be bound to a symbol in a different component; and the GP-relative address is not known at compile time.

Symbol preemption is a rarely used feature and has negative consequences for compiler optimization. For this reason, by default the compiler treats all global symbol definitions as non-preemptable (protected visibility). Global references to symbols defined in another compilation unit are assumed by default to be preemptable (default visibility). In those rare cases where all global definitions as well as references need to be preemptable, specify the `-fpic` option to override this default.

**Specifying Symbol Visibility Explicitly**

The Intel Fortran Compiler has the visibility attribute options that provide command-line control of the visibility attributes as well as a source syntax to set the complete range of these attributes. The options ensure immediate access to the feature without depending on header file modifications. The visibility options

cause all global symbols to get the visibility specified by the option. There are two variety of options to specify symbol visibility explicitly:

```
-fvisibility=keyword  
-fvisibility-keyword=file
```

The first form specifies the default visibility for global symbols. The second form specifies the visibility for symbols that are in a file (this form overrides the first form).

The *file* is the pathname of a file containing the list of symbols whose visibility you want to set; the symbols are separated by whitespace (spaces, tabs, or newlines).

In both options, the *keyword* is: *extern*, *default*, *protected*, *hidden*, and *internal*, see definitions above.

### **Note**

These two ways to explicitly set visibility are mutually exclusive: you may use the visibility attribute on the declaration, or specify the symbol name in a *file*, but not both.

The option `-fvisibility-keyword=file` specifies the same visibility attribute for a number of symbols using one of the five command line options corresponding to the *keyword*:

```
-fvisibility-extern=file  
-fvisibility-default=file  
-fvisibility-protected=file  
-fvisibility-hidden=file  
-fvisibility-internal=file
```

where *file* is the pathname of a file containing a list of the symbol names whose visibility you wish to set; the symbol names in the *file* are separated by either blanks, tabs, or newlines. For example, the command line option:

```
-fvisibility-protected=prot.txt
```

where file `prot.txt` contains symbols *a*, *b*, *c*, *d*, and *e* sets protected visibility for symbols *a*, *b*, *c*, *d*, and *e*. This has the same effect as declared attribute `visibility=protected` on the declaration for each of the symbols.

### Specifying Visibility without Symbol File, `-fvisibility=keyword`

This option sets the visibility for symbols not specified in a visibility list file and that do not have `visibility` attribute in their declaration. If no symbol file

option is specified, all symbols will get the specified attribute. Command line example:

```
ifort -fvisibility=protected a.f
```

You can set the default visibility for symbols using one of the following command line options:

```
-fvisibility=external
-fvisibility=default
-fvisibility=protected
-fvisibility=hidden
-fvisibility=internal
```

The above options are listed in the order of precedence: explicitly setting the visibility to `external`, by using either the attribute syntax or the command line option, overrides any setting to `default`, `protected`, `hidden`, or `internal`. Explicitly setting the visibility to `default` overrides any setting to `protected`, `hidden`, or `internal` and so on.

The visibility attribute `default` enables compiler to change the default symbol visibility and then set the default attribute on functions and variables that require the default setting. Since `internal` is a processor-specific attribute, it may not be desirable to have a general option for it.

In the combined command-line options

```
-fvisibility=protected -fvisibility-default=prot.txt
```

file `prot.txt` (see above) causes all global symbols except `a`, `b`, `c`, `d`, and `e` to have protected visibility. Those five symbols, however, will have default visibility and thus be preemptable.

## Visibility-related Options

### **-fminshared**

Directs to treat the compilation unit as a component of a main program and not to link it as a part of a shareable object.

Since symbols defined in the main program cannot be preempted, this enables the compiler to treat symbols declared with default visibility as though they have protected visibility. It means that

`-fminshared` implies `-fvisibility=protected`. The compiler need not generate position-independent code for the main program. It can use absolute addressing, which may reduce the size of the global offset table (GOT) and may reduce memory traffic.

**-fpic**

Specifies full symbol preemption. Global symbol definitions as well as global symbol references get default (that is, preemptable) visibility unless explicitly specified otherwise. Generates position-independent code.

**-fno\_common**

Instructs the compiler to treat common symbols as global definitions and to allocate memory for each symbol at compile time. This may permit the compiler to use the more efficient GP-relative addressing mode when accessing the symbol.

Normally a Fortran uninitialized common block declaration with no initializer and without the `extern` or `static` keywords. For example,

```
integer i
```

is represented as a common symbol. Such a symbol is treated as external reference, except that if no other compilation unit has a global definition for the name, the linker allocates memory for it.

## Optimizing Different Application Types

### Optimizing Different Application Types Overview

This section discusses the command-line options `-O0`, `-O1`, `-O2` (or `-O`), and `-O3`. The `-O0` option disables optimizations. Each of the other three turns on several compiler capabilities. To specify one of these optimizations, take into consideration the nature and structure of your application as indicated in the more detailed description of the options.

In general terms, `-O1`, `-O2` (or `-O`), and `-O3` optimize as follows:

`-O1` : code size and locality

`-O2` (or `-O`) : code speed; this is the default option

`-O3`: enables `-O2` with more aggressive optimizations.

`-fast`: enables `-O3` and `-ipo` to enhance speed across the entire program.

These options behave similarly on IA-32 and Itanium® architectures, with some specifics that are detailed in the sections that follow.

### Setting Optimizations with `-On` Options

The following table details the effects of the `-O0`, `-O1`, `-O2`, `-O3`, and `-fast` options. The table first describes the characteristics shared by both IA-32 and Itanium architectures and then explicitly describes the specifics (if any) of the `-On` and `-fast` options' behavior on each architecture.

Option	Effect
<code>-O0</code>	Disables <code>-On</code> optimizations. On IA-32 systems, this option sets the <code>-fp</code> option.
<code>-O1</code>	Optimizes to favor code size and code locality. Disables loop unrolling. May improve performance for applications with very large code size, many branches, and execution time not dominated by code within loops. In most cases, <code>-O2</code> is recommended over <code>-O1</code> . <b>IA-32 systems:</b> Disables intrinsics inlining to reduce code size. Enables optimizations for speed. Also disables intrinsic recognition and the <code>-fp</code> option. This option is the same as the <code>-O2</code> option.

	<p><b>Itanium-based systems:</b> Disables software pipelining and global code scheduling. Enables optimizations for server applications (straight-line and branch-like code with flat profile). Enables optimizations for speed, while being aware of code size. For example, this option disables software pipelining and loop unrolling.</p>
-O2, -O	<p>This option is the default for optimizations. However, if <code>-g</code> is specified, the default is <code>-O0</code>. Optimizes for code speed. This is the generally recommended optimization level.</p> <p>On <b>IA-32 systems</b>, this option is the same as the <code>-O1</code> option.</p> <p><b>Itanium-based systems:</b> Enables optimizations for speed, including global code scheduling, software pipelining, predication, and speculation.</p> <p>On these systems, the <code>-O2</code> option enables inlining of intrinsics. It also enables the following capabilities for performance gain: constant propagation, copy propagation, dead-code elimination, global register allocation, global instruction scheduling and control speculation, loop unrolling, optimized code selection, partial redundancy elimination, strength reduction/induction variable simplification, variable renaming, exception handling optimizations, tail recursions, peephole optimizations, structure assignment lowering and optimizations, and dead store elimination.</p>
-O3	<p>Enables <code>-O2</code> optimizations and in addition, enables more aggressive optimizations such as prefetching, scalar replacement, and loop and memory access transformations. Enables optimizations for maximum speed, but does not guarantee higher performance unless loop and memory access transformation take place. The <code>-O3</code> optimizations may slow down code in some cases compared to <code>-O2</code> optimizations. Recommended for applications that have loops that heavily use floating point calculations and process large data sets.</p> <p><b>IA-32 systems:</b> In conjunction with <code>-ax{K W N B P}</code> or -</p>

	<p><math>x\{K W N B P\}</math> options, this option causes the compiler to perform more aggressive data dependency analysis than for <code>-O2</code>. This may result in longer compilation times.</p> <p>On <b>Itanium-based systems</b>, enables optimizations for technical computing applications (loop-intensive code): loop optimizations and data prefetch.</p>
<code>-fast</code>	<p>This option is a single, simple method to enable a collection of optimizations for run-time performance. Sets the following options that can improve run-time performance:</p> <p><code>-O3</code>: maximum speed and high-level optimizations, see above</p> <p><code>-ipo</code>: enables interprocedural optimizations across files</p> <p><code>-static</code>: prevents linking with shared libraries</p> <p>Provides a shortcut that requests several important compiler optimizations. To override one of the options set by <code>-fast</code>, specify that option after the <code>-fast</code> option on the command line.</p> <p>The options set by the <code>-fast</code> option may change from release to release.</p> <p><b>IA-32 systems:</b></p> <p>In conjunction with <code>-ax{K W N B P}</code> or <code>-x{K W N B P}</code> options, this option provides the best run-time performance.</p>

## Restricting Optimizations

The following options restrict or preclude the compiler's ability to optimize your program:

<code>-O0</code>	Disables optimizations. Enables <code>-fp</code> option.
<code>-mp</code>	Restricts optimizations that cause some minor loss or gain of

	precision in floating-point arithmetic to maintain a declared level of precision and to ensure that floating-point arithmetic more nearly conforms to the ANSI and IEEE* standards. See <a href="#">-mp option</a> for more details.
-g	Specifying the -g option turns off the default -O2 option and makes -O0 the default unless -O2 (or -O1 or -O3) is explicitly specified in the command line together with -g. See <a href="#">Optimizations and Debugging</a> .
-nolib_inline	Disables inline expansion of intrinsic functions.

For more information on ways to restrict optimization, see [Using -ip with -Qoption specifiers](#).



## Floating-point Arithmetic Optimizations

### Options Used for IA-32 and Itanium® Architectures

The options described in this section all provide optimizations with varying degrees of precision in floating-point (FP) arithmetic for IA-32 and Itanium® compilers.

The `-mp1` (IA-32 only) and `-mp` options improve floating-point precision, but also affect the application performance. See more details about these options in [Improving/Restricting FP Arithmetic Precision](#).

The FP options provide optimizations with varying degrees of precision in floating-point arithmetic. The option that disables these optimizations is `-O0`.

#### `-mp` Option

Use `-mp` to limit floating-point optimizations and maintain declared precision. For example, the Intel® Fortran Compiler can change floating-point division computations into multiplication by the reciprocal of the denominator. This change can alter the results of floating point division computations slightly. The `-mp` switch may slightly reduce execution speed. See [Improving/Restricting FP Arithmetic Precision](#) for more detail.

#### `-mp1` Option (IA-32 Only)

Use the `-mp1` option to restrict floating-point precision to be closer to declared precision with less impact to performance than with the `-mp` option. The option will ensure the out-of-range check of operands of transcendental functions and improve accuracy of floating-point compares.

#### Flushing to Zero Denormal Values, `-ftz[-]`

Option `-ftz[-]` flushes denormal results to zero when the application is in the gradual underflow mode. Flushing the denormal values to zero with `-ftz` may improve performance of your application.

#### Default

The default status of `-ftz[-]` is OFF. By default, the compiler lets results gradually underflow. With the default `-O2` option, `-ftz[-]` is OFF.

#### `-ftz[-]` on Itanium-based systems

On Itanium-based systems only, the `-O3` option turns on `-ftz`.

If the `-ftz` option produces undesirable results of the numerical behavior of your program, you can turn the FTZ mode off by using `-ftz-` in the command line while still benefiting from the `-O3` optimizations:

```
ifort -O3 -ftz- myprog.f
```

## Usage

- Use this option if the denormal values are not critical to application behavior.
- `-ftz[-]` only needs to be used on the source that contains the `main()` program to turn the FTZ mode on. The initial thread, and any threads subsequently created by that process, will operate in FTZ mode.

## Results

The `-ftz[-]` option affects the results of floating underflow as follows:

- `-ftz-` results in **gradual** underflow to 0: the result of a floating underflow is a denormalized number or a zero.
- `-ftz` results in **abrupt** underflow to 0: the result of a floating underflow is set to zero and execution continues. `-ftz` also makes a denormal value used in a computation be treated as a zero so no floating invalid exception occurs. On Itanium-based systems, the `-O3` option sets the abrupt underflow to zero (`-ftz` is on). At lower optimization levels, gradual underflow to 0 is the default on the Itanium-based systems.

On IA-32, setting abrupt underflow by `-ftz` may improve performance of SSE/SSE2 instructions, while it does not affect either performance or numerical behavior of x87 instructions. Thus, `-ftz` will have no effect unless you select `-x{}` or `-ax{}` options, which activate instructions of the more recent IA-32 Intel processors.

On Itanium-based processors, gradual underflow to 0 can degrade performance. Using higher optimization levels to get the default abrupt underflow or explicitly setting `-ftz` improves performance.

`-ftz` may improve performance on Itanium® 2 processor, even in the absence of actual underflow, most frequently for single-precision code.

## Using the Floating-point Exception Handling, `-fpe n`

Use the `-fpe n` option to control the handling of exceptions. The `-fpe n` option controls floating-point exceptions according to the value of `n`.

The following are the kinds of floating-point exceptions:

- Floating overflow: the result of a computation is too large for the floating-point data type. The result is replaced with the exceptional value Infinity with the proper "+" or "-" sign. For example,  $1E30 * 1E30$  overflows single-precision floating-point value and results in a +Infinity;  $-1E30 * 1E30$  results in a -Infinity.
- Floating divide-by-zero: if the computation is  $0.0 / 0.0$ , the result is the exceptional value NaN (Not a Number), a value that means the computation was not successful. If the numerator is not 0.0, the result is a signed Infinity.
- Floating underflow: the result of a computation is too small for the floating-point type. Each floating-point type (32-, 64-, and 128-bit) has a denormalized range where very small numbers can be represented with some loss of precision. For example, the lower bound for normalized single-precision floating-point value is approximately  $1E-38$ ; the lower bound for denormalized single-precision floating-point value is  $1E-45$ .  $1E-30 / 1E10$  underflows the normalized range but not the denormalized range so the result is the denormal exceptional value  $1E-40$ .  $1E-30 / 1E30$  underflows the entire range and the result is zero. This is known as gradual underflow to 0.
- Floating invalid: when the exceptional value (signed Infinities, NaN, denormal) is used as input to a computation, the result is also a NaN.

The `-fpen` option allows some control over the results of floating-point exception handling at run time for the main program.

- `-fpe0` restricts floating-point exceptions as follows:
  - Floating overflow, floating divide-by-zero, and floating invalid cause the program to print an error message and abort.
  - If a floating underflow occurs, the result is set to zero and execution continues. This is called abrupt underflow to 0.
- `-fpe1` restricts only floating underflow:
  - Floating overflow, floating divide-by-zero, and floating invalid produce exceptional values (NaN and signed Infinities) and execution continues.
  - If a floating underflow occurs, the result is set to zero and execution continues.
- The default is `-fpe3` on both IA-32 and Itanium-based processors. This allows full floating-point exception behavior:
  - Floating overflow, floating divide-by-zero, and floating invalid produce exceptional values (NaN and signed Infinities) and execution continues.
  - Floating underflow is gradual: denormalized values are produced until the result becomes 0.

The `-fpen` only affects the Fortran main program. The floating-point exception behavior set by the Fortran main program is in effect throughout the execution of

the entire program. If the main program is not Fortran, you can use the Fortran intrinsic `FOR_SET_FPE` to set the floating-point exception behavior.

When compiling different routines in a program separately, you should use the same value of `n` in `-fpn`.

For more information, refer to the Intel Fortran Compiler User's Guide for Linux\*Systems, Volume I, section "Controlling Floating-point Exceptions."

## Floating-point Arithmetic Precision for IA-32 Systems

### `-prec_div` Option

The Intel® Fortran Compiler can change floating-point division computations into multiplication by the reciprocal of the denominator. Use `-prec_div` to disable floating point division-to-multiplication optimization resulting in more accurate division results. May have speed impact.

### `-pc{32|64|80}` Option

Use the `-pc{32|64|80}` option to enable floating-point significand precision control. Some floating-point algorithms, created for specific IA-32 and Itanium®-based systems, are sensitive to the accuracy of the significand or fractional part of the floating-point value. Use appropriate version of the option to round the significand to the number of bits as follows:

`-pc32`: 24 bits (single precision)

`-pc64`: 53 bits (double precision)

`-pc80`: 64 bits (extended precision)

The default version is `-pc80` for full floating-point precision.

This option enables full optimization. Using this option does not have the negative performance impact of using the `-mp` option because only the fractional part of the floating-point value is affected. The range of the exponent is not affected.

### Note

This option only has effect when the module being compiled contains the main program.

 **Caution**

A change of the default precision control or rounding mode (for example, by using the `-pc32` option or by user intervention) may affect the results returned by some of the mathematical functions.

### Rounding Control, `-rcd`, `-fp_port`

The Intel Fortran Compiler uses the `-rcd` option to disable changing of rounding mode for floating-point-to-integer conversions.

The system default floating-point rounding mode is round-to-nearest. This means that values are rounded during floating-point calculations. However, the Fortran language requires floating-point values to be truncated when a conversion to an integer is involved. To do this, the compiler must change the rounding mode to truncation before each floating-point conversion and change it back afterwards.

The `-rcd` option disables the change to truncation of the rounding mode for all floating-point calculations, including floating-point-to-integer conversions. Turning on this option can improve performance, but floating-point conversions to integer will not conform to Fortran semantics.

You can also use the `-fp_port` option to round floating-point results at assignments and casts. May cause some speed impact, but also makes sure that rounding to the user-declared precision at assignments is always done. The `-mp1` option implies `-fp_port`.

### Floating-point Arithmetic Precision for Itanium®-based Systems

The following Intel® Fortran Compiler options enable you to control the compiler optimizations for floating-point computations on Itanium®-based systems.

#### Contraction of FP Multiply and Add/Subtract Operations

`-IPF_fma[-]` enables or disables the contraction of floating-point multiply and add/subtract operations into a single operations. Unless `-mp` is specified, the compiler tries to contract these operations whenever possible. The `-mp` option disables the contractions.

`-IPF_fma` and `-IPF_fma-` can be used to override the default compiler behavior. For example, a combination of `-mp` and `-IPF_fma` enables the compiler to contract operations:

```
ifort -mp -IPF_fma myprog.f
```

## FP Speculation

`-IPF_fp_speculationmode` sets the compiler to speculate on floating-point operations in one of the following *modes*:

*fast*: sets the compiler to speculate on floating-point operations; this is the default.

*safe*: enables the compiler to speculate on floating-point operations only when it is safe;

*strict*: enables the compiler's speculation on floating-point operations preserving floating-point status in all situations. In the current version, this mode disables the speculation of floating-point operations (same as `off`).

*off*: disables the speculation on floating-point operations.

## FP Operations Evaluation

`-IPFflt_eval_method{0|2}` option directs the compiler to evaluate the expressions involving floating-point operands in the following way:

`-IPFflt_eval_method0` directs the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types declared in the program.

`-IPFflt_eval_method2` is not supported in the current version.

## Controlling Accuracy of the FP Results

`-IPFfltacc` disables the optimizations that affect floating-point accuracy. The default is `-IPFfltacc-` to enable such optimizations.

The Itanium® compiler may reassociate floating-point expressions to improve application performance. Use `-IPFfltacc` or `-mp` to disable or restrict these floating-point optimizations.

## Improving/Restricting FP Arithmetic Precision

The `-mp` and `-mp1` (`-mp1` is for IA-32 only) options maintain and restrict, respectively, floating-point precision, but also affect the application performance. The `-mp1` option causes less impact on performance than the `-mp` option. `-mp1` ensures the out-of-range check of operands of transcendental functions and improve accuracy of floating-point compares. For IA-32 systems, the `-mp` option implies `-mp1`; `-mp1` implies `-fp_port`. `-mp` slows down performance the most of these three, `-fp_port` the least of these three.

The `-mp` option restricts some optimizations to maintain declared precision and to ensure that floating-point arithmetic conforms more closely to the ANSI and IEEE\* standards. This option causes more frequent stores to memory, or disallow some data from being register candidates altogether. The Intel architecture normally maintains floating point results in registers. These registers are 80 bits long, and maintain greater precision than a double-precision number. When the results have to be stored to memory, rounding occurs. This can affect accuracy toward getting more of the "expected" result, but at a cost in speed. The `-pc{32|64|80}` option (IA-32 only) can be used to control floating point accuracy and rounding, along with setting various processor IEEE flags.

For most programs, specifying the `-mp` option adversely affects performance. If you are not sure whether your application needs this option, try compiling and running your program both with and without it to evaluate the effects on performance versus precision.

Specifying this option has the following effects on program compilation:

- On **IA-32 systems**, floating-point user variables declared as floating-point types are not assigned to registers.
- On **Itanium®-based systems**, floating-point user variables may be assigned to registers. The expressions are evaluated using precision of source operands. The compiler will not use Floating-point Multiply and Add (FMA) function to contract multiply and add/subtract operations in a single operation. The contractions can be enabled by using `-IPF_fma` option. The compiler will not speculate on floating-point operations that may affect the floating-point state of the machine. See Floating-point Arithmetic Precision for Itanium-based Systems.
- Floating-point arithmetic comparisons conform to IEEE 754.
- The exact operations specified in the code are performed. For example, division is never changed to multiplication by the reciprocal.
- The compiler performs floating-point operations in the order specified without reassociation.
- The compiler does not perform the constant folding on floating-point values. Constant folding also eliminates any multiplication by 1, division by 1, and addition or subtraction of 0. For example, code that adds 0.0 to a number is executed exactly as written. Compile-time floating-point arithmetic is not performed to ensure that floating-point exceptions are also maintained.

For IA-32 systems, whenever an expression is spilled, it is spilled as 80 bits (`EXTENDED PRECISION`), not 64 bits (`DOUBLE PRECISION`). Floating-point operations conform to IEEE 754. When assignments to type `REAL` and `DOUBLE PRECISION` are made, the precision is rounded from 80 bits (`EXTENDED`) down to 32 bits (`REAL`) or 64 bits (`DOUBLE PRECISION`).

When you do not specify `-O0`, the extra bits of precision are not always rounded away before the variable is reused.

- Even if vectorization is enabled by the `-xK | W | B | P` options, the compiler does not vectorize reduction loops (loops computing the dot product) and loops with mixed precision types. Similarly, the compiler does not enable certain loop transformations. For example, the compiler does not transform reduction loops to perform partial summation or loop interchange.



# Optimizing for Specific Processors

## Optimizing for Specific Processors Overview

This section describes targeting a processor and processor dispatch and extensions support options.

The options `-tpp{5|6|7}` optimize for the IA-32 processors, and the options `-tpp{1|2}` optimize for the Itanium® processor family. The options `-x{K|W|N|B|P}` and `-ax{K|W|N|B|P}` generate code that is specific to processor-instruction extensions.

Note that you can run your application on the latest processor-based systems, like Intel® Pentium® M processor or Intel processors code-named "Prescott" and still gear your code to any of the previous processors specified by `N/W` or `K` versions of the `-x` and `-ax` options.

## Targeting a Processor, `-tpp{n}`

The `-tpp{n}` optimizes your application's performance for specific Intel processors. This option generates code that is tuned for the processor associated with its version. For example, `-tpp7` generates code optimized for running on Intel® Pentium® 4, Intel® Xeon(TM), Intel® Pentium® M processors and Intel processors code-named "Prescott," and `-tpp2` generates code optimized for running on Itanium® 2 processor.

The `-tpp{n}` option always generates code that is backwards compatible with Intel® processors of the same family. This means that code generated with `-tpp7` will run correctly on Pentium Pro or Pentium III processors, possibly just not quite as fast as if the code had been compiled with `-tpp6`. Similarly, code generated with `-tpp2` will run correctly on Itanium processor, but possibly not quite as fast as if it had been generated with `-tpp1`.

## Processors for IA-32 Systems

The `-tpp5`, `-tpp6`, and `-tpp7` options optimize your application's performance for a **specific** Intel IA-32 processor as listed in the table below. The resulting binaries will also run correctly on any of the processors mentioned in the table.

Option	Optimizes your application for...
<code>-tpp5</code>	Intel® Pentium® and Pentium® with MMX(TM) technology processor
<code>-tpp6</code>	Intel® Pentium® Pro, Pentium® II and Pentium® III processors
<code>-tpp7</code>	Intel Pentium 4 processors, Intel® Xeon(TM) processors, Intel®

(default)	Pentium® M processors, and Intel processors code-named "Prescott"
-----------	---

### Example

The invocations listed below each result in a compiled binary of the source program `prog.f` optimized for Pentium 4 and Intel Xeon processors by default. The same binary will also run on Pentium, Pentium Pro, Pentium II, and Pentium III processors.

```
ifort prog.f
```

```
ifort -tpp7 prog.f
```

However if you intend to target your application specifically to the Intel Pentium and Pentium with MMX technology processors, use the `-tpp5` option:

```
ifort -tpp5 prog.f
```

### Processors for Itanium®-based Systems

The `-tpp1` and `-tpp2` options optimize your application's performance for a **specific** Intel IA-32 processor as listed in the table below. The resulting binaries will also run correctly on both processors mentioned in the table.

Option	Optimizes your application for...
<code>-tpp1</code>	Intel® Itanium® processor
<code>-tpp2</code> (default)	Intel® Itanium® 2 processor

### Example

The following invocation results in a compiled binary of the source program `prog.f` optimized for the Itanium 2 processor by default. The same binary will also run on Itanium processors.

```
ifort prog.f
```

```
ifort -tpp2 prog.f
```

However if you intend to target your application specifically to the Intel Itanium processor, use the `-tpp1` option:

```
ifort -tpp1 prog.f
```

## Processor-specific Optimization (IA-32 only)

The `-x{K|W|N|B|P}` options target your program to run on a specific Intel processor. The resulting code might contain unconditional use of features that are not supported on other processors.

Option	Optimizes for...
<code>-xK</code>	Intel® Pentium® III and compatible Intel processors.
<code>-xW</code>	Intel Pentium 4 and compatible Intel processors.
<code>-xN</code>	Intel Pentium 4 and compatible Intel Processors. When the main program is compiled with this option, it will detect non-compatible processors and generate an error message during execution. This option also enables new optimizations in addition to Intel processor specific-optimizations.
<code>-xB</code>	Intel® Pentium® M and compatible Intel processors. When the main program is compiled with this option, it will detect non-compatible processors and generate an error message during execution. This option also enables new optimizations in addition to Intel processor-specific optimizations.
<code>-xP</code>	Intel processors code-named "Prescott." When the main program is compiled with this option, it will detect non-compatible processors and generate an error message during execution. This option also enables new optimizations in addition to Intel processor-specific optimizations.

To execute a program on x86 processors not provided by Intel Corporation, do not specify the `-x{K|W|N|B|P}` option.

### Example

The invocation below compiles `myprog.f` for Intel Pentium 4 and compatible processors. The resulting binary might not execute correctly on Pentium, Pentium Pro, Pentium II, Pentium III, or Pentium with MMX technology processors, or on x86 processors not provided by Intel corporation.

```
ifort -xW myprog.f
```

### Caution

If a program compiled with `-x{K|W|N|B|P}` is executed on a non-compatible processor, it might fail with an illegal instruction exception, or display other unexpected behavior. Executing programs compiled with `-xN`, `-xB`, or `-xP` on unsupported processors (see table above) will display the following run-time error:

Fatal error: This program was not built to run on the processor in your system.

## Automatic Processor-specific Optimization (IA-32 only)

The `-ax{K|W|N|B|P}` options direct the compiler to find opportunities to generate separate versions of functions that take advantage of features that are specific to the specified Intel processor. If the compiler finds such an opportunity, it first checks whether generating a processor-specific version of a function is likely to result in a performance gain. If this is the case, the compiler generates both a processor-specific version of a function and a generic version of the function. The generic version will run on any IA-32 processor.

At run time, one of the versions is chosen to execute, depending on the Intel processor in use. In this way, the program can benefit from performance gains on more advanced Intel processors, while still working properly on older IA-32 processors.

The disadvantages of using `-ax{K|W|N|B|P}` are:

- The size of the compiled binary increases because it contains processor-specific versions of some of the code, as well as a generic version of the code.
- Performance is affected slightly by the run-time checks to determine which code to use.

### Note

Applications that you compile to optimize themselves for specific processors in this way will execute on any Intel IA-32 processor. If you specify both the `-x` and `-ax` options, the `-x` option forces the generic code to execute only on processors compatible with the processor type specified by the `-x` option.

Option	Optimizes Your Code for...
<code>-axK</code>	Intel® Pentium® III and compatible Intel processors.
<code>-axW</code>	Intel Pentium 4 and compatible Intel processors.
<code>-axN</code>	Intel Pentium 4 and compatible Intel processors. This option also enables new optimizations in addition to Intel processor-specific optimizations.
<code>-axB</code>	Intel Pentium M and compatible Intel processors. This option also enables new optimizations in addition to Intel processor-specific optimizations.

-axP	Intel processors code-named "Prescott." This option also enables new optimizations in addition to Intel processor-specific optimizations.
------	---

### Example

The compilation below generates a single executable that includes:

- a generic version for use on any IA-32 processor
- a version optimized for Intel Pentium III processors, as long as there is a performance benefit.
- a version optimized for Intel Pentium 4 processors, as long as there is a performance benefit.

```
ifort -axKW prog.f90
```

### Processor-specific Run-time Checks, IA-32 Systems

The Intel Fortran Compiler optimizations take effect at run-time. For IA-32 systems, the compiler enhances processor-specific optimizations by inserting in the main routine a code segment that performs run-time checks described below.

#### Check for Supported Processor with -xB , -xB, or -xP

To prevent from execution errors, the compiler inserts code in the main routine of the program to check for proper processor usage. Programs compiled with options -xN, -xB, or -xP check at run-time whether they are being executed on the Intel Pentium® 4, Intel® Pentium® M processor or the Intel processor code-named "Prescott," respectively, or a compatible Intel processor. If the program is not executed on one of these processors, the program terminates with an error.

### Example

To optimize a program `foo.f90` for an Intel processor code-named "Prescott," issue the following command:

```
ifort -xP foo.f90 -o foo.exe
```

`foo.exe` aborts if it is executed on a processor that is not validated to support the Intel processor code-named "Prescott," such as the Intel Pentium 4 processor (to account for the fact that "Prescott" may be a Pentium 4 processor with some feature enabling).

If you intend to run your programs on multiple IA-32 processors, do not use the `-x{ }` options that optimize for processor-specific features; consider using `-ax{ }` to attain processor-specific performance and portability among different processors.

## Setting FTZ and DAZ Flags

Previously, the default status of the flags flush-to-zero (FTZ) and denormals-are-zero (DAZ) for IA-32 processors were off by default. However, even at the cost of losing IEEE compliance, turning these flags on significantly increases the performance of programs with denormal floating-point values in the gradual underflow mode run on the most recent IA-32 processors. Hence, for the Intel Pentium III, Pentium 4, Pentium M, Intel processor code-named "Prescott," and compatible IA-32 processors, the compiler's default behavior is to turn these flags on. The compiler inserts code in the program to perform a run-time check for the processor on which the program runs to verify it is one of the afore-listed Intel processors.

- Executing a program on a Pentium III processor enables the FTZ flag, but not DAZ.
- Executing a program on an Intel Pentium M processor or "Prescott" processor enables both the FTZ and DAZ flags.

These flags are only turned on by Intel processors that have been validated to support them.

For non-Intel processors, the flags can be set manually by calling the following Intel Fortran intrinsic:

```
RESULT = FOR_SET_FPE (FOR_M_ABRUPT_UND).
```

## Interprocedural Optimizations (IPO)

### IPO Overview

Use `-ip` and `-ipo` to enable interprocedural optimizations (IPO), which enable the compiler to analyze your code to determine where you can benefit from the optimizations listed in tables that follow.

#### IA-32 and Itanium®-based applications

Optimization	Affected Aspect of Program
inline function expansion	calls, jumps, branches, and loops
interprocedural constant propagation	arguments, global variables, and return values
monitoring module-level static variables	further optimizations, loop invariant code
dead code elimination	code size
propagation of function characteristics	call deletion and call movement
multifile optimization	affects the same aspects as <code>-ip</code> , but across multiple files

#### IA-32 applications only

Optimization	Affected Aspect of Program
passing arguments in registers	calls, register usage
loop-invariant code motion	further optimizations, loop invariant code

Inline function expansion is one of the main optimizations performed by the interprocedural optimizer. For function calls that the compiler believes are frequently executed, the compiler might decide to replace the instructions of the call with code for the function itself.

With `-ip`, the compiler performs inline function expansion for calls to procedures defined within the current source file. However, when you use `-ipo` to specify multifile IPO, the compiler performs inline function expansion for calls to procedures defined in separate files.

To disable the IPO optimizations, use the `-O0` option.

## **Caution**

The `-ip` and `-ipo` options can in some cases significantly increase compile time and code size.

### **Option `-auto_ilp32` for Itanium Compiler**

On Itanium-based systems, the `-auto_ilp32` option requires interprocedural analysis over the whole program. This optimization allows the compiler to use 32-bit pointers whenever possible as long as the application does not exceed a 32-bit address space. Using the `-auto_ilp32` option on programs that exceed 32-bit address space might cause unpredictable results during program execution.

Because this optimization requires interprocedural analysis over the whole program, you must use the `-auto_ilp32` option with the `-ipo` option.

## **Multifile IPO**

### **Multifile IPO Overview**

Multifile IPO obtains potential optimization information from individual program modules of a multifile program. Using the information, the compiler performs optimizations across modules.

Building a program is divided into two phases: compilation and linkage. Multifile IPO performs different work depending on whether the compilation, linkage or both are performed.

### **Compilation Phase**

As each source file is compiled, multifile IPO stores an intermediate representation (IR) of the source code in the object file, which includes summary information used for optimization.

By default, the compiler produces "mock" object files during the compilation phase of multifile IPO. Generating mock files instead of real object files reduces the time spent in the multifile IPO compilation phase. Each mock object file contains the IR for its corresponding source file, but no real code or data. These mock objects must be linked using the `-ipo` option in `ifort` or using the `xild` tool. (See Creating a Multifile IPO Executable with `xild`.)

### **Note**

Failure to link "mock" objects with `ifort` and `-ipo` or `xild` will result in



linkage errors. There are situations where mock object files cannot be used. See [Compilation with Real Object Files](#) for more information.

## Linkage Phase

When you specify `-ipo`, the compiler is invoked a final time before the linker. The compiler performs multifile IPO across all object files that have an IR.

### Note

The compiler does not support multifile IPO for static libraries (`.a` files). See [Compilation with Real Object Files](#) for more information.

`-ipo` enables the driver and compiler to attempt detecting a whole program automatically. If a whole program is detected, the interprocedural constant propagation, stack frame alignment, data layout and padding of common blocks perform more efficiently, while more dead functions get deleted. This option is safe.

## Creating a Multifile IPO Executable with Command Line

Enable multifile IPO for compilations targeted for IA-32 architecture and for compilations targeted for Itanium® architecture as follows in the example below.

Compile your source files with `-ipo` as follows:

Compile source files to produce object files:

```
ifort -ipo -c a.f b.f c.f
```

Produces `a.o`, `b.o`, and `c.o` object files containing Intel compiler intermediate representation (IR) corresponding to the compiled source files `a.f`, `b.f`, and `c.f`. Using `-c` to stop compilation after generating `.o` files is required. You can now optimize interprocedurally.

Link object files to produce application executable:

```
ifort -oipo_file -ipo a.o b.o c.o
```

The `ifort` command performs IPO for objects containing IR and creates a new list of object(s) to be linked. The `ifort` command calls GCC `ld` to link the specified object files and produce `ipo_file` executable specified by the `-o` option. Multifile IPO is applied only to the source files that have an IR, otherwise the object file passes to link stage.

The `-oname` option stores the executable in `ipo_file`. Multifile IPO is applied only to the source files that have an IR, otherwise the object file passes to link stage.

For efficiency, combine steps 1 and 2:

```
ifort -ipo -oipo_file a.f b.f c.f
```

Instead of `ifort`, you can use the `xild` tool.

For a description of how to use multifile IPO with profile information for further optimization, see Example of Profile-Guided Optimization.

### Creating a Multifile IPO Executable Using `xild`

Use the Intel® linker, `xild`, instead of step 2 in Creating a Multifile IPO Executable with Command Line. The Intel linker `xild` performs the following steps:

1. Invokes the Intel compiler to perform multifile IPO if objects containing `IR` are found.
2. Invokes GCC `ld` to link the application.

The command-line syntax for `xild` is the same as that of the GCC linker:

```
prompt>xild [<options>] <LINK_commandline>
```

where:

- [`<options>`] (optional) may include any GCC linker options or options supported only by `xild`.
- `<LINK_commandline>` is your linker command line containing a set of valid arguments to the `ld`.

To place the multifile IPO executable in `ipo_file`, use the option `-ofilename`, for example:

```
prompt>xild -oipo_file a.o b.o c.o
```

`xild` calls Intel compiler to perform IPO for objects containing `IR` and creates a new list of object(s) to be linked. Then `xild` calls `ld` to link the object files that are specified in the new list and produce `ipo_file` executable specified by the `-ofilename` option.

#### **Note**

The `-ipo` option can reorder object files and linker arguments on the command line. Therefore, if your program relies on a precise order of arguments on the command line, `-ipo` can affect the behavior of your program.

## Usage Rules

You must use the Intel linker `xild` to link your application if:

- Your source files were compiled with multifile IPO enabled. Multifile IPO is enabled by specifying the `-ipo` command-line option
- You normally would invoke the GCC linker (`ld`) to link your application.

## The `xild` Options

The additional options supported by `xild` may be used to examine the results of multifile IPO. These options are described in the following table.

<code>-qipo_fa[file.s]</code>	Produces assembly listing for the multifile IPO compilation. You may specify an optional name for the listing file, or a directory (with the backslash) in which to place the file. The default listing name is <code>ipo_out.s</code> .
<code>-qipo_fo[file.o]</code>	Produces object file for the multifile IPO compilation. You may specify an optional name for the object file, or a directory (with the backslash) in which to place the file. The default object file name is <code>ipo_out.o</code> .
<code>-ipo_fcode-asm</code>	Add code bytes to assembly listing
<code>-ipo_fsource-asm</code>	Add high-level source code to assembly listing
<code>-ipo_fverbose-asm,</code> <code>-ipo_fnoverbose-asm</code>	Enable and disable, respectively, inserting comments containing version and options used in the assembly listing for <code>xild</code> .

## Compilation with Real Object Files

In certain situations you might need to generate real object files with `-ipo`. To force the compiler to produce real object files instead of "mock" ones with IPO, you must specify `-ipo_obj` in addition to `-ipo`.

Use of `-ipo_obj` is necessary under the following conditions:

- The objects produced by the compilation phase of `-ipo` will be placed in a static library without the use of `xiar`. The compiler does not support multifile IPO for static libraries, so all static libraries are passed to the linker. Linking with a static library that contains "mock" object files will result in linkage errors because the objects do not contain real code or

data. Specifying

`-ipo_obj` causes the compiler to generate object files that can be used in static libraries.

- Alternatively, if you create the static library using `xiar`, then the resulting static library will work as a normal library.
- The objects produced by the compilation phase of `-ipo` might be linked without the `-ipo` option and without the use of `xiar`.
- You want to generate an assembly listing for each source file (using `-S`) while compiling with `-ipo`. If you use `-ipo` with `-S`, but without `-ipo_obj`, the compiler issues a warning and an empty assembly file is produced for each compiled source file.

## Implementing the `.il` Files with Version Numbers

An IPO compilation consists of two parts: the compile phase and the link phase. In the compile phase, the compiler produces an intermediate language (IL) version of the users' code. In the link phase, the compiler reads the IL and completes the compilation, producing a real object file or executable.

Generally, different compiler versions produce IL based on different definitions, and therefore the ILs from different compilations can be incompatible. Intel Fortran Compiler assigns a unique version number with each compiler's IL definition. If a compiler attempts to read IL in a file with a version number other than its own, the compilation proceeds, but the IL is discarded and not used in the compilation. The compiler then issues a warning message about an incompatible IL detected and discarded.

### IL in Libraries: More Optimizations

The IL produced by the Intel compiler is stored in file with a suffix `.il`. Then the `.il` file is placed in the library. If this library is used in an IPO compilation invoked with the same compiler as produced the IL for the library, the compiler can extract the `.il` file from the library and use it to optimize the program. For example, it is possible to inline functions defined in the libraries into the users' source code.

## Creating a Library from IPO Objects

Normally, libraries are created using a library manager such as `ar`. Given a list of objects, the library manager will insert the objects into a named library to be used in subsequent link steps.

```
xiar cru user.a a.obj b.obj
```

The above command creates a library named `user.a` that contains the `a.o` and `b.o` objects.

If, however, the objects have been created using `-ipo -c`, then the objects will not contain a valid object but only the intermediate representation (IR) for that object file. For example:

```
ifort -ipo -c a.f b.f
```

will produce `a.o` and `b.o` that only contains IR to be used in a link time compilation. The library manager will not allow these to be inserted in a library.

In this case you must use the Intel library driver `xild -ar`. This program will invoke the compiler on the IR saved in the object file and generate a valid object that can be inserted in a library.

```
xild -lib cru user.a a.o b.o
```

See [Creating a Multifile IPO Executable Using `xild`](#).

### Analyzing the Effects of Multifile IPO

The `-ipo_c` and `-ipo_S` options are useful for analyzing the effects of multifile IPO, or when experimenting with multifile IPO between modules that do not make up a complete program.

Use the `-ipo_c` option to optimize across files and produce an object file. This option performs optimizations as described for `-ipo`, but stops prior to the final link stage, leaving an optimized object file. The default name for this file is `ipo_out.o`. You can use the `-o` option to specify a different name. For example:

```
ifort -tpp6 -ipo_c -ofilename a.f b.f c.f
```

Use the `-ipo_S` option to optimize across files and produce an assembly file. This option performs optimizations as described for `-ipo`, but stops prior to the final link stage, leaving an optimized assembly file. The default name for this file is `ipo_out.s`. You can use the `-o` option to specify a different name. For example:

```
ifort -tpp6 -ipo_S -ofilename a.f b.f c.f
```

For more information on inlining and the minimum inlining criteria, see [Criteria for Inline Function Expansion and Controlling Inline Expansion of User Functions](#).

### Using `-ip` with `-Q`option Specifiers

You can adjust the Intel® Fortran Compiler's optimization for a particular application by experimenting with memory and interprocedural optimizations.

Enter the `-Qoption` option with the applicable keywords to select particular inline expansions and loop optimizations. The option must be entered with a `-ip` or `-ipo` specification, as follows:

```
-ip[-Qoption,tool,opts]
```

where `tool` is Fortran (`f`) and `opts` are `-Qoption` specifiers (see below). Also refer to Criteria for Inline Function Expansion to see how these specifiers may affect the inlining heuristics of the compiler.

See Passing Options to Other Tools (`-Qoption,tool,opts`) for details about `-Qoption`.

### **-Qoption Specifiers**

If you specify `-ip` or `-ipo` without any `-Qoption` qualification, the compiler

- expands functions in line
- propagates constant arguments
- passes arguments in registers
- monitors module-level static variables.

You can refine interprocedural optimizations by using the following `-Qoption` specifiers. To have an effect, the `-Qoption` option must be entered with either `-ip` or `-ipo` also specified, as in this example:

```
-ip -Qoption,f,ip_specifier
```

where `ip_specifier` is one of the `-Qoption` specifiers described in the table that follows.

<b>-Qoption Specifiers</b>	
<code>-ip_args_in_regs=0</code>	Disables the passing of arguments in registers. By default, external functions can pass arguments in registers when called locally. Normally, only static functions can pass arguments in registers, provided the address of the function is not taken and the function does not use a variable number of arguments.
<code>-ip_ninl_max_stats=n</code>	Sets the valid number of intermediate language statements for a function that is expanded in line. The number <code>n</code>

	is a positive integer. The number of intermediate language statements usually exceeds the actual number of source language statements. The default value for <i>n</i> is 230.
<code>-ip_ninl_min_stats=n</code>	Sets the valid min number of intermediate language statements for a function that is expanded in line. The number <i>n</i> is a positive integer. The default value for <code>ip_ninl_min_stats</code> is: <b>IA-32 compiler:</b> <code>ip_ninl_min_stats = 7</code> <b>Itanium® compiler:</b> <code>ip_ninl_min_stats = 15</code>
<code>-ip_ninl_max_total_stats=n</code>	Sets the maximum increase in size of a function, measured in intermediate language statements, due to inlining. The number <i>n</i> is a positive integer. The default value for <i>n</i> is 2000.

The following command activates procedural and interprocedural optimizations on `source.f` and sets the maximum increase in the number of intermediate language statements to five for each function:

```
ifort -ip -Qoption,f,-ip_ninl_max_stats=5 source.f
```

## Inline Expansion of Functions

### Criteria for Inline Function Expansion

For a call to be considered for inlining, it has to meet certain minimum criteria. There are three main components of a call:

**Call-site** is the site of the call to the function that might be inlined.

**Caller** is the function that contains the call-site.

**Callee** is the function being called that might be inlined.

### Minimum call-site criteria:

- The number of actual arguments must match the number of formal arguments of the callee.
- The number of return values must match the number of return values of the callee.
- The data types of the actual and formal arguments must be compatible.
- No multilingual inlining is permitted. Caller and callee must be written in the same source language.

#### Minimum criteria for the caller:

- At most 2000 intermediate statements will be inlined into the caller from all the call-sites being inlined into the caller. You can change this value by specifying the option

`-Qoption,f,-ip_ninl_max_total_stats=new value`

- The function must be called if it is declared as static. Otherwise, it will be deleted.

#### Minimum criteria for the callee:

- Does not have variable argument list.
- Is not considered infrequent due to the name. Routines which contain the following substrings in their names are not inlined: `abort`, `alloca`, `denied`, `err`, `exit`, `fail`, `fatal`, `fault`, `halt`, `init`, `interrupt`, `invalid`, `quit`, `rare`, `stop`, `timeout`, `trace`, `trap`, and `warn`.
- Is not considered unsafe for other reasons.

### Selecting Routines for Inlining with or without PGO

Once the above criteria are met, the compiler picks the routines whose inline expansions will provide the greatest benefit to program performance. This is done using the default heuristics. The inlining heuristics used by the compiler differ based on whether you use profile-guided optimizations (`-prof_use`) or not.

**When you use profile-guided optimizations** with `-ip` or `-ipo`, the compiler uses the following heuristics:

- The default heuristic focuses on the most frequently executed call sites, based on the profile information gathered for the program.
- By default, the compiler does not inline functions with more than 230 intermediate statements. You can change this value by specifying the option

`-Qoption,f,-ip_ninl_max_stats=new value.`



- The default inline heuristic will stop inlining when direct recursion is detected.
- The default heuristic always inlines very small functions that meet the minimum inline criteria.
  - Default for Itanium®-based applications: `ip_ninl_min_stats = 15`.
  - Default for IA-32 applications: `ip_ninl_min_stats = 7`.

These limits can be modified with the option:

`-Qoption,f,-ip_ninl_min_stats=new value`.

See `-Qoption` Specifiers and Profile-Guided Optimization (PGO).

**When you do not use profile-guided optimizations** with `-ip` or `-ipo`, the compiler uses less aggressive inlining heuristics: it inlines a function if the inline expansion does not increase the size of the final program.

### Inlining and Preemption

Preemption of a function means that the code, which implements that function at run-time, is replaced by different code. When a function is preempted, the new version of this function is executed rather than the old version. Preemption can be used to replace an erroneous or inferior version of a function with a correct or improved version.

The compiler assumes that when `-ip` is on, any externally visible function might be preempted and therefore cannot be inlined. Currently, this means that all Fortran subprograms, except for internal procedures, are not inlinable when `-ip` is on.

However, if you use `-ipo` and `-ipo_obj` on a file-by-file basis, the functions can be inlined. See [Compilation with Real Object Files](#).

### Controlling Inline Expansion of User Functions

The compiler enables you to control the amount of inline function expansion, with the options shown in the following summary.

Option	Effect
<code>-ip_no_inlining</code>	This option is only useful if <code>-ip</code> or <code>-ipo</code> is also specified. In such case, <code>-ip_no_inlining</code> disables inlining that would result from the <code>-ip</code> interprocedural optimizations, but has no effect on other interprocedural optimizations.

- inline_debug_info	Preserve the source position of inlined code instead of assigning the call-site source position to inlined code.
-Ob{0 1 2}	Controls the compiler's inline expansion. The amount of inline expansion performed varies as follows:  -Ob0: disables inline expansion of user-defined functions; however, statement functions are always inlined.  -Ob1: disables inlining unless -ip or -Ob2 is specified. Enables inlining of functions (routines). This is the default.  -Ob2: enables inlining of any routine at the compiler's direction: the compiler decides which functions are inlined. This option enables interprocedural optimizations and has the same effect as specifying the -ip option.
<b>IA-32 only:</b> -ip_no_pinlining	Disables partial inlining; can be used if -ip or -ipo is also specified.

## Inline Expansion of Library Functions

By default, the compiler automatically expands (inlines) a number of standard and math library functions at the point of the call to that function, which usually results in faster computation.

However, the inlined library functions do not set the `errno` variable when being expanded inline. In code that relies upon the setting of the `errno` variable, you should use the `-nolib_inline` option. Also, if one of your functions has the same name as one of the compiler-supplied library functions, then when this function is called, the compiler assumes that the call is to the library function and replaces the call with an inlined version of the library function.

So, if the program defines a function with the same name as one of the known library routines, you must use the `-nolib_inline` option to ensure that the user-supplied function is used.

`-nolib_inline` disables inlining of all intrinsics.



### Note

Automatic inline expansion of library functions is not related to the inline expansion that the compiler does during interprocedural optimizations. For

example, the following command compiles the program `sum.f` without expanding the math library functions:

```
ifort -ip -nolib_inline sum.f
```

# Profile-guided Optimizations

## Profile-guided Optimizations Overview

Profile-guided optimizations (PGO) tell the compiler which areas of an application are most frequently executed. By knowing these areas, the compiler is able to be more selective and specific in optimizing the application. For example, the use of PGO often enables the compiler to make better decisions about function inlining, thereby increasing the effectiveness of interprocedural optimizations.

### Instrumented Program

Profile-guided Optimization creates an instrumented program from your source code and special code from the compiler. Each time this instrumented code is executed, the instrumented program generates a dynamic information file. When you compile a second time, the dynamic information files are merged into a summary file. Using the profile information in this file, the compiler attempts to optimize the execution of the most heavily travelled paths in the program.

Unlike other optimizations such as those strictly for size or speed, the results of IPO and PGO vary. This is due to each program having a different profile and different opportunities for optimizations. The guidelines provided help you determine if you can benefit by using IPO and PGO. You need to understand the principles of the optimizations and the unique aspects of your source code.

### Added Performance with PGO

In this version of the Intel® Fortran Compiler, PGO is improved in the following ways:

- Register allocation uses the profile information to optimize the location of spill code.
- For indirect function calls, branch prediction is improved by identifying the most likely targets. With the Intel® Pentium® 4 and Intel® Xeon(TM) processors' longer pipeline, improving branch prediction translates into high performance gains.
- The compiler detects and does not vectorize loops that execute only a small number of iterations, reducing the run time overhead that vectorization might otherwise add.

## Profile-guided Optimizations Methodology and Usage Model

PGO works best for code with many frequently executed branches that are difficult to predict at compile time. An example is the code with intensive error-checking in which the error conditions are false most of the time. The "cold" error-handling code can be placed such that the branch is hardly ever

mispredicted. Minimizing "cold" code interleaved into the "hot" code improves instruction cache behavior.

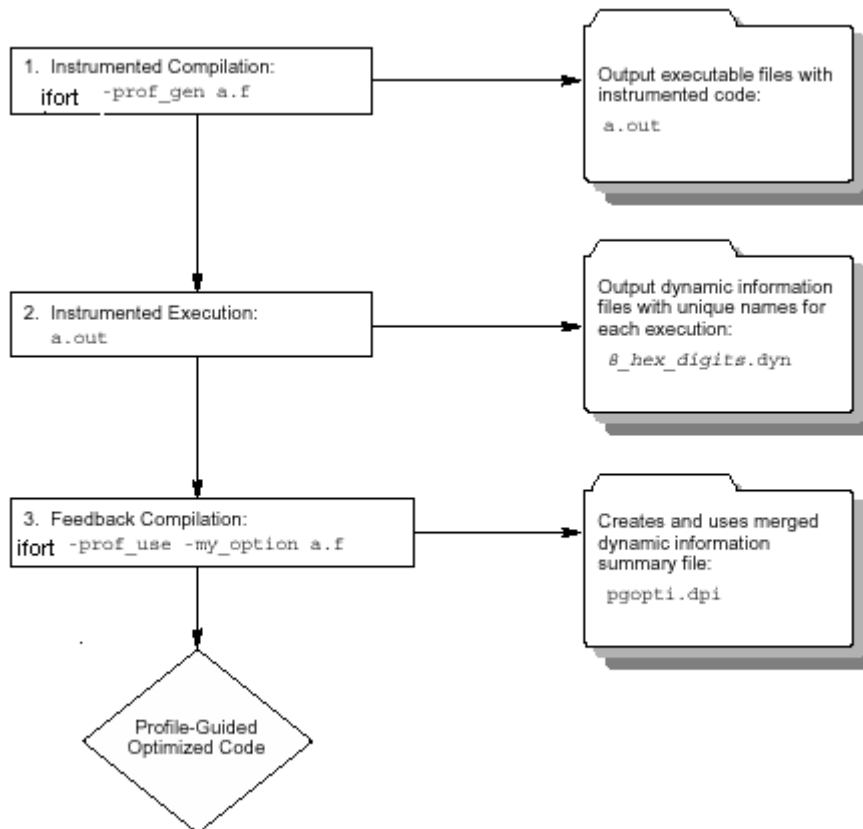
## PGO Phases

The PGO methodology requires three phases and options:

1. **Instrumentation compilation** and linking with `-prof_gen`
2. **Instrumented execution** by running the executable; as a result, the dynamic-information files (`.dyn`) are produced.
3. **Feedback compilation** with `-prof_use`

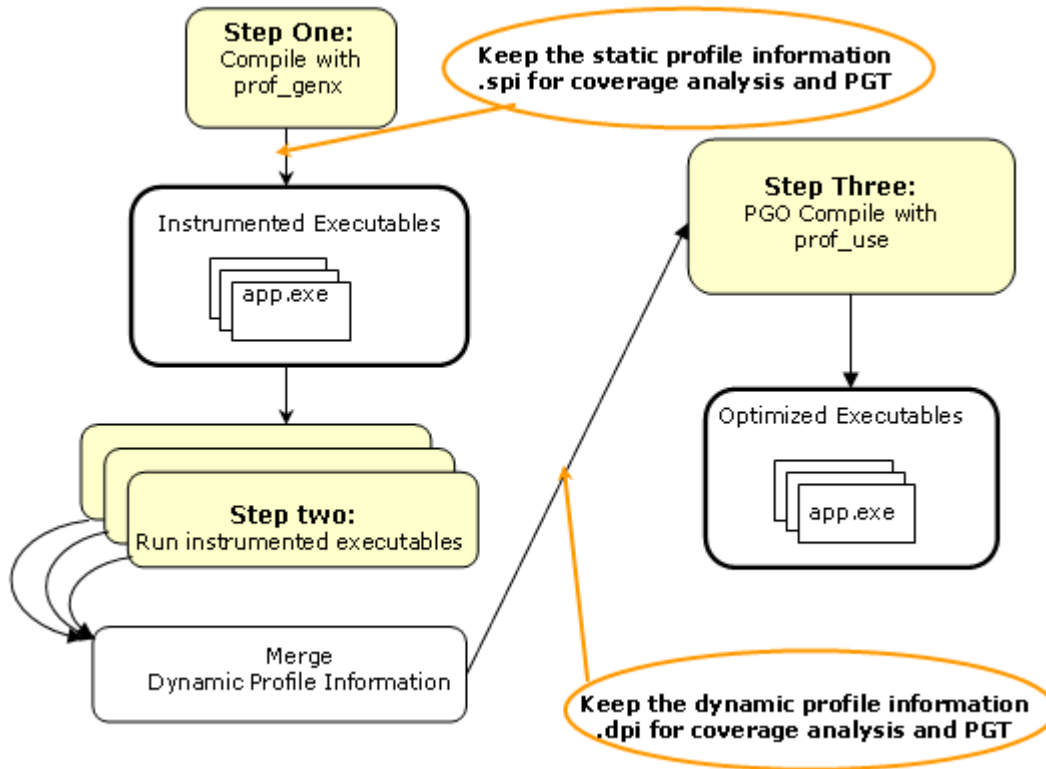
The flowcharts below illustrate this process for IA-32 compilation and Itanium®-based compilation. A key factor in deciding whether you want to use PGO lies in knowing which sections of your code are the most heavily used. If the data set provided to your program is very consistent and it elicits a similar behavior on every execution, then PGO can probably help optimize your program execution. However, different data sets can elicit different algorithms to be called. This can cause the behavior of your program to vary from one execution to the next.

### Phases of Basic Profile-guided Optimization



## PGO Usage Model

The chart that follows presents PGO usage model.



Here are the steps for a simple example (`myApp.f90`) for IA-32 systems.

1. Set

```
PROF_DIR=c:/myApp/prof_dir
```

2. Issue command

```
ifort -prof_genx myApp.f90
```

This command compiles the program and generates instrumented binary `myApp.exe` as well as the corresponding static profile information `pgopti.spi`.

3. Execute `myApp`

Each invocation of `myApp` runs the instrumented application and generates one or more new dynamic profile information files that have an extension `.dyn` in the directory specified by `PROF_DIR`.

#### 4. Issue command

```
ifort -prof_use myApp.f90
```

At this step, the compiler merges all the `.dyn` files into one `.dpi` file representing the total profile information of the application and generates the optimized binary. The default name of the `.dpi` file is `pgopti.dpi`.

### Basic PGO Options

The options used for basic PGO optimizations are:

- `-prof_gen` to generate instrumented code
- `-prof_use` to generate a profile-optimized executable
- `-prof_format_32` to produce 32-bit counters for `.dyn` and `.dpi` files

In cases where your code behavior differs greatly between executions, you have to ensure that the benefit of the profile information is worth the effort required to maintain up-to-date profiles. In the basic profile-guided optimization, the following options are used in the phases of the PGO:

#### Generating Instrumented Code, `-prof_gen`

The `-prof_gen` option instruments the program for profiling to get the execution count of each basic block. It is used in phase 1 of the PGO to instruct the compiler to produce instrumented code in your object files in preparation for instrumented execution. Parallel `make` is automatically supported for `-prof_gen` compilations.

#### Generating a Profile-optimized Executable, `-prof_use`

The `-prof_use` option is used in phase 3 of the PGO to instruct the compiler to produce a profile-optimized executable and merges available dynamic-information (`.dyn`) files into a `pgopti.dpi` file.



#### **Note:**

The dynamic-information files are produced in phase 2 when you run the instrumented executable.

If you perform multiple executions of the instrumented program, `-prof_use` merges the dynamic-information files again and overwrites the previous `pgopti.dpi` file.

#### Using 32-bit Counters, `-prof_format_32`

The Intel Fortran compiler by default produces profile data with 64-bit counters to handle large numbers of events in the `.dyn` and `.dpi` files. The `-prof_format_32` option produces 32-bit counters for compatibility with the earlier compiler versions. If the format of the `.dyn` and `.dpi` files is incompatible with the format used in the current compilation, the compiler issues the following message:

```
Error: xxx.dyn has old or incompatible file format - delete
file and redo instrumentation compilation/execution.
```

### Disabling Function Splitting, `-fnsplit-` (Itanium® Compiler only)

`-fnsplit-` disables function splitting. Function splitting is enabled by `-prof_use` in phase 3 to improve code locality by splitting routines into different sections: one section to contain the cold or very infrequently executed code and one section to contain the rest of the code (hot code).

You can use `-fnsplit-` to disable function splitting for the following reasons:

- Most importantly, to get improved debugging capability. In the debug symbol table, it is difficult to represent a split routine, that is, a routine with some of its code in the hot code section and some of its code in the cold code section.

The `-fnsplit-` option disables the splitting within a routine but enables function grouping, an optimization in which entire routines are placed either in the cold code section or the hot code section. Function grouping does not degrade debugging capability.

- Another reason can arise when the profile data does not represent the actual program behavior, that is, when the routine is actually used frequently rather than infrequently.



#### Note

For Itanium®-based applications, if you intend to use the `-prof_use` option with optimizations at the `-O3` level, the `-O3` option must be on. If you intend to use the `-prof_use` option with optimizations at the `-O2` level or lower, you can generate the profile data with the default options.

See an example of using PGO.

### Advanced PGO Options

The options controlling advanced PGO optimizations are:

- `-prof_dirdirname`



- `-prof_filefilename`.

### Specifying the Directory for Dynamic Information Files

Use the `-prof_dirdirname` option to specify the directory in which you intend to place the dynamic information (`.dyn`) files to be created. The default is the directory where the program is compiled. The specified directory must already exist.

You should specify `-prof_dirdirname` option with the same directory name for both the instrumentation and feedback compilations. If you move the `.dyn` files, you need to specify the new path.

### Specifying Profiling Summary File

The `-prof_filefilename` option specifies file name for profiling summary file.

### Guidelines for Using Advanced PGO

When you use PGO, consider the following guidelines:

- Minimize the changes to your program after instrumented execution and before feedback compilation. During feedback compilation, the compiler ignores dynamic information for functions modified after that information was generated.

#### Note

The compiler issues a warning that the dynamic information does not correspond to a modified function.

- Repeat the instrumentation compilation if you make many changes to your source files after execution and before feedback compilation.
- Specify the name of the profile summary file using the `-prof_filefilename` option

See PGO Environment Variables.

### PGO Environment Variables

The environment variables determine the directory in which to store dynamic information files or whether to overwrite `pgopti.dpi`. The PGO environment variables are described in the table below.

Variable	Description
PROF_DIR	Specifies the directory in which dynamic

	information files are created. This variable applies to all three phases of the profiling process.
PROF_DUMP_INTERVAL	Initiates interval profile dumping in an instrumented user application.
PROF_NO_CLOBBER	Alters the feedback compilation phase slightly. By default, during the feedback compilation phase, the compiler merges the data from all dynamic information files and creates a new <code>pgopti.dpi</code> file, even if one already exists. When this variable is set, the compiler does not overwrite the existing <code>pgopti.dpi</code> file. Instead, the compiler issues a warning and you must remove the <code>pgopti.dpi</code> file if you want to use additional dynamic information files.

See also the documentation for your operating system for instructions on how to specify environment variables and their values.

## Example of Profile-Guided Optimization

The following is an example of the basic PGO phases:

**1. Instrumentation Compilation and Linking**—Use `-prof_gen` to produce an executable with instrumented information. Use also the `-prof_dir` option as recommended for most programs, especially if the application includes the source files located in multiple directories. `-prof_dir` ensures that the profile information is generated in one consistent place. For example:

```
ifort -prof_gen -prof_dir/usr/profdata -c a1.f a2.f
a3.f
```

```
ifort -oa1 a1.o a2.o a3.o
```

In place of the second command, you could use the linker (`ld`) directly to produce the instrumented program. If you do this, make sure you link with the `libirc.a` library.

**2. Instrumented Execution**—Run your instrumented program with a representative set of data to create a dynamic information file.

```
prompt>a1
```

The resulting dynamic information file has a unique name and `.dyn` suffix every time you run `a1`. The instrumented file helps predict how the program runs with a

particular set of data. You can run the program more than once with different input data.

3. **Feedback Compilation**—Compile and link the source files with `-prof_use` to use the dynamic information to optimize your program according to its profile:

```
ifort -prof_use -prof_dir/usr/profdata -ipo a1.f a2.f
a3.f
```

Besides the optimization, the compiler produces a `pgopti.dpi` file. You typically specify the default optimizations (`-O2`) for phase 1, and specify more advanced optimizations (`-ip` or `-ipo`) for phase 3. This example used `-O2` in phase 1 and the `-ipo` in phase 3.

### Note

The compiler ignores the `-ip` or the `-ipo` options with `-prof_gen`.

See Basic PGO Options.

## Merging the .dyn Files

To merge the `.dyn` files, use the `profmerge` utility.

### The `profmerge` Utility

The compiler executes `profmerge` automatically during the feedback compilation phase when you specify `-prof_use`.

The command-line usage for `profmerge` is as follows:

```
profmerge [-nologo] [-prof_dirdirname]
```

where `-prof_dirdirname` is a `profmerge` utility option.

This merges all `.dyn` files in the current directory or the directory specified by `-prof_dir`, and produces the summary file `pgopti.dpi`.

The `-prof_filefilename` option enables you to specify the name of the `.dpi` file.

The command-line usage for `profmerge` with `-prof_filefilename` is as follows:

```
profmerge [-nologo] [-prof_filefilename]
```

where `/prof_filefilename` is a `profmerge` utility option.

 **Note**

The `profmerge` tool merges all the `.dyn` files that exist in the given directory. It is very important to make sure that unrelated `.dyn` files, oftentimes from previous runs, are not present in that directory. Otherwise, profile information will be based on invalid profile data. This can negatively impact the performance of optimized code as well as generate misleading coverage information.

 **Note**

The `.dyn` files can be merged to a `.dpi` file by the `profmerge` tool without recompiling the application.

**Dumping Profile Data**

This subsection provides an example of how to call the C PGO API routines from Fortran. For complete description of the PGO API support routines, see PGO API: Profile Information Generation Support.

As part of the instrumented execution phase of profile-guided optimization, the instrumented program writes profile data to the dynamic information file (`.dyn` file). The file is written after the instrumented program returns normally from `main()` or calls the standard exit function. Programs that do not terminate normally, can use the `_PGOPTI_Prof_Dump` function. During the instrumentation compilation (`-prof_gen`) you can add a call to this function to your program. Here is an example:

```
INTERFACE
SUBROUTINE PGOPTI_PROF_DUMP ( )
!DEC$ ATTRIBUTES C ,
ALIAS: 'PGOPTI_Prof_Dump' :: PGOPTI_PROF_DUMP
END SUBROUTINE
END INTERFACE
CALL PGOPTI_PROF_DUMP ( )
```

 **Note**

You must remove the call or comment it out prior to the feedback compilation with `-prof_use`.

**Using profmerge to Relocate the Source Files**

The compiler uses the full path to the source file for each routine to look up the profile summary information associated with that routine. By default, this prevents you from:

- Using the profile summary file (`.dpi`) if you move your application sources.
- Sharing the profile summary file with another user who is building identical application sources that are located in a different directory.

## Source Relocation

To enable the movement of application sources, as well as the sharing of profile summary files, use the `profmerge` with `-src_old` and `-src_new` options. For example:

```
prompt>profmerge -prof_dir c:/work -src_old c:/work/sources
-src_new d:/project/src
```

The above command will read the `c:/work/pgopti.dpi` file. For each routine represented in the `pgopti.dpi` file, whose source path begins with the `c:/work/sources` prefix, `profmerge` replaces that prefix with `d:/project/src`. The `c:/work/pgopti.dpi` file is updated with the new source path information.



### Notes

- You can execute `profmerge` more than once on a given `pgopti.dpi` file. You may need to do this if the source files are located in multiple directories. For example:

```
profmerge -src_old "c:/program files" -src_new
"e:/program files"
```

```
profmerge -src_old c:/proj/application -src_new
d:/app
```

- In the values specified for `-src_old` and `-src_new`, uppercase and lowercase characters are treated as identical. Likewise, forward slash (/) and backward slash (\) characters are treated as identical.
- Because the source relocation feature of `profmerge` modifies the `pgopti.dpi` file, you may wish to make a backup copy of the file prior to performing the source relocation.

## Code-coverage Tool

The Intel® Compilers Code-coverage tool can be used for both IA-32 and Itanium® architectures, in a number of ways to improve development efficiency, reduce defects, and increase application performance. The major features of the Intel Compilers code-coverage tool are:

- Visual presentation of the application's code coverage information with the code-coverage coloring scheme
- Display of the dynamic execution counts of each basic block of the application
- Differential coverage, or comparison of the profiles of the application's two runs

### Command-line Syntax

The syntax for this tool is as follows:

```
codecov [-codecov_option]
```

where `-codecov_option` is a tool option you choose to run the code coverage with. If you do not use any option, the tool will provide the top level code coverage for your whole program.

### Tool Options

The tool uses options that are listed in the table that follows.

Option	Description	Default
<code>-help</code>	Prints all the options of the code-coverage tool.	
<code>-spi file</code>	Sets the path name of the static profile information file <code>.spi</code> .	<code>pgopti.spi</code>
<code>-dpi file</code>	Sets the path name of the dynamic profile information file <code>.dpi</code> .	<code>pgopti.dpi</code>
<code>-prj</code>	Sets the project name.	
<code>-counts</code>	Generates dynamic execution counts.	
<code>-nopartial</code>	Treats partially covered code as fully covered code.	
<code>-comp</code>	Sets the filename that contains the list of files of interest.	
<code>-ref</code>	Finds the differential coverage with respect to <code>ref_dpi_file</code> .	
<code>-demang</code>	Demangles both function names and their arguments.	
<code>-mname</code>	Sets the name of the web-page owner.	

-maddr	Sets the email address of the web-page owner.	
-bcolor	Sets the html color name or code of the uncovered blocks.	#ffff99
-fcolor	Sets the html color name or code of the uncovered functions.	#ffcccc
-pcolor	Sets the html color name or code of the partially covered code.	#fafad2
-ccolor	Sets the html color name or code of the covered code.	#ffffff
-ucolor	Sets the html color name or code of the unknown code.	#ffffff

## Visual Presentation of the Application's Code Coverage

Based on the profile information collected from running the instrumented binaries when testing an application, Intel® Compiler creates HTML files using a code-coverage tool. These HTML files indicate portions of the source code that were or were not exercised by the tests. When applied to the profile of the performance workloads, the code-coverage information shows how well the training workload covers the application's critical code. High coverage of performance-critical modules is essential to taking full advantage of the profile-guided optimizations.

The code-coverage tool can create two levels of coverage:

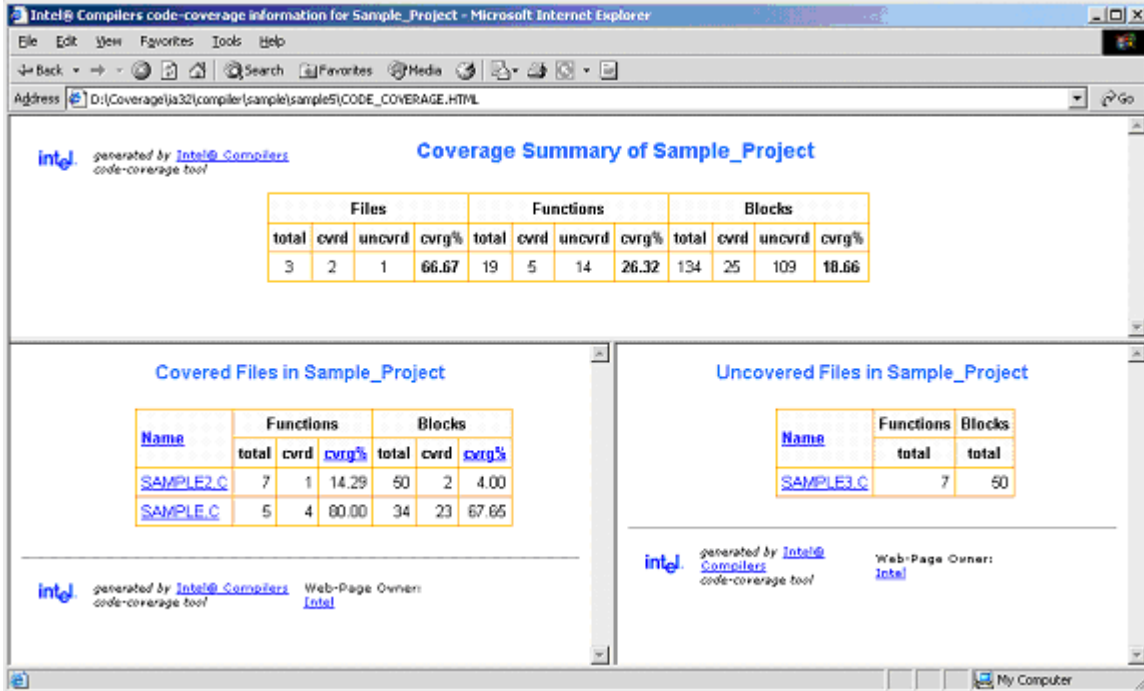
- Top level: for a group of selected modules
- Individual module source view

### Top Level Coverage

The top-level coverage reports the overall code coverage of the modules that were selected. The following options are provided:

- You can select the modules of interest
- For the selected modules, the tool generates a list with their coverage information. The information includes the total number of functions and blocks in a module and the portions that were covered.
- By clicking on the title of columns in the reported tables, the lists may be sorted in ascending or descending order based on:
  - basic block coverage
  - function coverage
  - function name.

The screenshot that follows shows a sample top-level coverage summary for a project. By clicking on a module name (for example, `SAMPLE.C`), the browser will display the coverage source view of that particular module.



## Browsing the Frames

The coverage tool creates frames that facilitate browsing through the code to identify uncovered code. The top frame displays the list of uncovered functions while the bottom frame displays the list of covered functions. For uncovered functions, the total number of basic blocks of each function is also displayed. For covered functions, both the total number of blocks and the number of covered blocks as well as their ratio (that is, the coverage rate) are displayed.

For example, 66.67(4/6) indicates that four out of the six blocks of the corresponding function were covered. The block coverage rate of that function is thus 66.67%. These lists can be sorted based on the coverage rate, number of blocks, or function names. Function names are linked to the position in source view where the function body starts. So, just by one click, the user can see the least-covered function in the list and by another click the browser displays the body of the function. The user can then scroll down in the source view and browse through the function body.

## Individual Module Source View

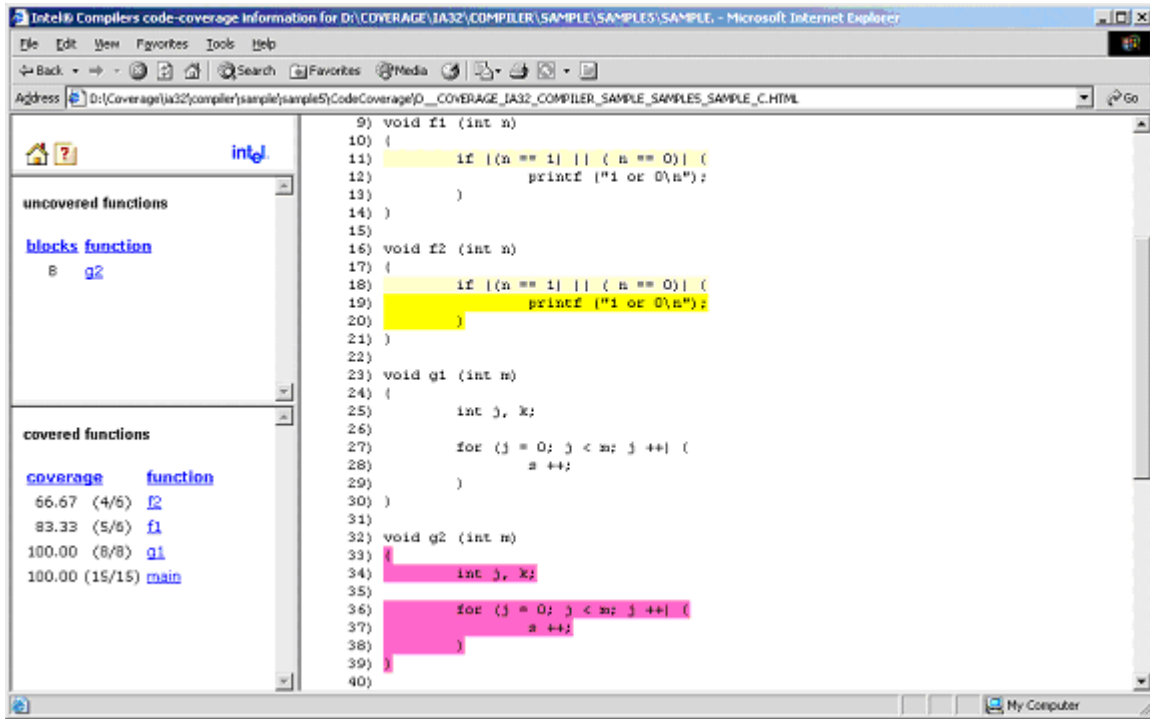
Within the individual module source views, the tool provides the list of uncovered functions as well as the list of covered functions. The lists are reported in two distinct frames that provide easy navigation of the source code. The lists can be sorted based on:

- the number of blocks within uncovered functions
- the block coverage in the case of covered functions



- the function names.

The following screen shows the coverage source view of SAMPLE.C.



## Setting the Coloring Scheme for the Code Coverage

The tool provides a visible coloring distinction of the following coverage categories:

- covered code
- uncovered basic blocks
- uncovered functions
- partially covered code
- unknown.

The default colors that the tool uses for presenting the coverage information are shown in the tables that follows.

This color	Means
Covered code	The portion of code colored in this color was exercised by the tests. The default color can be overridden with the <code>-ccolor</code> option.
Uncovered basic block	Basic blocks that are colored in this color were not exercised by any of the tests.  They were, however, within functions that were executed

	during the tests.  The default color can be overridden with the <code>-bcolor</code> option.
Uncovered function	Functions that are colored in this color were never called during the tests. The default color can be overridden with the <code>-fcolor</code> option.
Partially covered code	More than one basic block was generated for the code at this position.  Some of the blocks were covered while some were not. The default color can be overridden with the <code>-pcolor</code> option.
Unknown	No code was generated for this source line. Most probably, the source at this position is a comment, a header-file inclusion, or a variable declaration. The default color can be overridden with the <code>-ucolor</code> option.

The default colors can be customized to be any valid HTML by using the options mentioned for each coverage category in the table above.

For code-coverage colored presentation, the coverage tool uses the following heuristic. Source characters are scanned until reaching a position in the source that is indicated by the profile information as the beginning of a basic block. If the profile information for that basic block indicates that a coverage category changes, then the tool changes the color corresponding to the coverage condition of that portion of the code, and the coverage tool inserts the appropriate color change in the HTML files.

### Note

You need to interpret the colors in the context of the code. For instance, comment lines that follow a basic block that was never executed would be colored in the same color as the uncovered blocks. Another example is the closing brackets in C/C++ applications.

## Coverage Analysis of a Modules Subset

One of the capabilities of the Intel Compilers code-coverage tool is efficient coverage analysis of an application's subset of modules. This analysis is accomplished based on the selected option `-comp` of the tool's execution.

You can generate the profile information for the whole application, or a subset of it, and then break the covered modules into different components and use the coverage tool to obtain the coverage information of each individual component. If only a subset of the application modules is compiled with the `-prof_genx` option, then the coverage information is generated only for those modules that are involved with this compiler option, thus avoiding the overhead incurred for profile generation of other modules.

To specify the modules of interest, use the tool's `-comp` option. This option takes the name of a file as its argument. That file must be a text file that includes the name of modules or directories you would like to analyze. Here is an example:

```
codecov -prj Project_Name -comp component1
```

### Note

Each line of component file should include one, and only one, module name.

Any module of the application whose full path name has an occurrence of any of the names in the component file will be selected for coverage analysis. For example, if a line of file `component1` in the above example contains `mod1.f90`, then all modules in the application that have such a name will be selected. The user can specify a particular module by giving more specific path information. For instance, if the line contains `/cmp1/mod1.f90`, then only those modules with the name `mod1.c` will be selected that are in a directory named `cmp1`. If no component file is specified, then all files that have been compiled with `-prof_genx` are selected for coverage analysis.

## Dynamic Counters

This feature displays the dynamic execution count of each basic block of the application, and as such it is useful for both coverage and performance tuning.

The coverage tool can be configured to generate the information about the dynamic execution counts. This configuration requires using the `-counts` option. The counts information is displayed under the code after a `^` sign precisely under the source position where the corresponding basic block begins. If more than one basic block is generated for the code at a source position (for example, for macros), then the total number of such blocks and the number of the blocks that were executed are also displayed in front of the execution count.

For example, line 11 in the code is an `IF` statement:

```
11  IF ((N .EQ. 1).OR. (N .EQ. 0))
^ 10 (1/2)
12  PRINT N
  ^ 7
```

The coverage lines under code lines 11 and 12 contain the following information:

- The `IF` statement in line 11 was executed 10 times.
- Two basic blocks were generated for the `IF` statement in line 11.
- Only one of the two blocks was executed, hence the partial coverage color.

- Only seven out of the ten times variable `n` had a value of 0 or 1.

In certain situations, it may be desirable to consider all the blocks generated for a single source position as one entity. In such cases, it is necessary to assume that all blocks generated for one source position are covered when at least one of the blocks is covered. This assumption can be configured with the `-nopartial` option. When this option is specified, decision coverage is disabled, and the related statistics are adjusted accordingly. The code lines 11 and 12 indicate that the `PRINT` statement in line 12 was covered. However, only one of the conditions in line 11 was ever true. With the `-nopartial` option, the tool treats the partially covered code (like the code on line 11) as covered.

## Differential Coverage

Using the code-coverage tool, you can compare the profiles of the application's two runs: a reference run and a new run identifying the code that is covered by the new run but not covered by the reference run. This feature can be used to find the portion of the application's code that is not covered by the application's tests but is executed when the application is run by a customer. It can also be used to find the incremental coverage impact of newly added tests to an application's test space.

The dynamic profile information of the reference run for differential coverage is specified by the `-ref` option, such as in the following command:

```
codecov -prj Project_Name -dpi customer.dpi -ref  
appTests.dpi
```

The coverage statistics of a differential-coverage run shows the percentage of the code that was exercised on a new run but was missed in the reference run. In such cases, the coverage tool shows only the modules that included the code that was uncovered.

The coloring scheme in the source views also should be interpreted accordingly. The code that has the same coverage property (covered or not covered) on both runs is considered as covered code. Otherwise, if the new run indicates that the code was executed while in the reference run the code was not executed, then the code is treated as uncovered. On the other hand, if the code is covered in the reference run but not covered in the new run, the differential-coverage source view shows the code as covered.

## Running for Differential Coverage

### Files Required

To run the Intel Compilers code-coverage tool for differential coverage, the following files are required:

- The application sources
- The `.spi` file generated by Intel Compilers when compiling the application for the instrumented binaries with the `-prof_genx` option.
- The `.dpi` file generated by Intel Compilers `profmerge` utility as the result of merging the dynamic profile information `.dyn` files or the `.dpi` file generated implicitly by Intel Compilers when compiling the application with the `-prof_use` option.

See Usage Model of the Profile-guided Optimizations.

## Running

Once the required files are available, the coverage tool may be launched from this command line:

```
codecov -prj Project_Name -spi pgopti.spi -dpi pgopti.dpi
```

The `-spi` and `-dpi` options specify the paths to the corresponding files.

The coverage tool also has the following additional options for generating a link at the bottom of each HTML page to send an electronic message to a named contact by using `-mname` and `-maddr` options.

```
codecov -prj Project_Name -mname John_Smith -maddr  
js@company.com
```

## Test Prioritization Tool

The Intel® Compilers Test-prioritization tool enables the profile-guided optimizations to select and prioritize application's tests based on prior execution profiles of the application. The tool offers a potential of significant time saving in testing and developing large-scale applications where testing is the major bottleneck. The tool can be used for both IA-32 and Itanium® architectures.

This tool enables the users to select and prioritize the tests that are most relevant for any subset of the application's code. When certain modules of an application are changed, the test-prioritization tool suggests the tests that are most probably affected by the change. The tool analyzes the profile data from previous runs of the application, discovers the dependency between the application's components and its tests, and uses this information to guide the process of testing.

### Features and Benefits

The tool provides an effective testing hierarchy based on the application's code coverage. The advantages of the tool usage can be summarized as follows:

- Minimizing the number of tests that are required to achieve a given overall coverage for any subset of the application: the tool defines the smallest subset of the application tests that achieve exactly the same code coverage as the entire set of tests.
- Reducing the turn-around time of testing: instead of spending a long time on finding a possibly large number of failures, the tool enables the users to quickly find a small number of tests that expose the defects associated with the regressions caused by a change set.
- Selecting and prioritizing the tests to achieve certain level of code coverage in a minimal time based on the data of the tests' execution time.

## Command-line Syntax

The syntax for this tool is as follows:

```
tselect -dpi_list file
```

where `-dpi_list` is a required tool option that sets the path to the DPI list *file* that contains the list of the `.dpi` files of the tests you need to prioritize.

## Tool Options

The tool uses options that are listed in the table that follows.

Option	Description	Default
<code>-help</code>	Prints all the options of the test-prioritization tool.	
<code>-spi file</code>	Sets the path name of the static profile information file <code>.spi</code> .	<code>pgopti.spi</code>
<code>-dpi_list file</code>	Sets the path name of the file that contains the name of the dynamic profile information ( <code>.dpi</code> ) files. Each line of the file should contain one <code>.dpi</code> name optionally followed by its execution time. The name must uniquely identify the test.	
<code>-prof_dpi file</code>	Sets the path name of the output report file.	
<code>-comp</code>	Sets the filename that contains the list of files of interest.	
<code>-cutoff value</code>	Terminates when the cumulative block coverage reaches <i>value</i> % of pre-computed total coverage. <i>value</i> must be greater than 0.0 (for example, 99.00). It may be set to 100.	
<code>-nototal</code>	Does not pre-compute the total coverage.	
<code>-mintime</code>	Minimizes testing execution time. The	

	execution time of each test must be provided on the same line of <code>dpi_list</code> file after the test name in <code>dd:hh:mm:ss</code> format.	
<code>-verbose</code>	Generates more logging information about the program progress.	

## Usage Requirements

To run the test-prioritization tool on an application's tests, the following files are required:

- The `.spi` file generated by Intel Compilers when compiling the application for the instrumented binaries with the `-prof_genx` option.
- The `.dpi` files generated by Intel Compilers `profmerge` tool as a result of merging the dynamic profile information `.dyn` files of each of the application tests. The user needs to apply the `profmerge` tool to all `.dyn` files that are generated for each individual test and name the resulting `.dpi` in a fashion that uniquely identifies the test. The `profmerge` tool merges all the `.dyn` files that exist in the given directory.

### Note

It is very important that the user makes sure that unrelated `.dyn` files, oftentimes from previous runs or from other tests, are not present in that directory. Otherwise, profile information will be based on invalid profile data. This can negatively impact the performance of optimized code as well as generate misleading coverage information.

- User-generated file containing the list of tests to be prioritized.

### Note

For successful tool execution, you should:

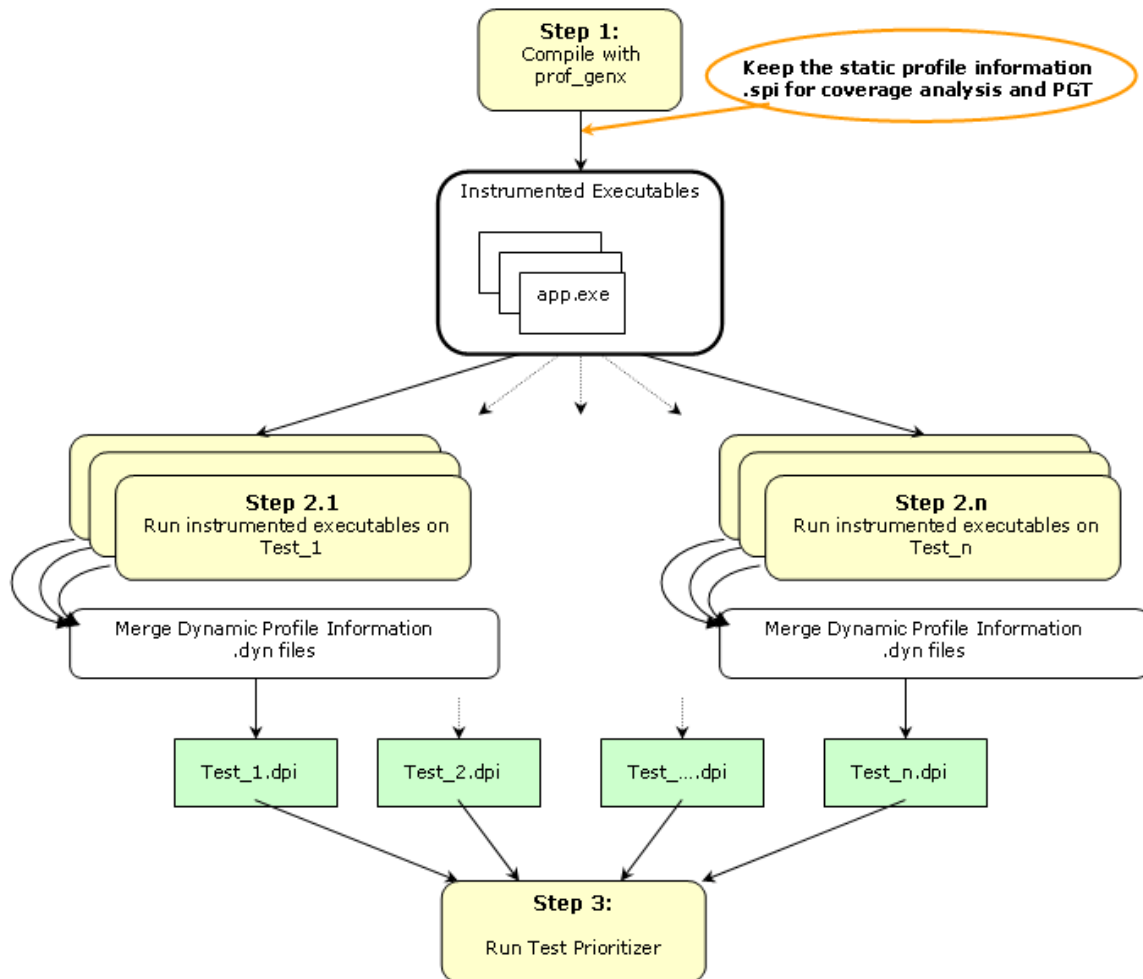
- § Name each test `.dpi` file so that the file names uniquely identify each test.
- § Create a DPI list file: a text file that contains the names of all `.dpi` test files. The name of this file serves as an input for the test-prioritization tool execution command. Each line of the DPI list file should include one, and only one, `.dpi` file name. The name can optionally be followed by the duration of the execution time for a corresponding test in the `dd:hh:mm:ss` format.

For example: `Test1.dpi 00:00:60:35` informs that Test1 lasted 0 days, 0 hours, 60 minutes and 35 seconds.

The execution time is optional. However, if it is not provided, then the tool will not prioritize the test for minimizing execution time. It will prioritize to minimize the number of tests only.

## Usage Model

The chart that follows presents the test-prioritization tool usage model.



Here are the steps for a simple example (`myApp.f90`) for IA-32 systems.

1. Set

```
PROF_DIR=c:/myApp/prof_dir
```

2. Issue command

```
ifort -prof_genx myApp.f90
```



This command compiles the program and generates instrumented binary `myApp` as well as the corresponding static profile information `pgopti.spi`.

3. Issue command

```
rm PROF_DIR /*.dyn
```

Make sure that there are no unrelated `.dyn` files present.

4. Issue command

```
myApp < data1
```

Invocation of this command runs the instrumented application and generates one or more new dynamic profile information files that have an extension `.dyn` in the directory specified by `PROF_DIR`.

5. Issue command

```
profmerge -prof_dpi Test1.dpi
```

At this step, the `profmerge` tool merges all the `.dyn` files into one file (`Test1.dpi`) that represents the total profile information of the application on `Test1`.

6. Issue command

```
rm PROF_DIR /*.dyn
```

Make sure that there are no unrelated `.dyn` files present.

7. Issue command

```
myApp < data2
```

This command runs the instrumented application and generates one or more new dynamic profile information files that have an extension `.dyn` in the directory specified by `PROF_DIR`.

8. Issue command

```
profmerge -prof_dpi Test2.dpi
```

At this step, the `profmerge` tool merges all the `.dyn` files into one file (`Test2.dpi`) that represents the total profile information of the application on `Test2`.

## 9. Issue command

```
rm PROF_DIR /*.dyn
```

Make sure that there are no unrelated `.dyn` files present.

## 10. Issue command

```
myApp < data3
```

This command runs the instrumented application and generates one or more new dynamic profile information files that have an extension `.dyn` in the directory specified by `PROF_DIR`.

## 11. Issue Command

```
profmerge -prof_dpi Test3.dpi
```

At this step, the `profmerge` tool merges all the `.dyn` files into one file (`Test3.dpi`) that represents the total profile information of the application on `Test3`.

12. Create a file named `tests_list` with three lines. The first line contains `Test1.dpi`, the second line contains `Test2.dpi`, and the third line contains `Test3.dpi`.

When these items are available, the test-prioritization tool may be launched from the command line in `PROF_DIR` directory as described in the following examples. *In all examples, the discussion references the same set of data.*

**Example 1 Minimizing the Number of Tests**

```
tselect -dpi_list tests_list -spi pgopti.spi
```

where the `/spi` option specifies the path to the `.spi` file.

Here is a sample output from this run of the test-prioritization tool.

```
Total number of tests    = 3
Total block coverage     ~ 52.17
Total function coverage  ~ 50.00
```

Num	%RatCvrg	%BlkCvrg	%FncCvrg	Test Name @ Options
1	87.50	45.65	37.50	Test3.dpi
2	100.00	52.17	50.00	Test2.dpi

In this example, the test-prioritization tool has provided the following information:

- By running all three tests, we achieve 52.17% block coverage and 50.00% function coverage.
- Test3 by itself covers 45.65% of the basic blocks of the application, which is 87.50% of the total block coverage that can be achieved from all three tests.
- By adding Test2, we achieve a cumulative block coverage of 52.17% or 100% of the total block coverage of Test1, Test2, and Test3.
- Elimination of Test1 has no negative impact on the total block coverage.

## Example 2 Minimizing Execution Time

Suppose we have the following execution time of each test in the `tests_list` file.

```
Test1.dpi 00:00:60:35
```

```
Test2.dpi 00:00:10:15
```

```
Test3.dpi 00:00:30:45
```

The following command executes the test-prioritization tool to minimize the execution time with the `-mintime` option:

```
tselect -dpi_list tests_list -spi pgopti.spi -mintime
```

Here is a sample output.

```
Total number of tests   = 3
Total block coverage    ~ 52.17
Total function coverage ~ 50.00
Total execution time    = 1:41:35
```

num	elapsedTime	%RatCvrg	%BlkCvrg	%FncCvrg	Test Name @ Options
1	10:15	75.00	39.13	25.00	Test2.dpi
2	41:00	100.00	52.17	50.00	Test3.dpi

In this case, the results indicate that the running all tests sequentially would require one hour, 45 minutes, and 35 seconds, while the selected tests would achieve the same total block coverage in only 41 minutes.

## Note

The order of tests when prioritization is based on minimizing time (first Test2, then Test3) could be different than when prioritization is done based on minimizing the number of tests. See example above: first Test3, then Test2. In Example 2, Test2 is the test that gives the highest coverage per execution time. So, it is picked as the first test to run.

## Using Other Options

The `-cutoff` option enables the test-prioritization tool to exit when it reaches a given level of basic block coverage.

```
tselect -dpi_list tests_list -spi pgopti.spi -cutoff 85.00
```

If the tool is run with the cutoff value of 85.00 in the above example, only Test3 will be selected, as it achieves 45.65% block coverage, which corresponds to 87.50% of the total block coverage that is reached from all three tests.

The test-prioritization tool does an initial merging of all the profile information to figure out the total coverage that is obtained by running all the tests. The `-nototal` option enables you to skip this step. In such a case, only the absolute coverage information will be reported, as the overall coverage remains unknown.

# PGO API: Profile Information Generation Support

## PGO API Support Overview

The Profile Information Generation Support (Profile IGS) enables you to control the generation of profile information during the instrumented execution phase of profile-guided optimizations.

Normally, profile information is generated by an instrumented application when it terminates by calling the standard `exit()` function.

To ensure that profile information is generated, the functions described in this section may be necessary or useful in the following situations:

- The instrumented application exits using a non-standard exit routine.
- The instrumented application is a non-terminating application: `exit()` is never called.
- The application requires control of when the profile information is generated.

A set of functions and an environment variable comprise the Profile IGS.

## The Profile IGS Functions

The Profile IGS functions are available to your application by inserting a header file at the top of any source file where the functions may be used.

```
#include "pgouser.h"
```



### Note

The Profile IGS functions are written in C language. Fortran applications need to call C functions.

The rest of the topics in this section describe the Profile IGS functions.



### Note

Without instrumentation, the Profile IGS functions cannot provide PGO API support.

## The Profile IGS Environment Variable

The environment variable for Profile IGS is `PROF_DUMP_INTERVAL`. This environment variable may be used to initiate Interval Profile Dumping in an instrumented user application. See the recommended usage of `_PGOPTI_Set_Interval_Prof_Dump()` for more information.

## Dumping Profile Information

The `_PGOPTI_Prof_Dump()` function dumps the profile information collected by the instrumented application and has the following prototype:

```
void _PGOPTI_Prof_Dump(void);
```

The profile information is generated in a `.dyn` file (generated in [phase 2](#) of the PGO).

## Recommended usage

Insert a single call to this function in the body of the function which terminates the user application. Normally, `_PGOPTI_Prof_Dump()` should be called just once.

It is also possible to use this function in conjunction with the `_PGOPTI_Prof_Reset()` function to generate multiple `.dyn` files (presumably from multiple sets of input data).

Example
! selectively collect profile information

```
! for the portion of the application
! involved in processing input data

input_data = get_input_data()
do while (input_data)
  call _PGOPTI_Prof_Reset()
  call process_data(input_data)
  call _PGOPTI_Prof_Dump();
  input_data = get_input_data();
end do
```

## Resetting the Dynamic Profile Counters

The `_PGOPTI_Prof_Reset()` function resets the dynamic profile counters and has the following prototype:

```
void _PGOPTI_Prof_Reset(void);
```

### Recommended usage

Use this function to clear the profile counters prior to collecting profile information on a section of the instrumented application. See the example under [\\_PGOPTI\\_Prof\\_Dump\(\)](#).

## Dumping and Resetting Profile Information

The `_PGOPTI_Prof_Dump_And_Reset()` function dumps the profile information to a new `.dyn` file and then resets the dynamic profile counters. Then the execution of the instrumented application continues. The prototype of this function is:

```
void _PGOPTI_Prof_Dump_And_Reset(void);
```

This function is used in non-terminating applications and may be called more than once.

### Recommended usage

Periodic calls to this function enables a non-terminating application to generate one or more profile information files (`.dyn` files). These files are merged during the feedback phase ([phase 3](#)) of profile-guided optimizations. The direct use of this function enables your application to control precisely when the profile information is generated.

## Interval Profile Dumping

The `_PGOPTI_Set_Interval_Prof_Dump()` function activates Interval Profile Dumping and sets the approximate frequency at which dumps occur. The prototype of the function call is:

```
void _PGOPTI_Set_Interval_Prof_Dump(int interval);
```

This function is used in non-terminating applications.

The *interval* parameter specifies the time interval at which profile dumping occurs and is measured in milliseconds. For example, if interval is set to 5000, then a profile dump and reset will occur approximately every 5 seconds. The interval is approximate because the time-check controlling the dump and reset is only performed upon entry to any instrumented function in your application.



### Notes

1. Setting interval to zero or a negative number will disable interval profile dumping.
2. Setting a very small value for interval may cause the instrumented application to spend nearly all of its time dumping profile information. Be sure to set interval to a large enough value so that the application can perform actual work and substantial profile information is collected.

### Recommended usage

This function may be called at the start of a non-terminating user application, to initiate Interval Profile Dumping. Note that an alternative method of initiating Interval Profile Dumping is by setting the environment variable, `PROF_DUMP_INTERVAL`, to the desired interval value prior to starting the application.

The intention of Interval Profile Dumping is to allow a non-terminating application to be profiled with minimal changes to the application source code.

## High-level Language Optimizations (HLO)

### HLO Overview

High-level optimizations exploit the properties of source code constructs (for example, loops and arrays) in the applications developed in high-level programming languages, such as Fortran and C++. The high-level optimizations include loop interchange, loop fusion, loop unrolling, loop distribution, unroll-and-jam, blocking, data prefetch, scalar replacement, data layout optimizations and loop unrolling techniques.

The option that turns on the high-level optimizations is `-O3`. The scope of optimizations turned on by `-O3` is different for IA-32 and Itanium®-based applications. See Setting Optimization Levels.

### IA-32 and Itanium®-based Applications

The `-O3` option enables `-O2` option and adds more aggressive optimizations; for example, loop transformation and prefetching. `-O3` optimizes for maximum speed, but may not improve performance for some programs.

### IA-32 Applications

In conjunction with the vectorization options, `-ax{K|W|N|B|P}` and `-x{K|W|N|B|P}`, the `-O3` option causes the compiler to perform more aggressive data dependency analysis than for default `-O2`. This may result in longer compilation times.

### Itanium-based Applications

The `-ivdep_parallel` option asserts there is no loop-carried dependency in the loop where `IVDEP` directive is specified. This is useful for sparse matrix applications.

### Loop Transformations

The loop transformation techniques include:

- loop normalization
- loop reversal
- loop interchange and permutation
- loop skewing
- loop distribution
- loop fusion
- scalar replacement



The loop transformations listed above are supported by data dependence. The loop transformation techniques also include:

- induction variable elimination
- constant propagation
- copy propagation
- forward substitution
- and dead code elimination.

In addition to the loop transformations listed for both IA-32 and Itanium® architectures above, the Itanium architecture enables implementation of the collapsing techniques.

### Scalar Replacement (IA-32 Only)

The goal of scalar replacement is to reduce memory references. This is done mainly by replacing array references with register references.

While the compiler replaces some array references with register references when `-O1` or `-O2` is specified, more aggressive replacement is performed when `-O3` (`-scalar_rep`) is specified. For example, with `-O3` the compiler attempts replacement when there are loop-carried dependences or when data-dependence analysis is required for memory disambiguation.

<code>-scalar_rep[-]</code>	Enables (default) or disables scalar replacement performed during loop transformations (requires <code>-O3</code> ).
-----------------------------	--

### Loop Unrolling with `-unroll[n]`

The `-unroll[n]` option is used in the following way:

- `-unrolln` specifies the maximum number of times you want to unroll a loop. The following example unrolls a loop at most four times:

```
ifort -unroll4 a.f
```

To disable loop unrolling, specify `n` as 0. The following example disables loop unrolling:

```
ifort -unroll0 a.f
```

- `-unroll` (`n` omitted) lets the compiler decide whether to perform unrolling or not.

- `-unroll0` ( $n = 0$ ) disables unroller.

Itanium® compiler currently uses only  $n = 0$ ; any other value is NOP.

## Benefits and Limitations of Loop Unrolling

The benefits are:

- Unrolling eliminates branches and some of the code.
- Unrolling enables you to aggressively schedule (or pipeline) the loop to hide latencies if you have enough free registers to keep variables live.
- The Intel® Pentium® 4 or Intel® Xeon(TM) processors can correctly predict the exit branch for an inner loop that has 16 or fewer iterations, if that number of iterations is predictable and there are no conditional branches in the loop. Therefore, if the loop body size is not excessive, and the probable number of iterations is known, unroll inner loops for:
  - Pentium 4 or Intel Xeon processor, until they have a maximum of 16 iterations
  - Pentium III or Pentium II processors, until they have a maximum of 4 iterations

The potential costs are:

- Excessive unrolling, or unrolling of very large loops can lead to increased code size.
- If the number of iterations of the unrolled loop is 16 or less, the branch predictor should be able to correctly predict branches in the loop body that alternate direction.

For more information on how to optimize with `-unroll[n]`, refer to *Intel® Pentium® 4 and Intel® Xeon(TM) Processor Optimization Reference Manual*.

## Absence of Loop-carried Memory Dependency with IVDEP Directive

For Itanium®-based applications, the `-ivdep_parallel` option indicates there is absolutely no loop-carried memory dependency in the loop where `IVDEP` directive is specified. This technique is useful for some sparse matrix applications.

For example, the following loop requires `-ivdep_parallel` in addition to the directive `IVDEP` to indicate there is no loop-carried dependencies.

```
!DIR$ IVDEP
do i=1,n
  e(ix(2,i))=e(ix(2,i))+1.0
```

```
e(ix(3,i))=e(ix(3,i))+2.0
enddo
```

The following example shows that using this option and the `IVDEP` directive ensures there is no loop-carried dependency for the store into `a()`.

```
!DIR$IVDEP
do j=1,n
a(b(j)) = a(b(j))+1
enddo
```

See `IVDEP` directive for Vectorization Support.

## Prefetching

The goal of `-prefetch` insertion is to reduce cache misses by providing hints to the processor about when data should be loaded into the cache. The prefetching optimizations implement the following options:

<code>-prefetch[-]</code>	Enable or disable ( <code>-prefetch-</code> ) prefetch insertion. This option requires that <code>-O3</code> be specified. The default with <code>-O3</code> is <code>-prefetch</code> .
---------------------------	--

To facilitate compiler optimization:

- Minimize use of global variables and pointers.
- Minimize use of complex control flow.
- Choose data types carefully and avoid type casting.

For more information on how to optimize with `-prefetch[-]`, refer to *Intel® Pentium® 4 and Intel® Xeon(TM) Processor Optimization Reference Manual*.

In addition to the `-prefetch` option, an intrinsic subroutine, `MM_PREFETCH`, is also available. This intrinsic subroutine prefetches data from the specified address on one memory cache line. For details, refer to the *Intel® Fortran Language Reference*.

# Parallel Programming with Intel® Fortran

## Parallelism: an Overview

This section discusses the three major features of parallel programming supported by the Intel® Fortran compiler: OpenMP\*, Auto-parallelization, and Auto-vectorization. Each of these features contributes to the application performance depending on the number of processors, target architecture (IA-32 or Itanium® architecture), and the nature of the application. The three features OpenMP, Auto-parallelization and Auto-vectorization, can be combined arbitrarily to contribute to the application performance.

Parallel programming can be **explicit**, that is, defined by a programmer using OpenMP directives. Parallel programming can be **implicit**, that is, detected automatically by the compiler. Implicit parallelism implements Auto-parallelization of outer-most loops and Auto-vectorization of innermost loops.

Parallelism defined with OpenMP and Auto-parallelization directives is based on **thread-level** parallelism (**TLP**). Parallelism defined with Auto-vectorization techniques is based on **instruction-level** parallelism (**ILP**).

The Intel Fortran compiler supports OpenMP and Auto-parallelization on both IA-32 and Itanium architectures for multiprocessor systems as well as on single IA-32 processors with Hyper-Threading Technology (for Hyper-Threading Technology, refer to the IA-32 Intel® Architecture Optimization Reference Manual). Auto-vectorization is supported on the families of the Pentium®, Pentium with MMX(TM) technology, Pentium II, Pentium III, and Pentium 4 processors. To enhance the compilation of the code with Auto-vectorization, the users can also add vectorizer directives to their program. A closely related technique that is available on the Itanium-based systems is software pipelining (SWP).

The table below summarizes the different ways in which parallelism can be exploited with the Intel Fortran compiler.

Parallelism	
Explicit	Implicit
Parallelism programmed by the user	Parallelism generated by compiler and by user-supplied hints

<b>OpenMP* (TLP)</b> IA-32 and Itanium architectures	<b>Auto-parallelization (TLP)</b> of outer-most loops  IA-32 and Itanium architectures	<b>Auto-vectorization (ILP)</b> of inner-most loops  IA-32 only  Software pipelining for Itanium architecture
<b>Supported on</b>		<b>Supported on</b>
IA-32 or Itanium-based Multiprocessor systems;  IA-32 Hyper-Threading Technology-enabled systems.		Pentium®, Pentium with MMX™ Technology, Pentium II, Pentium III, and Pentium 4 processors

## Parallel Program Development

The Intel Fortran Compiler supports the OpenMP Fortran version 2.0 API specification available from the [www.openmp.org](http://www.openmp.org) web site. The OpenMP directives relieve the user from having to deal with the low-level details of iteration space partitioning, data sharing, and thread scheduling and synchronization.

The Auto-parallelization feature of the Intel Fortran Compiler automatically translates serial portions of the input program into semantically equivalent multithreaded code. Automatic parallelization determines the loops that are good worksharing candidates, performs the dataflow analysis to verify correct parallel execution, and partitions the data for threaded code generation as is needed in programming with OpenMP directives. The OpenMP and Auto-parallelization applications provide the performance gains from shared memory on multiprocessor systems and IA-32 processors with the Hyper-Threading Technology.

Auto-vectorization detects low-level operations in the program that can be done in parallel, and then converts the sequential program to process 2, 4, 8 or up to 16 elements in one operation, depending on the data type. In some cases auto-parallelization and vectorization can be combined for better performance results. For example, in the code below, TLP can be exploited in the outermost loop, while ILP can be exploited in the innermost loop.

```
DO I = 1, 100           ! execute groups of
iterations in different
                      ! threads (TLP)
DO J = 1, 32           ! execute in SIMD style with
multimedia
```

```

                ! extension (ILP)
A(J,I) = A(J,I) + 1
ENDDO
ENDDO

```

Auto-vectorization can help improve performance of an application that runs on the systems based on Pentium®, Pentium with MMX(TM) technology, Pentium II, Pentium III, and Pentium 4 processors.

The following table lists the options that enable Auto-vectorization, Auto-parallelization, and OpenMP support.

<b>Auto-vectorization, IA-32 only</b>	
<code>-x{K W N B P}</code>	Generates specialized code to run exclusively on processors with the extensions specified by {K W N B P}.
<code>-ax{K W N B P}</code>	Generates, in a single binary, code specialized to the extensions specified by {K W N B P} and also generic IA-32 code. The generic code is usually slower.
<code>-vec_report{0 1 2 3 4 5}</code>	Controls the diagnostic messages from the vectorizer, see subsection that follows the table.
<b>Auto-parallelization, IA-32 and Itanium architectures</b>	
<code>-parallel</code>	Enables the auto-parallelizer to generate multithreaded code for loops that can be safely executed in parallel. Default: OFF.
<code>-par_threshold{n}</code>	Sets a threshold for the auto-parallelization of loops based on the probability of profitable execution of the loop in parallel, n=0 to 100. n=0 implies "always." Default: n=75.
<code>-par_report{0 1 2 3}</code>	Controls the auto-parallelizer's diagnostic levels. Default: <code>-par_report1</code> .
<b>OpenMP, IA-32 and Itanium architectures</b>	
<code>-openmp</code>	Enables the parallelizer to generate multithreaded code based on the OpenMP directives. Default: OFF.
<code>-openmp_report{0 1 2}</code>	Controls the OpenMP parallelizer's diagnostic levels. Default: <code>/Qopenmp_report1</code> .
<code>-openmp_stubs</code>	Enables compilation of OpenMP programs in sequential mode. The OpenMP directives are ignored and a stub OpenMP library is linked. Default: OFF.

 **Note**

When both `-openmp` and `-parallel` are specified on the command line, the `-parallel` option is only honored in routines that do not contain OpenMP Directives. For routines that contain OpenMP directives, only the `-openmp` option is honored.

With the right choice of options, the programmers can:

- increase the performance of your application with minimum effort
- use compiler features to develop multithreaded programs faster

With a relatively small effort of adding the OpenMP directives to their code, the programmers can transform a sequential program into a parallel program. The following are examples of the OpenMP directives within the code:

```
!OMP$ PARALLEL PRIVATE(NUM), SHARED (X,A,B,C)
      !Defines a parallel region
!OMP$ PARALLEL DO ! Specifies a parallel region that
      ! implicitly contains a single DO directive
DO I = 1, 1000
NUM = FOO(B(i), C(I))
X(I) = BAR(A(I), NUM)
      ! Assume FOO and BAR have no side effects
ENDDO
```

See examples of the Auto-parallelization and Auto-vectorization directives in the respective sections.

## Auto-vectorization (IA-32 Only)

### Vectorization Overview

The vectorizer is a component of the Intel® Fortran Compiler that automatically uses SIMD instructions in the MMX(TM), SSE, and SSE2 instruction sets. The vectorizer detects operations in the program that can be done in parallel, and then converts the sequential operations like one SIMD instruction that processes 2, 4, 8 or up to 16 elements in parallel, depending on the data type.

This section provides options description, guidelines, and examples for Intel Fortran Compiler vectorization implemented by IA-32 compiler only. For additional information, see Publications on Compiler Optimizations.

The following list summarizes this section contents.

- Descriptions of compiler options to control vectorization
- Vectorization Key Programming Guidelines
- Discussion and general guidelines on vectorization levels:

—automatic vectorization

—vectorization with user intervention

- Examples demonstrating typical vectorization issues and resolutions

The Intel compiler supports a variety of directives that can help the compiler to generate effective vector instructions. See compiler directives supporting vectorization.

### Vectorizer Options

Vectorization is an IA-32-specific feature and can be summarized by the command line options described in the following tables. Vectorization depends upon the compiler's ability to disambiguate memory references. Certain options may enable the compiler to do better vectorization. These options can enable other optimizations in addition to vectorization. When an `-x{K|W|N|B|P}` or `-ax{K|W|N|B|P}` is used and `-O2` (which is ON by default) is also in effect, the vectorizer is enabled. The `-x{K|W|N|B|P}` or `-ax{K|W|N|B|P}` options enable vectorizer with `-O1` and `-O3` options also.

<code>-x{K W N B P}</code>	Generate specialized code to run exclusively on the processors supporting the extensions indicated by <code>{K W N B P}</code> . See Processor-
----------------------------	---



	Specific Exclusive Specialized Code (IA-32 only) for details.
<code>-ax{K W N B P}</code>	Generates, in a single binary, code specialized to the extensions specified by <code>{K W N B P}</code> and also generic IA-32 code. The generic code is usually slower. See Processor Automatic Non-Exclusive Specialized Code (IA-32 only) for details.
<code>-vec_report</code> <code>{0 1 2 3 4 5}</code> Default: <code>-vec_report1</code>	Controls the diagnostic messages from the vectorizer, see subsection that follows the table.

## Vectorization Reports

The `-vec_report{0|1|2|3|4|5}` options directs the compiler to generate the vectorization reports with different level of information as follows:

`-vec_report0`: no diagnostic information is displayed

`-vec_report1`: display diagnostics indicating loops successfully vectorized (default)

`-vec_report2`: same as `-vec_report1`, plus diagnostics indicating loops not successfully vectorized

`-vec_report3`: same as `-vec_report2`, plus additional information about any proven or assumed dependences

`-vec_report4`: indicate non-vectorized loops

`-vec_report5`: indicate non-vectorized loops and the reason why they were not vectorized.

## Usage with Other Options

The vectorization reports are generated in the final compilation phase when executable is generated. Therefore if you use the `-c` option and a `-vec_report{n}` option in the command line, no report will be generated.

If you use `-c`, `-ipo` and `-x{K|W|N|B|P}` or `-ax{K|W|N|B|P}` and `-vec_report{n}`, the compiler issues a warning and no report is generated.

To produce a report when using the above mentioned options, you need to add the `-ipo_obj` option. The combination of `-c` and `-ipo_obj` produces a single file compilation, and hence does generate object code, and eventually a report is generated.

The following commands generate vectorization report:

```
ifort -x{K|W|N|B|P} -vec_report3 file.f
```

```
ifort -x{K|W|N|B|P} -ipo -ipo_obj -vec_report3 file.f
```

```
ifort -c -x{K|W|N|B|P} -ipo -ipo_obj -vec_report3 file.f
```

## Loop Parallelization and Vectorization

Combining the `-parallel` and `-x{K|W|N|B|P}` options instructs the compiler to attempt both automatic loop parallelization and automatic loop vectorization in the same compilation. In most cases, the compiler will consider outermost loops for parallelization and innermost loops for vectorization. If deemed profitable, however, the compiler may even apply loop parallelization and vectorization to the same loop. See Guidelines for Effective Auto-parallelization Usage and Vectorization Key Programming Guidelines.

Note that in some rare cases successful loop parallelization (either automatically or by means of OpenMP\* directives) may affect the messages reported by the compiler for a non-vectorizable loop in a non-intuitive way.

## Vectorization Key Programming Guidelines

The goal of vectorizing compilers is to exploit single-instruction multiple data (SIMD) processing automatically. Users can help however by supplying the compiler with additional information; for example, directives. Review these guidelines and restrictions, see code examples in further topics, and check them against your code to eliminate ambiguities that prevent the compiler from achieving optimal vectorization.

### Guidelines

You will often need to make some changes to your loops.

#### For loop bodies -

##### Use:

- Straight-line code (a single basic block)
- Vector data only; that is, arrays and invariant expressions on the right hand side of assignments. Array references can appear on the left hand side of assignments.
- Only assignment statements

## Avoid:

- Function calls
- Unvectorizable operations (other than mathematical)
- Mixing vectorizable types in the same loop
- Data-dependent loop exit conditions
- Loop unrolling (compiler does it)
- Decomposing one loop with several statements in the body into several single-statement loops.

## Restrictions

Vectorization depends on the two major factors:

- **Hardware.** The compiler is limited by restrictions imposed by the underlying hardware. In the case of Streaming SIMD Extensions, the vector memory operations are limited to `stride-1` accesses with a preference to 16-byte-aligned memory references. This means that if the compiler abstractly recognizes a loop as vectorizable, it still might not vectorize it for a distinct target architecture.
- **Style.** The style in which you write source code can inhibit optimization. For example, a common problem with global pointers is that they often prevent the compiler from being able to prove that two memory references refer to distinct locations. Consequently, this prevents certain reordering transformations.

Many stylistic issues that prevent automatic vectorization by compilers are found in loop structures. The ambiguity arises from the complexity of the keywords, operators, data references, and memory operations within the loop bodies.

However, by understanding these limitations and by knowing how to interpret diagnostic messages, you can modify your program to overcome the known limitations and enable effective vectorization. The following sections summarize the capabilities and restrictions of the vectorizer with respect to loop structures.

## Data Dependence

Data dependence relations represent the required ordering constraints on the operations in serial loops. Because vectorization rearranges the order in which operations are executed, any auto-vectorizer must have at its disposal some form of data dependence analysis.

An example where data dependencies prohibit vectorization is shown below. In this example, the value of each element of an array is dependent on the value of its neighbor that was computed in the previous iteration.

**Data-dependent Loop**

```

REAL DATA(0:N)
INTEGER I
DO I=1, N-1
DATA(I) =DATA(I-1)*0.25+DATA(I)*0.5+DATA(I+1)*0.25
END DO

```

The loop in the above example is not vectorizable because the WRITE to the current element `DATA(I)` is dependent on the use of the preceding element `DATA(I-1)`, which has already been written to and changed in the previous iteration. To see this, look at the access patterns of the array for the first two iterations as shown below.

**Data Dependence Vectorization Patterns**

```

I=1: READ DATA (0)
READ DATA (1)
READ DATA (2)
WRITE DATA (1)
I=2: READ DATA(1)
READ DATA (2)
READ DATA (3)
WRITE DATA (2)

```

In the normal sequential version of this loop, the value of `DATA(1)` read from during the second iteration was written to in the first iteration. For vectorization, it must be possible to do the iterations in parallel, without changing the semantics of the original loop.

**Data Dependence Analysis**

Data dependence analysis involves finding the conditions under which two memory accesses may overlap. Given two references in a program, the conditions are defined by:

- whether the referenced variables may be aliases for the same (or overlapping) regions in memory, and, for array references
- the relationship between the subscripts

For IA-32, data dependence analyzer for array references is organized as a series of tests, which progressively increase in power as well as in time and space costs. First, a number of simple tests are performed in a dimension-by-dimension manner, since independence in any dimension will exclude any dependence relationship. Multidimensional arrays references that may cross their declared dimension boundaries can be converted to their linearized form before the tests are applied. Some of the simple tests that can be used are the fast greatest common divisor (GCD) test and the extended bounds test. The GCD

test proves independence if the GCD of the coefficients of loop indices cannot evenly divide the constant term. The extended bounds test checks for potential overlap of the extreme values in subscript expressions. If all simple tests fail to prove independence, we eventually resort to a powerful hierarchical dependence solver that uses Fourier-Motzkin elimination to solve the data dependence problem in all dimensions. For more details of data dependence theory and data dependence analysis, refer to the Publications on Compiler Optimizations.

## Loop Constructs

Loops can be formed with the usual DO-ENDDO and DO WHILE, or by using a GOTO and a label. However, the loops must have a single entry and a single exit to be vectorized. Following are the examples of correct and incorrect usages of loop constructs.

Correct Usage
<pre> SUBROUTINE FOO (A, B, C) DIMENSION A(100),B(100), C(100) INTEGER I I = 1 DO WHILE (I .LE. 100) A(I) = B(I) * C(I) IF (A(I) .LT. 0.0) A(I) = 0.0 I = I + 1 ENDDO RETURN END </pre>
Incorrect Usage
<pre> SUBROUTINE FOO (A, B, C) DIMENSION A(100),B(100), C(100) INTEGER I I = 1 DO WHILE (I .LE. 100) A(I) = B(I) * C(I) C The next statement allows early C exit from the loop and prevents C vectorization of the loop. IF (A(I) .LT. 0.0) GOTO 10 I = I + 1 ENDDO 10 CONTINUE RETURN END </pre>

## Loop Exit Conditions

Loop exit conditions determine the number of iterations that a loop executes. For example, fixed indexes for loops determine the iterations. The loop iterations must be countable; that is, the number of iterations must be expressed as one of the following:

- a constant
- a loop invariant term
- a linear function of outermost loop indices

Loops whose exit depends on computation are not countable. Examples below show countable and non-countable loop constructs.

<b>Correct Usage for Countable Loop, Example 1</b>
<pre> SUBROUTINE FOO (A, B, C, N, LB) DIMENSION A(N),B(N),C(N) INTEGER N, LB, I, COUNT ! Number of iterations is "N - LB + 1" COUNT = N DO WHILE (COUNT .GE. LB) A(I) = B(I) * C(I) COUNT = COUNT - 1 I = I + 1 ENDDO ! LB is not defined within loop RETURN END </pre>
<b>Correct Usage for Countable Loop, Example 2</b>
<pre> ! Number of iterations is (N-M+2) /2 SUBROUTINE FOO (A, B, C, M, N, LB) DIMENSION A(N),B(N),C(N) INTEGER I, L, M, N I = 1; DO L = M,N,2 A(I) = B(I) * C(I) I = I + 1 ENDDO RETURN END </pre>
<b>Incorrect Usage for Non-countable Loop</b>
<pre> ! Number of iterations is dependent on A(I) SUBROUTINE FOO (A, B, C) DIMENSION A(100),B(100),C(100) </pre>

```

INTEGER I
I = 1
DO WHILE (A(I) .GT. 0.0)
A(I) = B(I) * C(I)
I = I + 1
ENDDO
RETURN
END

```

## Types of Loop Vectorized

For integer loops, the 64-bit MMX(TM) technology and 128-bit Streaming SIMD Extensions (SSE) provide SIMD instructions for most arithmetic and logical operators on 32-bit, 16-bit, and 8-bit integer data types. Vectorization may proceed if the final precision of integer wrap-around arithmetic will be preserved. A 32-bit shift-right operator, for instance, is not vectorized in 16-bit mode if the final stored value is a 16-bit integer. Because the MMX(TM) and SSE instruction sets are not fully orthogonal (shifts on byte operands, for instance, are not supported), not all integer operations can actually be vectorized.

For loops that operate on 32-bit single-precision and 64-bit double-precision floating-point numbers, SSE provides SIMD instructions for the arithmetic operators '+', '-', '\*', and '/'. In addition, SSE provides SIMD instructions for the binary `MIN` and `MAX` and unary `SQRT` operators. SIMD versions of several other mathematical operators (like the trigonometric functions `SIN`, `COS`, `TAN`) are supported in software in a vector mathematical run-time library that is provided with the Intel® Fortran Compiler, of which the compiler takes advantage.

## Strip-mining and Cleanup

Strip-mining, also known as loop sectioning, is a loop transformation technique for enabling SIMD-encodings of loops, as well as providing a means of improving memory performance. By fragmenting a large loop into smaller segments or strips, this technique transforms the loop structure in two ways:

- It increases the temporal and spatial locality in the data cache if the data are reusable in different passes of an algorithm.
- It reduces the number of iterations of the loop by a factor of the length of each "vector," or number of operations being performed per SIMD operation. In the case of Streaming SIMD Extensions, this vector or strip-length is reduced by 4 times: four floating-point data items per single Streaming SIMD Extensions single-precision floating-point SIMD operation are processed.

First introduced for vectorizers, this technique consists of the generation of code when each vector operation is done for a size less than or equal to the maximum vector length on a given vector machine.

The compiler automatically strip-mines your loop and generates a cleanup loop.

### Stripmining and Cleanup Loops

#### Before Vectorization

```
i = 1
do while (i<=n)
a(i) = b(i) + c(i) ! Original loop
code
i = i + 1
end do
```

#### After Vectorization

```
!The vectorizer generates the
following two loops
i = 1
do while (i < (n - mod(n,4)))
! Vector strip-mined loop.
a(i:i+3) = b(i:i+3) + c(i:i+3)
i = i + 4
end do
do while (i <= n)
a(i) = b(i) + c(i)      !Scalar
clean-up loop
i = i + 1
end do
```

### Loop Blocking

It is possible to treat loop blocking as strip-mining in two or more dimensions. Loop blocking is a useful technique for memory performance optimization. The main purpose of loop blocking is to eliminate as many cache misses as possible. This technique transforms the memory domain into smaller chunks rather than sequentially traversing through the entire memory domain. Each chunk should be small enough to fit all the data for a given computation into the cache, thereby maximizing data reuse.

Consider the following example. The two-dimensional array *A* is referenced in the *j* (column) direction and then in the *i* (row) direction (column-major order); array *B* is referenced in the opposite manner (row-major order). Assume the memory layout is in column-major order; therefore, the access strides of array *A* and *B* for the code would be 1 and *MAX*, respectively.

In the **B.** example: *BS* = *block\_size*; *MAX* must be evenly divisible by *BS*.

### Loop Blocking of Arrays

#### A. Original loop



```

REAL A(MAX,MAX), B(MAX,MAX)
DO I =1, MAX
DO J = 1, MAX
  A(I,J) = A(I,J) + B(J,I)
ENDDO
ENDDO

```

### B. Transformed Loop after blocking

```

REAL A(MAX,MAX), B(MAX,MAX)
DO I =1, MAX, BS
DO J = 1, MAX, BS
  DO II = I, I+MAX, BS-1
    DO J = J, J+MAX, BS-1
      A(II,JJ) = A(II,JJ) +
B(JJ,II)
    ENDDO
  ENDDO
ENDDO
ENDDO

```

## Statements in the Loop Body

The vectorizable operations are different for floating point and integer data.

### Floating-point Array Operations

The statements within the loop body may be `REAL` operations (typically on arrays). Arithmetic operations supported are addition, subtraction, multiplication, division, negation, square root, `MAX`, `MIN`, and mathematical functions such as `SIN` and `COS`. Note that conversion to/from some types of floats is not valid. Operation on `DOUBLE PRECISION` types is not valid, unless optimizing for an Intel®

Pentium® 4 and Intel® Xeon(TM) processors' system, and Intel® Pentium® M processor, using the `-xW` or `-axW` compiler option.

### Integer Array Operations

The statements within the loop body may be arithmetic or logical operations (again, typically for arrays). Arithmetic operations are limited to such operations as addition, subtraction, `ABS`, `MIN`, and `MAX`. Logical operations include bitwise `AND`, `OR` and `XOR` operators. You can mix data types only if the conversion can be done without a loss of precision. Some example operators where you can mix data types are multiplication, shift, or unary operators.

## Other Operations

No statements other than the preceding floating-point and integer operations are permitted. The loop body cannot contain any function calls other than the ones described above.

## Vectorization Examples

This section contains simple examples of some common issues in vector programming.

### Argument Aliasing: A Vector Copy

The loop in the example of a vector copy operation does not vectorize because the compiler cannot prove that `DEST(A(I))` and `DEST(B(I))` are distinct.

#### Unvectorizable Copy Due to Unproven Distinction

```


SUBROUTINE
VEC_COPY(DEST, A, B, LEN)
DIMENSION DEST(*)
INTEGER A(*), B(*)
INTEGER LEN, I
DO I=1,LEN
DEST(A(I)) = DEST(B(I))
END DO
RETURN
END
    
```

## Data Alignment


A 16-byte or greater data structure or array should be aligned so that the beginning of each structure or array element is aligned in a way that its base address is a multiple of 16.

The Misaligned Data Crossing 16-Byte Boundary figure shows the effect of a data cache unit (DCU) split due to misaligned data. The code loads the misaligned data across a 16-byte boundary, which results in an additional memory access causing a six- to twelve-cycle stall. You can avoid the stalls if you know that the data is aligned and you specify to assume alignment.

#### Misaligned Data Crossing 16-Byte Boundary



After vectorization, the loop is executed as shown in figure below.



Both the vector iterations  $A(1:4) = B(1:4)$ ; and  $A(5:8) = B(5:8)$ ; can be implemented with aligned moves if both the elements  $A(1)$  and  $B(1)$  are 16-byte aligned.

### Caution

If you specify the vectorizer with incorrect alignment options, the compiler will generate code with unexpected behavior. Specifically, using aligned moves on unaligned data, will result in an illegal instruction exception!

### Alignment Strategy

The compiler has at its disposal several alignment strategies in case the alignment of data structures is not known at compile-time. A simple example is shown below (several other strategies are supported as well). If in the loop shown below the alignment of  $A$  is unknown, the compiler will generate a prelude loop that iterates until the array reference, that occurs the most, hits an aligned address. This makes the alignment properties of  $A$  known, and the vector loop is optimized accordingly. In this case, the vectorizer applies dynamic loop peeling, a specific Intel® Fortran feature.

#### Data Alignment Example

##### Original loop:

```
SUBROUTINE DOIT(A)
REAL A(100)          ! alignment of argument
A is unknown
DO I = 1, 100
A(I) = A(I) + 1.0
ENDDO
END SUBROUTINE
```

##### Aligning Data

```
! The vectorizer will apply dynamic loop
peeling as follows:
SUBROUTINE DOIT(A)
REAL A(100)
! let P be (A%16)where A is address of
A(1)
IF (P .NE. 0) THEN
P = (16 - P) / 4      ! determine run-time
```

```

peeling
                                ! factor
DO I = 1, P
A(I) = A(I) + 1.0
ENDDO
ENDIF
! Now this loop starts at a 16-byte
boundary,
! and will be vectorized accordingly
DO I = P + 1, 100
A(I) = A(I) + 1.0
ENDDO
END SUBROUTINE

```

## Loop Interchange and Subscripts: Matrix Multiply

Matrix multiplication is commonly written as shown in the following example.

```

DO I=1, N
DO J=1, N
DO K=1, N
  C(I,J) = C(I,J) +
A(I,K)*B(K,J)
END DO
END DO
END DO

```

The use of `B(K,J)`, is not a `stride-1` reference and therefore will not normally be vectorizable. If the loops are interchanged, however, all the references will become `stride-1` as in the Matrix Multiplication with Stride-1 example that follows.

### Note

Interchanging is not always possible because of dependencies, which can lead to different results.

#### Matrix Multiplication with Stride-1

```

DO J=1,N
DO K=1,N
DO I=1,N
  C(I,J) = C(I,J) +
A(I,K)*B(K,J)
ENDDO
ENDDO
ENDDO

```

For additional information, see Publications on Compiler Optimizations.

# Auto-parallelization

## Auto-parallelization Overview

The auto-parallelization feature of the Intel® Fortran Compiler automatically translates serial portions of the input program into equivalent multithreaded code. The auto-parallelizer analyzes the dataflow of the program's loops and generates multithreaded code for those loops which can be safely and efficiently executed in parallel. This enables the potential exploitation of the parallel architecture found in symmetric multiprocessor (SMP) systems.

Automatic parallelization relieves the user from:

- having to deal with the details of finding loops that are good worksharing candidates
- performing the dataflow analysis to verify correct parallel execution
- partitioning the data for threaded code generation as is needed in programming with OpenMP\* directives.

The parallel run-time support provides the same run-time features as found in OpenMP, such as handling the details of loop iteration modification, thread scheduling, and synchronization.

While OpenMP directives enable serial applications to transform into parallel applications quickly, the programmer must explicitly identify specific portions of the application code that contain parallelism and add the appropriate compiler directives. Auto-parallelization triggered by the `-parallel` option automatically identifies those loop structures, which contain parallelism. During compilation, the compiler automatically attempts to decompose the code sequences into separate threads for parallel processing. No other effort by the programmer is needed.

The following example illustrates how a loop's iteration space can be divided so that it can be executed concurrently on two threads:

### Original Serial Code

```
do i=1,100
  a(i) = a(i) + b(i) * c(i)
enddo
```

### Transformed Parallel Code

Thread 1
<pre>do i=1,50   a(i) = a(i) + b(i) * c(i) enddo</pre>

### Thread 2

```
do i=50,100
  a(i) = a(i) + b(i) * c(i)
enddo
```

## Programming with Auto-parallelization

Auto-parallelization feature implements some concepts of OpenMP, such as worksharing construct (with the `PARALLEL DO` directive). See Programming with OpenMP for worksharing construct. This section provides specifics of auto-parallelization.

### Guidelines for Effective Auto-parallelization Usage

A loop is parallelizable if:

- The loop is countable at compile time: this means that an expression representing how many times the loop will execute (also called "the loop trip count") can be generated just before entering the loop.
- There are no `FLOW (READ after WRITE)`, `OUTPUT (WRITE after READ)` or `ANTI (WRITE after READ)` loop-carried data dependences. A loop-carried data dependence occurs when the same memory location is referenced in different iterations of the loop. At the compiler's discretion, a loop may be parallelized if any assumed inhibiting loop-carried dependencies can be resolved by run-time dependency testing.

The compiler may generate a run-time test for the profitability of executing in parallel for loop with loop parameters that are not compile-time constants.

### Coding Guidelines

Enhance the power and effectiveness of the auto-parallelizer by following these coding guidelines:

- Expose the trip count of loops whenever possible; specifically use constants where the trip count is known and save loop parameters in local variables.
- Avoid placing structures inside loop bodies that the compiler may assume to carry dependent data, for example, procedure calls, ambiguous indirect references or global references.
- Insert the `!DEC$ PARALLEL` directive to disambiguate assumed data dependencies.
- Insert the `!DEC$ NOPARALLEL` directive before loops known to have insufficient work to justify the overhead of sharing among threads.

## Auto-parallelization Data Flow

For auto-parallelization processing, the compiler performs the following steps:

**Data flow analysis ---> Loop classification ---> Dependence analysis ---> High-level parallelization --> Data partitioning ---> Multi-threaded code generation.**

These steps include:

- Data flow analysis: compute the flow of data through the program
- Loop classification: determine loop candidates for parallelization based on correctness and efficiency as shown by threshold analysis
- Dependence analysis: compute the dependence analysis for references in each loop nest
- High-level parallelization:
  - analyze dependence graph to determine loops which can execute in parallel.
  - compute run-time dependency
- Data partitioning: examine data reference and partition based on the following types of access: `SHARED`, `PRIVATE`, and `FIRSTPRIVATE`
- Multi-threaded code generation:
  - modify loop parameters
  - generate entry/exit per threaded task
  - generate calls to parallel run-time routines for thread creation and synchronization

## Auto-parallelization: Enabling, Options, Directives, and Environment Variables

To enable the auto-parallelizer, use the `-parallel` option. The `-parallel` option detects parallel loops capable of being executed safely in parallel and automatically generates multithreaded code for these loops. An example of the command using auto-parallelization is as follows:

```
ifort -c -parallel myprog.f
```

## Auto-parallelization Options

The `-parallel` option enables the auto-parallelizer if the `-O2` (or `-O3`) optimization option is also on (the default is `-O2`). The `-parallel` option detects parallel loops capable of being executed safely in parallel and automatically generates multithreaded code for these loops.

<code>-parallel</code>	Enables the auto-parallelizer
<code>-par_threshold{1-100}</code>	Controls the work threshold needed for auto-parallelization, see later subsection.
<code>-par_report{1 2 3}</code>	Controls the diagnostic messages from the auto-parallelizer, see later subsection.

## Auto-parallelization Directives

Auto-parallelization uses two specific directives, `!DEC$ PARALLEL` and `!DEC$ NOPARALLEL`.

## Auto-parallelization Directives Format and Syntax

The format of Intel Fortran auto-parallelization compiler directive is:

```
<prefix> <directive>
```

where the brackets above mean:

- `<xxx>`: the prefix and directive are required

**For fixed form** source input, the prefix is `!DEC$` or `CDEC$`

**For free form** source input, the prefix is `!DEC$` only.

The prefix is followed by the directive name; for example:

```
!DEC$ PARALLEL
```

Since auto-parallelization directives begin with an exclamation point, the directives take the form of comments if you omit the `-parallel` option.

## Examples

The `!DEC$ PARALLEL` directive instructs the compiler to ignore dependencies which it assumes may exist and which would prevent correct parallelization in the



immediately following loop. However, if dependencies are proven, they are not ignored.

The `!DEC$ NOPARALLEL` directive disables auto-parallelization for the immediately following loop.

```

program main
parameter (n=100)
integer x(n),a(n)

!DEC$ NOPARALLEL
do i=1,n
x(i) = i
enddo

!DEC$ PARALLEL
do i=1,n
a( x(i) ) = i
enddo
end

```

### Auto-parallelization Environment Variables

Option	Description	Default
OMP_NUM_THREADS	Controls the number of threads used.	Number of processors currently installed in the system while generating the executable
OMP_SCHEDULE	Specifies the type of run-time scheduling.	static

### Auto-parallelization Threshold Control and Diagnostics

#### Threshold Control

The `-par_threshold{n}` option sets a threshold for auto-parallelization of loops based on the probability of profitable execution of the loop in parallel. The value of `n` can be from 0 to 100. The default value is 100. The `-par_threshold{n}` option should be used when the computation work in loops cannot be determined at compile-time.

The meaning for various values of `n` is as follows:

- `n = 100`. Parallelization will only proceed when performance gains are predicted based on the compiler analysis data. This is the default. This

value is used when `-par_threshold{n}` is not specified on the command line or is used without specifying a value of `n`.

- `n = 0`, `-par_threshold0` is specified. The loops get auto-parallelized regardless of computation work volume, that is, parallelize always.
- The intermediate 1 to 99 values represent the percentage probability for profitable speed-up. For example, `n=50` would mean: parallelize only if there is a 50% probability of the code speeding up if executed in parallel.

The compiler applies a heuristic that tries to balance the overhead of creating multiple threads versus the amount of work available to be shared amongst the threads.

## Diagnostics

The `-par_report{0|1|2|3}` option controls the auto-parallelizer's diagnostic levels 0, 1, 2, or 3 as follows:

`-par_report0` = no diagnostic information is displayed.

`-par_report1` = indicates loops successfully auto-parallelized (default). Issues a "LOOP AUTO-PARALLELIZED" message for parallel loops.

`-par_report2` = indicates successfully auto-parallelized loops as well as unsuccessful loops.

`-par_report3` = same as 2 plus additional information about any proven or assumed dependences inhibiting auto-parallelization (reasons for not parallelizing).

## Example of Parallelization Diagnostics Report

Example below shows an output generated by `-par_report3` as a result from the command:

```
ifort -c -parallel -par_report3 myprog.f90
```

where the program `myprog.f90` is as follows:

```
program myprog
  integer a(10000), q
C Assumed side effects
  do i=1,10000
    a(i) = foo(i)
  enddo
C Actual dependence
  do i=1,10000
    a(i) = a(i-1) + i
```

```
enddo  
end
```

#### Example of `-par_report` Output

```
program myprog  
  procedure: myprog  
  serial loop: line 5: not a parallel candidate  
    due to statement at line 6  
  serial loop: line 9  
    flow data dependence from line 10 to line  
10, due to "a"  
12 Lines Compiled
```

#### Troubleshooting Tips

- Use `-par_threshold0` to see if the compiler assumed there was not enough computational work
- Use `-par_report3` to view diagnostics
- Use `!DIR$ PARALLEL` directive to eliminate assumed data dependencies
- Use `-ipo` to eliminate assumed side-effects done to function calls.

## Parallelization with OpenMP\*

### Parallelization with OpenMP\* Overview

The Intel® Fortran Compiler supports the OpenMP\* Fortran version 2.0 API specification, except for the `WORKSHARE` directive. OpenMP provides symmetric multiprocessing (SMP) with the following major features:

- Relieves the user from having to deal with the low-level details of iteration space partitioning, data sharing, and thread scheduling and synchronization.
- Provides the benefit of the performance available from shared memory, multiprocessor systems; and, for IA-32 systems, from Hyper-Threading Technology-enabled systems (for Hyper-Threading Technology, refer to the IA-32 Intel® Architecture Optimization Reference Manual).

The Intel Fortran Compiler performs transformations to generate multithreaded code based on the user's placement of OpenMP directives in the source program making it easy to add threading to existing software. The Intel compiler supports all of the current industry-standard OpenMP directives, except `workshare`, and compiles parallel programs annotated with OpenMP directives.

In addition, the Intel Fortran Compiler provides Intel-specific extensions to the OpenMP Fortran version 2.0 specification including run-time library routines and environment variables.

#### Note

As with many advanced features of compilers, you must properly understand the functionality of the OpenMP directives in order to use them effectively and avoid unwanted program behavior.

See parallelization options summary for all options of the OpenMP feature in the Intel Fortran Compiler. For complete information on the OpenMP standard, visit the [www.openmp.org](http://www.openmp.org) web site. For complete Fortran language specifications, see the OpenMP Fortran version 2.0 specifications.

### Parallel Processing with OpenMP

To compile with OpenMP, you need to prepare your program by annotating the code with OpenMP directives in the form of the Fortran program comments. The Intel Fortran Compiler first processes the application and produces a multithreaded version of the code which is then compiled. The output is a Fortran executable with the parallelism implemented by threads that execute parallel regions or constructs. See Programming with OpenMP.

## Performance Analysis

For performance analysis of your program, you can use the VTune(TM) analyzer and/or the Intel® Threading Tools to show performance information. You can obtain detailed information about which portions of the code that require the largest amount of time to execute and where parallel performance problems are located.

## Programming with OpenMP

The Intel® Fortran Compiler accepts a Fortran program containing OpenMP directives as input and produces a multithreaded version of the code. When the parallel program begins execution, a single thread exists. This thread is called the master thread. The master thread will continue to process serially until it encounters a parallel region.

### Parallel Region

A parallel region is a block of code that must be executed by a team of threads in parallel. In the OpenMP Fortran API, a parallel construct is defined by placing OpenMP directives `parallel` at the beginning and `end parallel` at the end of the code segment. Code segments thus bounded can be executed in parallel.

A structured block of code is a collection of one or more executable statements with a single point of entry at the top and a single point of exit at the bottom.

The Intel Fortran Compiler supports worksharing and synchronization constructs. Each of these constructs consists of one or two specific OpenMP directives and sometimes the enclosed or following structured block of code. For complete definitions of constructs, see the OpenMP Fortran version 2.0 specifications.

At the end of the parallel region, threads wait until all team members have arrived. The team is logically disbanded (but may be reused in the next parallel region), and the master thread continues serial execution until it encounters the next parallel region.

### Worksharing Construct

A worksharing construct divides the execution of the enclosed code region among the members of the team created on entering the enclosing parallel region. When the master thread enters a parallel region, a team of threads is formed. Starting from the beginning of the parallel region, code is replicated (executed by all team members) until a `worksharing` construct is encountered. A worksharing construct divides the execution of the enclosed code among the members of the team that encounter it.

The OpenMP `sections` or `do` constructs are defined as `worksharing` constructs because they distribute the enclosed work among the threads of the current team. A `worksharing` construct is only distributed if it is encountered during dynamic execution of a parallel region. If the `worksharing` construct occurs lexically inside of the parallel region, then it is always executed by distributing the work among the team members. If the `worksharing` construct is not lexically (explicitly) enclosed by a parallel region (that is, it is `orphaned`), then the `worksharing` construct will be distributed among the team members of the closest dynamically-enclosing parallel region, if one exists. Otherwise, it will be executed serially.

When a thread reaches the end of a `worksharing` construct, it may wait until all team members within that construct have completed their work. When all of the work defined by the `worksharing` construct is finished, the team exits the `worksharing` construct and continues executing the code that follows.

A combined parallel/`worksharing` construct denotes a parallel region that contains only one `worksharing` construct.

## Parallel Processing Directive Groups

The parallel processing directives include the following groups:

### Parallel Region

- `PARALLEL` and `END PARALLEL`

### Worksharing Construct

- The `DO` and `END DO` directives specify parallel execution of loop iterations.
- The `SECTIONS` and `END SECTIONS` directives specify parallel execution for arbitrary blocks of sequential code. Each `SECTION` is executed once by a thread in the team.
- The `SINGLE` and `END SINGLE` directives define a section of code where exactly one thread is allowed to execute the code; threads not chosen to execute this section ignore the code.

## Combined Parallel/Worksharing Constructs

The combined parallel/`worksharing` constructs provide an abbreviated way to specify a parallel region that contains a single `worksharing` construct. The combined parallel/`worksharing` constructs are:

- `PARALLEL DO` and `END PARALLEL DO`
- `PARALLEL SECTIONS` and `END PARALLEL SECTIONS`

## Synchronization and MASTER

Synchronization is the interthread communication that ensures the consistency of shared data and coordinates parallel execution among threads. Shared data is consistent within a team of threads when all threads obtain the identical value when the data is accessed. A synchronization construct is used to insure this consistency of the shared data.

- The OpenMP synchronization directives are `CRITICAL`, `ORDERED`, `ATOMIC`, `FLUSH`, and `BARRIER`.
  - Within a parallel region or a `worksharing` construct only one thread at a time is allowed to execute the code within a `CRITICAL` construct.
  - The `ORDERED` directive is used in conjunction with a `DO` or `SECTIONS` construct to impose a serial order on the execution of a section of code.
  - The `ATOMIC` directive is used to update a memory location in an uninterruptable fashion.
  - The `FLUSH` directive is used to insure that all threads in a team have a consistent view of memory.
  - A `BARRIER` directive forces all team members to gather at a particular point in code. Each team member that executes a `BARRIER` waits at the `BARRIER` until all of the team members have arrived. A `BARRIER` cannot be used within `worksharing` or other synchronization constructs due to the potential for deadlock.
- The `MASTER` directive is used to force execution by the master thread.

See the list of OpenMP Directives and Clauses.

## Data Sharing

Data sharing is specified at the start of a parallel region or `worksharing` construct by using the `shared` and `private` clauses. All variables in the `shared` clause are shared among the members of a team. It is the application's responsibility to:

- synchronize access to these variables. All variables in the `private` clause are private to each team member. For the entire parallel region, assuming  $t$  team members, there are  $t+1$  copies of all the variables in the `private` clause: one global copy that is active outside parallel regions and a `private` copy for each team member.
- initialize `private` variables at the start of a parallel region, unless the `firstprivate` clause is specified. In this case, the `private` copy is initialized from the global copy at the start of the construct at which the `firstprivate` clause is specified.

- update the global copy of a `private` variable at the end of a parallel region. However, the `lastprivate` clause of a `DO` directive enables updating the global copy from the team member that executed serially the last iteration of the loop.

In addition to `shared` and `private` variables, individual variables and entire common blocks can be privatized using the `threadprivate` directive.

## Orphaned Directives

OpenMP contains a feature called orphaning which dramatically increases the expressiveness of parallel directives. Orphaning is a situation when directives related to a parallel region are not required to occur lexically within a single program unit. Directives such as `critical`, `barrier`, `sections`, `single`, `master`, and `do`, can occur by themselves in a program unit, dynamically “binding” to the enclosing parallel region at run time.

Orphaned directives enable parallelism to be inserted into existing code with a minimum of code restructuring. Orphaning can also improve performance by enabling a single parallel region to bind with multiple `do` directives located within called subroutines. Consider the following code segment:

```
...
!$omp parallel
call phase1
call phase2
!$omp end parallel
...

subroutine phase1
!$omp do private(i)
shared(n)
do i = 1, n
call some_work(i)
end do
!$omp end do
end

subroutine phase2
!$omp do private(j)
shared(n)
do j = 1, n
call more_work(j)
end do
!$omp end do
end
```



## Orphaned Directives Usage Rules

- An orphaned `worksharing` construct (`section`, `single`, `do`) is executed by a team consisting of one thread, that is, serially.
- Any collective operation (`worksharing` construct or `barrier`) executed inside of a `worksharing` construct is illegal.
- It is illegal to execute a collective operation (`worksharing` construct or `barrier`) from within a synchronization region (`critical/ordered`).
- The opening and closing directives of a directive pair (for example, `do - end do`) must occur in a single block of the program.
- Private scoping of a variable can be specified at a `worksharing` construct. Shared scoping must be specified at the parallel region. For complete details, see the OpenMP Fortran version 2.0 specifications.

## Preparing Code for OpenMP Processing

The following are the major stages and steps of preparing your code for using OpenMP. Typically, the first two stages can be done on uniprocessor or multiprocessor systems; later stages are typically done only on multiprocessor systems.

### Before Inserting OpenMP Directives

Before inserting any OpenMP parallel directives, verify that your code is safe for parallel execution by doing the following:

- Place local variables on the stack. This is the default behavior of the Intel Fortran Compiler when `-openmp` is used.
- Use `-auto` or similar (`-auto_scalar`) compiler option to make the locals automatic. This is the default behavior of the Intel Fortran Compiler when `-openmp` is used. Avoid using compiler options that inhibit stack allocation of local variables. By default (`-auto_scalar`) local scalar variables become shared across threads, so you may need to add synchronization code to ensure proper access by threads.

### Analyze

The analysis includes the following major actions:

- Profile the program to find out where it spends most of its time. This is the part of the program that benefits most from parallelization efforts. This stage can be accomplished using VTune(TM) analyzer or basic PGO options.
- Wherever the program contains nested loops, choose the outer-most loop, which has very few cross-iteration dependencies.

## Restructure

- To restructure your program for successful OpenMP implementation, you can perform some or all of the following actions:
  1. If a chosen loop is able to execute iterations in parallel, introduce a `parallel do` construct around this loop.
  2. Try to remove any cross-iteration dependencies by rewriting the algorithm.
  3. Synchronize the remaining cross-iteration dependencies by placing `critical` constructs around the uses and assignments to variables involved in the dependencies.
  4. List the variables that are present in the loop within appropriate `shared`, `private`, `lastprivate`, `firstprivate`, or `reduction` clauses.
  5. List the `do` index of the parallel loop as `private`. This step is optional.
  6. `common` block elements must not be placed on the `private` list if their global scope is to be preserved. The `threadprivate` directive can be used to privatize to each thread the `common` block containing those variables with global scope. `threadprivate` creates a copy of the `common` block for each of the threads in the team.
  7. Any I/O in the parallel region should be synchronized.
  8. Identify more parallel loops and restructure them.
  9. If possible, merge adjacent `parallel do` constructs into a single parallel region containing multiple `do` directives to reduce execution overhead.

## Tune

The tuning process should include minimizing the sequential code in critical sections and load balancing by using the `schedule` clause or the `omp_schedule` environment variable.

### Note

This step is typically performed on a multiprocessor system.

## Parallel Processing Thread Model

This topic explains the processing of the parallelized program and adds more definitions of the terms used in the parallel programming.

## The Execution Flow

As mentioned in previous topic, a program containing OpenMP Fortran API compiler directives begins execution as a single process, called the **master**

thread of execution. The master thread executes sequentially until the first **parallel construct** is encountered.

In OpenMP Fortran API, the `PARALLEL` and `END PARALLEL` directives define the parallel construct. When the master thread encounters a parallel construct, it creates a **team** of threads, with the master thread becoming the master of the team. The program statements enclosed by the parallel construct are executed in parallel by each thread in the team. These statements include routines called from within the enclosed statements.

The statements enclosed lexically within a construct define the **static extent** of the construct. The **dynamic extent** includes the static extent as well as the routines called from within the construct. When the `END PARALLEL` directive is encountered, the threads in the team synchronize at that point, the team is dissolved, and only the master thread continues execution. The other threads in the team enter a wait state.

You can specify any number of parallel constructs in a single program. As a result, thread teams can be created and dissolved many times during program execution.

### Using Orphaned Directives

In routines called from within parallel constructs, you can also use directives. Directives that are not in the lexical extent of the parallel construct, but are in the dynamic extent, are called orphaned directives. Orphaned directives allow you to execute major portions of your program in parallel with only minimal changes to the sequential version of the program. Using this functionality, you can code parallel constructs at the top levels of your program call tree and use directives to control execution in any of the called routines. For example:

```
subroutine F
...
!$OMP
parallel...
...
  call G
...
  subroutine G
  ...
  !$OMP DO...
  ...
```

The `!$OMP DO` is an orphaned directive because the parallel region it will execute in is not lexically present in `G`.

## Data Environment Directive

A data environment directive controls the data environment during the execution of parallel constructs.

You can control the data environment within parallel and worksharing constructs. Using directives and data environment clauses on directives, you can:

- Privatize named common blocks by using `THREADPRIVATE` directive
- Control data scope attributes by using the `THREADPRIVATE` directive's clauses.

The data scope attribute clauses are:

- `COPYIN`
- `DEFAULT`
- `PRIVATE`
- `FIRSTPRIVATE`
- `LASTPRIVATE`
- `REDUCTION`
- `SHARED`

You can use several directive clauses to control the data scope attributes of variables for the duration of the construct in which you specify them. If you do not specify a data scope attribute clause on a directive, the default is `SHARED` for those variables affected by the directive.

For detailed descriptions of the clauses, see the OpenMP Fortran version 2.0 specifications.

## Pseudo Code of the Parallel Processing Model

A sample program using some of the more common OpenMP directives is shown in the code example that follows. This example also indicates the difference between serial regions and parallel regions.

```
program main      ! Begin Serial Execution
...              ! Only the master thread executes
!$omp parallel    ! Begin a Parallel Construct,
                 ! form a team
...              ! This is Replicated Code where
                 ! each team ! member executes the
                 ! same code
!$omp sections    ! Begin a Worksharing Construct
!$omp section     ! One unit of work
...              !
```

```

!$omp section      ! Another unit of work
...              !
!$omp end do      ! Wait until both units of work
sections         complete
...              ! More Replicated Code
!$omp do          ! Begin a Worksharing Construct,
do              ! each iteration is a unit of work
...            ! Work is distributed among the
end do          team
!$omp end do      ! End of Worksharing Construct,
nowait          nowait is
...            ! specified
!$omp end parallel ! More Replicated Code
parallel        ! End of Parallel Construct,
...            disband team ! and continue with
end            serial execution
              ! Possibly more Parallel
              Constructs
              ! End serial execution

```

## Compiling with OpenMP, Directive Format, and Diagnostics

To run the Intel® Fortran Compiler in OpenMP mode, you need to invoke the Intel compiler with the `-openmp` option:

```
ifort -openmp input_file(s)
```

Before you run the multithreaded code, you can set the number of desired threads to the OpenMP environment variable, `OMP_NUM_THREADS`. See the OpenMP Environment Variables section for further information. The Intel Extension Routines topic describes the OpenMP extensions to the specification that have been added by Intel in the Intel® Fortran Compiler.

### **-openmp Option**

The `-openmp` option enables the parallelizer to generate multithreaded code based on the OpenMP directives. The code can be executed in parallel on both uniprocessor and multiprocessor systems.

The `-openmp` option works with both `-O0` (no optimization) and any optimization level of `-O1`, `-O2` (default) and `-O3`. Specifying `-O0` with `-openmp` helps to debug OpenMP applications.

When you use the `-openmp` option, the compiler sets the `-auto` option (causes all variables to be allocated on the stack, rather than in local static storage.) for the compiler unless you specified it on the command line.

## OpenMP Directive Format and Syntax

The OpenMP directives use the following format:

```
<prefix> <directive> [<clause> [[,] <clause> . . .]]
```

where the brackets above mean:

- `<xxx>`: the prefix and directive are required
- `[<xxx>]`: if a directive uses one clause or more, the clause(s) is required
- `[,]`: commas between the `<clause>`s are optional.

**For fixed form** source input, the prefix is `!$omp` or `c$omp`

**For free form** source input, the prefix is `!$omp` only.

The prefix is followed by the directive name; for example:

```
!$omp parallel
```

Since OpenMP directives begin with an exclamation point, the directives take the form of comments if you omit the `-openmp` option.

## Syntax for Parallel Regions in the Source Code

The OpenMP constructs defining a parallel region have one of the following syntax forms:

```
!$omp <directive>  
<structured block of code>  
!$omp end <directive>
```

or

```
!$omp <directive>  
<structured block of code>
```

or

```
!$omp <directive>
```

where `<directive>` is the name of a particular OpenMP directive.

## OpenMP Diagnostic Reports

The `-openmp_report{0|1|2}` option controls the OpenMP parallelizer's diagnostic levels 0, 1, or 2 as follows:

`-openmp_report0` = no diagnostic information is displayed.

`-openmp_report1` = display diagnostics indicating loops, regions, and sections successfully parallelized.


`-openmp_report2` = same as `-openmp_report1` plus diagnostics indicating master constructs, single constructs, critical constructs, ordered constructs, atomic directives, etc. successfully handled.

The default is `-openmp_report1`.

## OpenMP Directives and Clauses Summary

This topic provides a summary of the OpenMP directives and clauses. For detailed descriptions, see the OpenMP Fortran version 2.0 specifications.

### OpenMP Directives

Directive	Description
<code>parallel</code> <code>end parallel</code>	Defines a parallel region.
<code>do</code> <code>end do</code>	Identifies an iterative worksharing construct in which the iterations of the associated loop should be executed in parallel.
<code>sections</code> <code>end sections</code>	Identifies a non-iterative worksharing construct that specifies a set of structured blocks that are to be divided among threads in a team.
<code>section</code>	Indicates that the associated structured block should be executed in parallel as part of the enclosing sections construct.
<code>single</code> <code>end single</code>	Identifies a construct that specifies that the associated structured block is executed by only one thread in the team.
<code>parallel do</code> <code>end parallel</code> <code>do</code>	<p>A shortcut for a <code>parallel</code> region that contains a single <code>do</code> directive.</p> <p> <b>Note</b> The <code>parallel do</code> or <code>do</code> OpenMP directive must be immediately followed by a <code>do</code> statement (<code>do-stmt</code> as defined by R818 of the ANSI Fortran standard). If you place another statement or an OpenMP directive</p>

	between the <code>parallel do</code> or <code>do</code> directive and the <code>do</code> statement, the Intel Fortran Compiler issues a syntax error.
<code>parallel sections</code> <code>end parallel sections</code>	Provides a shortcut form for specifying a parallel region containing a single <code>sections</code> construct.
<code>master</code> <code>end master</code>	Identifies a construct that specifies a structured block that is executed by only the <code>master</code> thread of the team.
<code>critical[lock]</code> <code>end critical[lock]</code>	Identifies a construct that restricts execution of the associated structured block to a single thread at a time. Each thread waits at the beginning of the critical construct until no other thread is executing a critical construct with the same <code>lock</code> argument.
<code>barrier</code>	Synchronizes all the threads in a team. Each thread waits until all of the other threads in that team have reached this point.
<code>atomic</code>	Ensures that a specific memory location is updated atomically, rather than exposing it to the possibility of multiple, simultaneously writing threads.
<code>flush [(list)]</code>	Specifies a "cross-thread" sequence point at which the implementation is required to ensure that all the threads in a team have a consistent view of certain objects in memory. The optional <code>list</code> argument consists of a comma-separated list of variables to be flushed.
<code>ordered</code> <code>end ordered</code>	The structured block following an <code>ordered</code> directive is executed in the order in which iterations would be executed in a sequential loop.
<code>threadprivate (list)</code>	Makes the named <code>common</code> blocks or variables private to a thread. The <code>list</code> argument consists of a comma-separated list of <code>common</code> blocks or variables.

## OpenMP Clauses

Clause	Description
<code>private (list)</code>	Declares variables in <code>list</code> to be <code>private</code> To each thread in a team.
<code>firstprivate (list)</code>	Same as <code>private</code> , but the copy of each variable in the <code>list</code> is initialized using the value of the



	original variable existing before the construct.
<code>lastprivate (list)</code>	Same as <code>private</code> , but the original variables in <code>list</code> are updated using the values assigned to the corresponding <code>private</code> variables in the last iteration in the <code>do</code> construct loop or the last <code>section</code> construct.
<code>copyprivate (list)</code>	Uses <code>private</code> variables in <code>list</code> to broadcast values, or pointers to shared objects, from one member of a team to the other members at the end of a single construct.
<code>nowait</code>	Specifies that threads need not wait at the end of <code>worksharing</code> constructs until they have completed execution. The threads may proceed past the end of the <code>worksharing</code> constructs as soon as there is no more work available for them to execute.
<code>shared (list)</code>	Shares variables in <code>list</code> among all the threads in a team.
<code>default (mode)</code>	Determines the default data-scope attributes of variables not explicitly specified by another clause. Possible values for <code>mode</code> are <code>private</code> , <code>shared</code> , or <code>none</code> .
<code>reduction ( {operator intrinsic} : list )</code>	Performs a reduction on variables that appear in <code>list</code> with the operator <code>operator</code> or the intrinsic procedure name <code>intrinsic</code> ; <code>operator</code> is one of the following: <code>+</code> , <code>*</code> , <code>.and.</code> , <code>.or.</code> , <code>.eqv.</code> , <code>.neqv.</code> ; <code>intrinsic</code> refers to one of the following: <code>max</code> , <code>min</code> , <code>iand</code> , <code>ior</code> , or <code>ieor</code> .
<code>ordered end ordered</code>	Used in conjunction with a <code>do</code> or <code>sections</code> construct to impose a serial order on the execution of a section of code. If <code>ordered</code> constructs are contained in the dynamic extent of the <code>do</code> construct, the <code>ordered</code> clause must be present

	on the <code>do</code> directive.
<code>if (scalar_logical_expression)</code>	The enclosed parallel region is executed in parallel only if the <code>scalar_logical_expression</code> evaluates to <code>.true.</code> ; otherwise the parallel region is serialized.
<code>num_threads(scalar_integer_expression)</code>	Requests the number of threads specified by <code>scalar_integer_expression</code> for the parallel region.
<code>schedule (type[,chunk])</code>	Specifies how iterations of the <code>do</code> construct are divided among the threads of the team. Possible values for the <code>type</code> argument are <code>static</code> , <code>dynamic</code> , <code>guided</code> , and <code>runtime</code> . The optional <code>chunk</code> argument must be a positive scalar integer expression.
<code>copyin (list)</code>	Specifies that the master thread's data values be copied to the <code>threadprivate</code> 's copies of the common blocks or variables specified in <code>list</code> at the beginning of the parallel region.

### Directives and Clauses Cross-reference

Directive	Uses These Clauses
PARALLEL END PARALLEL	COPYIN, DEFAULT, PRIVATE, FIRSTPRIVATE, REDUCTION, SHARED
DO END DO	PRIVATE, FIRSTPRIVATE, LASTPRIVATE, REDUCTION, SCHEDULE
SECTIONS END SECTIONS	PRIVATE, FIRSTPRIVATE, LASTPRIVATE, REDUCTION
SECTION	PRIVATE, FIRSTPRIVATE, LASTPRIVATE, REDUCTION
SINGLE END SINGLE	PRIVATE, FIRSTPRIVATE
PARALLEL DO END PARALLEL DO	COPYIN, DEFAULT, PRIVATE, FIRSTPRIVATE, LASTPRIVATE, REDUCTION, SHARED, SCHEDULE
PARALLEL SECTIONS END PARALLEL SECTIONS	COPYIN, DEFAULT, PRIVATE, FIRSTPRIVATE, LASTPRIVATE, REDUCTION, SHARED

MASTER END MASTER	None
CRITICAL[ <i>lock</i> ] END CRITICAL[ <i>lock</i> ]	None
BARRIER	None
ATOMIC	None
FLUSH [( <i>list</i> )]	None
ORDERED END ORDERED	None
THREADPRIVATE ( <i>list</i> )	None

## OpenMP Directive Descriptions

### Parallel Region Directives

The `PARALLEL` and `END PARALLEL` directives define a parallel region as follows:

```
!$OMP PARALLEL
! parallel region
!$OMP END PARALLEL
```

When a thread encounters a parallel region, it creates a team of threads and becomes the master of the team. You can control the number of threads in a team by the use of an environment variable or a run-time library call, or both.

### Clauses Used

The `PARALLEL` directive takes an optional comma-separated list of clauses that specify as follows:

- `IF`: whether the statements in the parallel region are executed in parallel by a team of threads or serially by a single thread.
- `PRIVATE`, `FIRSTPRIVATE`, `SHARED`, or `REDUCTION`: variable types
- `DEFAULT`: variable data scope attribute
- `COPYIN`: master thread common block values are copied to `THREADPRIVATE` copies of the common block

### Changing the Number of Threads

Once created, the number of threads in the team remains constant for the duration of that parallel region. To explicitly change the number of threads used in the next parallel region, call the `OMP_SET_NUM_THREADS` run-time library routine from a serial portion of the program. This routine overrides any value you may have set using the `OMP_NUM_THREADS` environment variable.

Assuming you have used the `OMP_NUM_THREADS` environment variable to set the number of threads to 6, you can change the number of threads between parallel regions as follows:

```
CALL OMP_SET_NUM_THREADS(3)
!$OMP PARALLEL
. . .
!$OMP END PARALLEL
CALL OMP_SET_NUM_THREADS(4)
!$OMP PARALLEL DO
. . .
!$OMP END PARALLEL DO
```

### Setting Units of Work

Use the worksharing directives such as `DO`, `SECTIONS`, and `SINGLE` to divide the statements in the parallel region into units of work and to distribute those units so that each unit is executed by one thread.

In the following example, the `!$OMP DO` and `!$OMP END DO` directives and all the statements enclosed by them comprise the static extent of the parallel region:

```
!$OMP PARALLEL
!$OMP DO
DO I=1,N
  B(I) = (A(I) + A(I-1))/ 2.0
END DO
!$OMP END DO
!$OMP END PARALLEL
```

In the following example, the `!$OMP DO` and `!$OMP END DO` directives and all the statements enclosed by them, including all statements contained in the `WORK` subroutine, comprise the dynamic extent of the parallel region:

```
!$OMP PARALLEL
DEFAULT(SHARED)
!$OMP DO
DO I=1,N
  CALL WORK(I,N)
END DO
!$OMP END DO
!$OMP END PARALLEL
```

### Setting Conditional Parallel Region Execution

When an `IF` clause is present on the `PARALLEL` directive, the enclosed code region is executed in parallel only if the scalar logical expression evaluates to

.TRUE.. Otherwise, the parallel region is serialized. When there is no IF clause, the region is executed in parallel by default.

In the following example, the statements enclosed within the !\$OMP DO and !\$OMP END DO directives are executed in parallel only if there are more than three processors available. Otherwise the statements are executed serially:

```
!$OMP PARALLEL IF (OMP_GET_NUM_PROCS() .GT.
3)
!$OMP DO
DO I=1,N
  Y(I) = SQRT(Z(I))
END DO
!$OMP END DO
!$OMP END PARALLEL
```

If a thread executing a parallel region encounters another parallel region, it creates a new team and becomes the master of that new team. By default, **nested parallel regions** are always executed by a team of one thread.

### Note

To achieve better performance than sequential execution, a parallel region must contain one or more worksharing constructs so that the team of threads can execute work in parallel. It is the contained worksharing constructs that lead to the performance enhancements offered by parallel processing.

## Worksharing Construct Directives

A worksharing construct must be enclosed dynamically within a parallel region if the worksharing directive is to execute in parallel. No new threads are launched and there is no implied barrier on entry to a worksharing construct.

The worksharing constructs are:

- DO and END DO directives
- SECTIONS, SECTION, and END SECTIONS directives
- SINGLE and END SINGLE directives

## DO and END DO

The DO directive specifies that the iterations of the immediately following DO loop must be dispatched across the team of threads so that each iteration is executed by a single thread. The loop that follows a DO directive cannot be a DO WHILE or a DO loop that does not have loop control. The iterations of the DO loop are dispatched among the existing team of threads.

The DO directive optionally lets you:

- Control data scope attributes (see Controlling Data Scope Attributes)
- Use the SCHEDULE clause to specify schedule type and chunk size (see Specifying Schedule Type and Chunk Size)

### Clauses Used

The clauses for DO directive specify:

- Whether variables are PRIVATE, FIRSTPRIVATE, LASTPRIVATE, or REDUCTION
- How loop iterations are SCHEDULED onto threads
- In addition, the ORDERED clause must be specified if the ORDERED directive appears in the dynamic extent of the DO directive.
- If you do not specify the optional NOWAIT clause on the END DO directive, threads synchronize at the END DO directive. If you specify NOWAIT, threads do not synchronize, and threads that finish early proceed directly to the instructions following the END DO directive.

### Usage Rules

- You cannot use a GOTO statement, or any other statement, to transfer control onto or out of the DO construct.
- If you specify the optional END DO directive, it must appear immediately after the end of the DO loop. If you do not specify the END DO directive, an END DO directive is assumed at the end of the DO loop, and threads synchronize at that point.
- The loop iteration variable is private by default, so it is not necessary to declare it explicitly.

### SECTIONS, SECTION and END SECTIONS

Use the noniterative worksharing SECTIONS directive to divide the enclosed sections of code among the team. Each section is executed just one time by one thread.

Each section should be preceded with a SECTION directive, except for the first section, in which the SECTION directive is optional. The SECTION directive must appear within the lexical extent of the SECTIONS and END SECTIONS directives.

The last section ends at the END SECTIONS directive. When a thread completes its section and there are no undispatched sections, it waits at the END SECTION directive unless you specify NOWAIT.

The `SECTIONS` directive takes an optional comma-separated list of clauses that specifies which variables are `PRIVATE`, `FIRSTPRIVATE`, `LASTPRIVATE`, or `REDUCTION`.

The following example shows how to use the `SECTIONS` and `SECTION` directives to execute subroutines `X_AXIS`, `Y_AXIS`, and `Z_AXIS` in parallel. The first `SECTION` directive is optional:

```
!$OMP PARALLEL
!$OMP SECTIONS
!$OMP SECTION
  CALL X_AXIS
!$OMP SECTION
  CALL Y_AXIS
!$OMP SECTION
  CALL Z_AXIS
!$OMP END SECTIONS
!$OMP END PARALLEL
```

## **SINGLE and END SINGLE**

Use the `SINGLE` directive when you want just one thread of the team to execute the enclosed block of code.

Threads that are not executing the `SINGLE` directive wait at the `END SINGLE` directive unless you specify `NOWAIT`.

The `SINGLE` directive takes an optional comma-separated list of clauses that specifies which variables are `PRIVATE` or `FIRSTPRIVATE`.

When the `END SINGLE` directive is encountered, an implicit barrier is erected and threads wait until all threads have finished. This can be overridden by using the `NOWAIT` option.

In the following example, the first thread that encounters the `SINGLE` directive executes subroutines `OUTPUT` and `INPUT`:

```
!$OMP PARALLEL
  DEFAULT(SHARED)
  CALL WORK(X)
!$OMP BARRIER
!$OMP SINGLE
  CALL OUTPUT(X)
  CALL INPUT(Y)
!$OMP END SINGLE
  CALL WORK(Y)
!$OMP END PARALLEL
```

## Combined Parallel/Worksharing Constructs

The combined parallel/worksharing constructs provide an abbreviated way to specify a parallel region that contains a single worksharing construct. The combined parallel/worksharing constructs are:

- PARALLEL DO
- PARALLEL SECTIONS

### PARALLEL DO and END PARALLEL DO

Use the `PARALLEL DO` directive to specify a parallel region that implicitly contains a single `DO` directive.

You can specify one or more of the clauses for the `PARALLEL` and the `DO` directives.

The following example shows how to parallelize a simple loop. The loop iteration variable is private by default, so it is not necessary to declare it explicitly. The `END PARALLEL DO` directive is optional:

```
!$OMP PARALLEL DO
  DO I=1,N
    B(I) = (A(I) + A(I-1)) / 2.0
  END DO
!$OMP END PARALLEL DO
```

### PARALLEL SECTIONS and END PARALLEL SECTIONS

Use the `PARALLEL SECTIONS` directive to specify a parallel region that implicitly contains a single `SECTIONS` directive.

You can specify one or more of the clauses for the `PARALLEL` and the `SECTIONS` directives.

The last section ends at the `END PARALLEL SECTIONS` directive.

In the following example, subroutines `X_AXIS`, `Y_AXIS`, and `Z_AXIS` can be executed concurrently. The first `SECTION` directive is optional. Note that all `SECTION` directives must appear in the lexical extent of the `PARALLEL SECTIONS/END PARALLEL SECTIONS` construct:

```
!$OMP PARALLEL SECTIONS
!$OMP SECTION
  CALL X_AXIS
!$OMP SECTION
  CALL Y_AXIS
```



```
!$OMP SECTION
  CALL Z_AXIS
!$OMP END PARALLEL SECTIONS
```

## Synchronization Constructs

Synchronization constructs are used to ensure the consistency of shared data and to coordinate parallel execution among threads.

The synchronization constructs are:

- `ATOMIC` directive
- `BARRIER` directive
- `CRITICAL` directive
- `FLUSH` directive
- `MASTER` directive
- `ORDERED` directive

### ATOMIC Directive

Use the `ATOMIC` directive to ensure that a specific memory location is updated atomically instead of exposing the location to the possibility of multiple, simultaneously writing threads.

This directive applies only to the immediately following statement, which must have one of the following forms:

```
x = x operator expr
```

```
x = expr operator x
```

```
x = intrinsic (x, expr)
```

```
x = intrinsic (expr, x)
```

In the preceding statements:

- *x* is a scalar variable of intrinsic type
- *expr* is a scalar expression that does not reference *x*
- *intrinsic* is either `MAX`, `MIN`, `IAND`, `IOR`, or `IEOR`
- *operator* is either `+`, `*`, `-`, `/`, `.AND.`, `.OR.`, `.EQV.`, or `.NEQV.`

This directive permits optimization beyond that of a critical section around the assignment. An implementation can replace all `ATOMIC` directives by enclosing the statement in a critical section. All of these critical sections must use the same unique name.

Only the load and store of  $x$  are atomic; the evaluation of  $expr$  is not atomic. To avoid race conditions, all updates of the location in parallel must be protected by using the `ATOMIC` directive, except those that are known to be free of race conditions. The function *intrinsic*, the operator *operator*, and the assignment must be the intrinsic function, operator, and assignment.

This restriction applies to the `ATOMIC` directive: All references to storage location  $x$  must have the same type parameters.

In the following example, the collection of  $Y$  locations is updated atomically:

```
!$OMP ATOMIC
  Y = Y + B(I)
```

### BARRIER Directive

To synchronize all threads within a parallel region, use the `BARRIER` directive. You can use this directive only within a parallel region defined by using the `PARALLEL` directive. You cannot use the `BARRIER` directive within the `DO`, `PARALLEL DO`, `SECTIONS`, `PARALLEL SECTIONS`, and `SINGLE` directives.

When encountered, each thread waits at the `BARRIER` directive until all threads have reached the directive.

In the following example, the `BARRIER` directive ensures that all threads have executed the first loop and that it is safe to execute the second loop:

```
c$OMP PARALLEL
c$OMP DO PRIVATE(i)
      DO i = 1, 100
        b(i) = i
      END DO
c$OMP BARRIER
c$OMP DO PRIVATE(i)
      DO i = 1, 100
        a(i) = b(101-i)
      END DO
c$OMP END PARALLEL
```

### CRITICAL and END CRITICAL

Use the `CRITICAL` and `END CRITICAL` directives to restrict access to a block of code, referred to as a critical section, to one thread at a time.

A thread waits at the beginning of a critical section until no other thread in the team is executing a critical section having the same name.

When a thread enters the critical section, a latch variable is set to closed and all other threads are locked out. When the thread exits the critical section at the `END CRITICAL` directive, the latch variable is set to open, allowing another thread access to the critical section.

If you specify a critical section name in the `CRITICAL` directive, you must specify the same name in the `END CRITICAL` directive. If you do not specify a name for the `CRITICAL` directive, you cannot specify a name for the `END CRITICAL` directive.

All unnamed `CRITICAL` directives map to the same name. Critical section names are global to the program.

The following example includes several `CRITICAL` directives, and illustrates a queuing model in which a task is dequeued and worked on. To guard against multiple threads dequeuing the same task, the dequeuing operation must be in a critical section. Because there are two independent queues in this example, each queue is protected by `CRITICAL` directives having different names, `X_AXIS` and `Y_AXIS`, respectively:

```
!$OMP PARALLEL
DEFAULT( PRIVATE , SHARED( X, Y )
!$OMP CRITICAL( X_AXIS )
    CALL DEQUEUE( IX_NEXT, X )
!$OMP END CRITICAL( X_AXIS )
    CALL WORK( IX_NEXT, X )
!$OMP CRITICAL( Y_AXIS )
    CALL DEQUEUE( IY_NEXT, Y )
!$OMP END CRITICAL( Y_AXIS )
    CALL WORK( IY_NEXT, Y )
!$OMP END PARALLEL
```

Unnamed critical sections use the global lock from the Pthread package. This allows you to synchronize with other code by using the same lock. Named locks are created and maintained by the compiler and can be significantly more efficient.

### FLUSH Directive

Use the `FLUSH` directive to identify a synchronization point at which a consistent view of memory is provided. Thread-visible variables are written back to memory at this point.

To avoid flushing all thread-visible variables at this point, include a list of comma-separated named variables to be flushed.

The following example uses the `FLUSH` directive for point-to-point synchronization between thread 0 and thread 1 for the variable `ISYNC`:

```
!$OMP PARALLEL DEFAULT(PRIVATE),SHARED(ISYNC)
    IAM = OMP_GET_THREAD_NUM()
    ISYNC(IAM) = 0
!$OMP BARRIER
    CALL WORK()
! I Am Done With My Work, Synchronize With My
Neighbor
    ISYNC(IAM) = 1
!$OMP FLUSH(ISYNC)
! Wait Till Neighbor Is Done
    DO WHILE (ISYNC(NEIGH) .EQ. 0)
!$OMP FLUSH(ISYNC)
    END DO
!$OMP END PARALLEL
```

## MASTER and END MASTER

Use the `MASTER` and `END MASTER` directives to identify a block of code that is executed only by the master thread.

The other threads of the team skip the code and continue execution. There is no implied barrier at the `END MASTER` directive.

In the following example, only the master thread executes the routines `OUTPUT` and `INPUT`:

```
!$OMP PARALLEL
DEFAULT(SHARED)
    CALL WORK(X)
!$OMP MASTER
    CALL OUTPUT(X)
    CALL INPUT(Y)
!$OMP END MASTER
    CALL WORK(Y)
!$OMP END PARALLEL
```

## ORDERED and END ORDERED

Use the `ORDERED` and `END ORDERED` directives within a `DO` construct to allow work within an ordered section to execute sequentially while allowing work outside the section to execute in parallel.

When you use the `ORDERED` directive, you must also specify the `ORDERED` clause on the `DO` directive.

Only one thread at a time is allowed to enter the ordered section, and then only in the order of loop iterations.

In the following example, the code prints out the indexes in sequential order:

```

!$OMP DO ORDERED, SCHEDULE(DYNAMIC)
  DO I=LB,UB,ST
    CALL WORK(I)
  END DO
  SUBROUTINE WORK(K)
!$OMP ORDERED
  WRITE(*,*) K
!$OMP END ORDERED

```

## THREADPRIVATE Directive

You can make named common blocks private to a thread, but global within the thread, by using the `THREADPRIVATE` directive.

Each thread gets its own copy of the common block with the result that data written to the common block by one thread is not directly visible to other threads. During serial portions and `MASTER` sections of the program, accesses are to the master thread copy of the common block.

You cannot use a thread private common block or its constituent variables in any clause other than the `COPYIN` clause.

In the following example, common blocks `BLK1` and `FIELDS` are specified as thread private:

```

COMMON /BLK1/ SCRATCH
      COMMON /FIELDS/ XFIELD, YFIELD,
ZFIELD
!$OMP THREADPRIVATE(/BLK1/, /FIELDS/)

```

## OpenMP Clause Descriptions

### Controlling Data Scope

#### Data Scope Attribute Clauses Overview

You can use several directive clauses to control the data scope attributes of variables for the duration of the construct in which you specify them. If you do not specify a data scope attribute clause on a directive, the default is `SHARED` for those variables affected by the directive.

Each of the data scope attribute clauses accepts a list, which is a comma-separated list of named variables or named common blocks that are accessible in the scoping unit. When you specify named common blocks, they must appear between slashes (*/name/*).

Not all of the clauses are allowed on all directives, but the directives to which each clause applies are listed in the clause descriptions.

The data scope attribute clauses are:

- COPYIN
- DEFAULT
- PRIVATE
- FIRSTPRIVATE
- LASTPRIVATE
- REDUCTION
- SHARED

### **COPYIN Clause**

Use the `COPYIN` clause on the `PARALLEL`, `PARALLEL DO`, and `PARALLEL SECTIONS` directives to copy the data in the master thread common block to the thread private copies of the common block. The copy occurs at the beginning of the parallel region. The `COPYIN` clause applies only to common blocks that have been declared `THREADPRIVATE`.

You do not have to specify a whole common block to be copied in; you can specify named variables that appear in the `THREADPRIVATE` common block. In the following example, the common blocks `BLK1` and `FIELDS` are specified as thread private, but only one of the variables in common block `FIELDS` is specified to be copied in:

```
COMMON /BLK1/ SCRATCH
COMMON /FIELDS/ XFIELD, YFIELD, ZFIELD
!$OMP THREADPRIVATE(/BLK1/, /FIELDS/)
!$OMP PARALLEL
DEFAULT(PRIVATE), COPYIN(/BLK1/, ZFIELD)
```

### **DEFAULT Clause**

Use the `DEFAULT` clause on the `PARALLEL`, `PARALLEL DO`, and `PARALLEL SECTIONS` directives to specify a default data scope attribute for all variables within the lexical extent of a parallel region. Variables in `THREADPRIVATE` common blocks are not affected by this clause. You can specify only one `DEFAULT` clause on a directive. The default data scope attribute can be one of the following:

- PRIVATE

Makes all named objects in the lexical extent of the parallel region private to a thread. The objects include common block variables, but exclude `THREADPRIVATE` variables.

- SHARED

Makes all named objects in the lexical extent of the parallel region shared among all the threads in the team.

- NONE

Declares that there is no implicit default as to whether variables are `PRIVATE` or `SHARED`. You must explicitly specify the scope attribute for each variable in the lexical extent of the parallel region.

If you do not specify the `DEFAULT` clause, the default is `DEFAULT ( SHARED )`. However, loop control variables are always `PRIVATE` by default.

You can exempt variables from the default data scope attribute by using other scope attribute clauses on the parallel region as shown in the following example:

```
!$OMP PARALLEL DO DEFAULT(PRIVATE) ,  
FIRSTPRIVATE ( I ) , SHARED ( X ) ,  
!$OMP& SHARED ( R ) LASTPRIVATE ( I )
```

## PRIVATE, FIRSTPRIVATE, and LASTPRIVATE Clauses

### PRIVATE

Use the `PRIVATE` clause on the `PARALLEL`, `DO`, `SECTIONS`, `SINGLE`, `PARALLEL DO`, and `PARALLEL SECTIONS` directives to declare variables to be private to each thread in the team.

The behavior of variables declared `PRIVATE` is as follows:

- A new object of the same type and size is declared once for each thread in the team, and the new object is no longer storage associated with the original object.
- All references to the original object in the lexical extent of the directive construct are replaced with references to the private object.
- Variables defined as `PRIVATE` are undefined for each thread on entering the construct, and the corresponding shared variable is undefined on exit from a parallel construct.
- Contents, allocation state, and association status of variables defined as `PRIVATE` are undefined when they are referenced outside the lexical extent, but inside the dynamic extent, of the construct unless they are passed as actual arguments to called routines.

In the following example, the values of `I` and `J` are undefined on exit from the parallel region:

```

INTEGER I, J
      I = 1
      J = 2
!$OMP PARALLEL PRIVATE(I) FIRSTPRIVATE(J)
      I = 3
      J = J + 2
!$OMP END PARALLEL
      PRINT *, I, J

```

## FIRSTPRIVATE

Use the `FIRSTPRIVATE` clause on the `PARALLEL`, `DO`, `SECTIONS`, `SINGLE`, `PARALLEL DO`, and `PARALLEL SECTIONS` directives to provide a superset of the `PRIVATE` clause functionality.

In addition to the `PRIVATE` clause functionality, private copies of the variables are initialized from the original object existing before the parallel construct.

## LASTPRIVATE

Use the `LASTPRIVATE` clause on the `DO`, `SECTIONS`, `PARALLEL DO`, and `PARALLEL SECTIONS` directives to provide a superset of the `PRIVATE` clause functionality.

When the `LASTPRIVATE` clause appears on a `DO` or `PARALLEL DO` directive, the thread that executes the sequentially last iteration updates the version of the object it had before the construct.

When the `LASTPRIVATE` clause appears on a `SECTIONS` or `PARALLEL SECTIONS` directive, the thread that executes the lexically last section updates the version of the object it had before the construct.

Subobjects that are not assigned a value by the last iteration of the `DO` loop or the lexically last `SECTION` directive are undefined after the construct.

Correct execution sometimes depends on the value that the last iteration of a loop assigns to a variable. You must list all such variables as arguments to a `LASTPRIVATE` clause so that the values of the variables are the same as when the loop is executed sequentially. As shown in the following example, the value of `I` at the end of the parallel region is equal to `N+1`, as it would be with sequential execution.

```

!$OMP PARALLEL
!$OMP DO LASTPRIVATE(I)
      DO I=1,N
          A(I) = B(I) + C(I)
      END DO
!$OMP END PARALLEL
      CALL REVERSE(I)

```



## REDUCTION Clause

Use the REDUCTION clause on the PARALLEL, DO, SECTIONS, PARALLEL DO, and PARALLEL SECTIONS directives to perform a reduction on the specified variables by using an operator or intrinsic as shown:

```
REDUCTION (
  operator
  or
  intrinsic
  :list )
```

*Operator* can be one of the following: +, \*, -, .AND., .OR., .EQV., or .NEQV..

*Intrinsic* can be one of the following: MAX, MIN, IAND, IOR, or Ieor.

The specified variables must be named scalar variables of intrinsic type and must be SHARED in the enclosing context. A private copy of each specified variable is created for each thread as if you had used the PRIVATE clause. The private copy is initialized to a value that depends on the operator or intrinsic as shown in the Table Operators/Intrinsics and Initialization Values for Reduction Variables. The actual initialization value is consistent with the data type of the reduction variable.

### Operators/Intrinsics and Initialization Values for Reduction Variables

Operator/Intrinsic	Initialization Value
+	0
*	1
-	0
.AND.	.TRUE.
.OR.	.FALSE.
.EQV.	.TRUE.
.NEQV.	.FALSE.
MAX	Largest representable number
MIN	Smallest representable number
IAND	All bits on
IOR	0
IEOR	0

At the end of the construct to which the reduction applies, the shared variable is updated to reflect the result of combining the original value of the `SHARED` reduction variable with the final value of each of the private copies using the specified operator.

Except for subtraction, all of the reduction operators are associative and the compiler can freely reassociate the computation of the final value. The partial results of a subtraction reduction are added to form the final value.

The value of the shared variable becomes undefined when the first thread reaches the clause containing the reduction, and it remains undefined until the reduction computation is complete. Normally, the computation is complete at the end of the `REDUCTION` construct. However, if you use the `REDUCTION` clause on a construct to which `NOWAIT` is also applied, the shared variable remains undefined until a barrier synchronization has been performed. This ensures that all of the threads have completed the `REDUCTION` clause.

The `REDUCTION` clause is intended to be used on a region or worksharing construct in which the reduction variable is used only in reduction statements having one of the following forms:

```
x = x operator expr
x = expr operator x (except for subtraction)
x = intrinsic (x,expr)
x = intrinsic (expr, x)
```

Some reductions can be expressed in other forms. For instance, a `MAX` reduction might be expressed as follows:

```
IF (x .LT. expr) x = expr
```

Alternatively, the reduction might be hidden inside a subroutine call. Be careful that the operator specified in the `REDUCTION` clause matches the reduction operation.

Any number of reduction clauses can be specified on the directive, but a variable can appear only once in a `REDUCTION` clause for that directive as shown in the following example:

```
!$OMP DO REDUCTION(+: A, Y), REDUCTION(.OR.: AM)
```

The following example shows how to use the `REDUCTION` clause:

```
!$OMP PARALLEL DO
DEFAULT( PRIVATE ), SHARED( A, B, REDUCTION(+: A, B)
DO I=1, N
CALL WORK( ALOCAL, BLOCAL )
A = A + ALOCAL
```

```

      B = B + BLOCAL
    END DO
!$OMP END PARALLEL DO

```

## SHARED Clause

Use the `SHARED` clause on the `PARALLEL`, `PARALLEL DO`, and `PARALLEL SECTIONS` directives to make variables shared among all the threads in a team.

In the following example, the variables `X` and `NPOINTS` are shared among all the threads in the team:

```

!$OMP PARALLEL
DEFAULT( PRIVATE ) , SHARED( X , NPOINTS )
  IAM = OMP_GET_THREAD_NUM( )
  NP = OMP_GET_NUM_THREADS( )
  IPOINTS = NPOINTS/NP
  CALL SUBDOMAIN( X , IAM , IPOINTS )
!$OMP END PARALLEL

```

## Specifying Schedule Type and Chunk Size

The `SCHEDULE` clause of the `DO` or `PARALLEL DO` directive specifies a scheduling algorithm that determines how iterations of the `DO` loop are divided among and dispatched to the threads of the team. The `SCHEDULE` clause applies only to the current `DO` or `PARALLEL DO` directive.

Within the `SCHEDULE` clause, you must specify a **schedule type** and, optionally, a **chunk size**. A **chunk** is a contiguous group of iterations dispatched to a thread. Chunk size must be a scalar integer expression.

The following list describes the schedule types and how the chunk size affects scheduling:

- `STATIC`

The iterations are divided into pieces having a size specified by chunk. The pieces are statically dispatched to threads in the team in a round-robin manner in the order of thread number.

When chunk is not specified, the iterations are first divided into contiguous pieces by dividing the number of iterations by the number of threads in the team. Each piece is then dispatched to a thread before loop execution begins.

- `DYNAMIC`

The iterations are divided into pieces having a size specified by chunk. As each thread finishes its currently dispatched piece of the iteration space, the next piece is dynamically dispatched to the thread.

When no chunk is specified, the default is 1.

- GUIDED

The chunk size is decreased exponentially with each succeeding dispatch. Chunk specifies the minimum number of iterations to dispatch each time. If there are less than chunk number of iterations remaining, the rest are dispatched.

When no chunk is specified, the default is 1.

- RUNTIME

The decision regarding scheduling is deferred until run time. The schedule type and chunk size can be chosen at run time by using the `OMP_SCHEDULE` environment variable.

When you specify `RUNTIME`, you cannot specify a chunk size.

The following list shows which schedule type is used, in priority order:

1. The schedule type specified in the `SCHEDULE` clause of the current `DO` or `PARALLEL DO` directive
2. If the schedule type for the current `DO` or `PARALLEL DO` directive is `RUNTIME`, the default value specified in the `OMP_SCHEDULE` environment variable
3. The compiler default schedule type of `STATIC`

The following list shows which chunk size is used, in priority order:

1. The chunk size specified in the `SCHEDULE` clause of the current `DO` or `PARALLEL DO` directive
2. For `RUNTIME` schedule type, the value specified in the `OMP_SCHEDULE` environment variable
3. For `DYNAMIC` and `GUIDED` schedule types, the default value 1
4. If the schedule type for the current `DO` or `PARALLEL DO` directive is `STATIC`, the loop iteration space divided by the number of threads in the team.

## OpenMP Support Libraries

The Intel Fortran Compiler with OpenMP support provides a production support library, `libguide.a`. This library enables you to run an application under different execution modes. It is used for normal or performance-critical runs on applications that have already been tuned.

### Execution modes

The compiler with OpenMP enables you to run an application under different execution modes that can be specified at run time. The libraries support the serial, turnaround, and throughput modes. These modes are selected by using the `kmp_library` environment variable at run time.

#### Turnaround

In a multi-user environment where the load on the parallel machine is not constant or where the job stream is not predictable, it may be better to design and tune for throughput. This minimizes the total time to run multiple jobs simultaneously. In this mode, the worker threads will yield to other threads while waiting for more parallel work.

The throughput mode is designed to make the program aware of its environment (that is, the system load) and to adjust its resource usage to produce efficient execution in a dynamic environment. This mode is the default.

After completing the execution of a parallel region, threads wait for new parallel work to become available. After a certain period of time has elapsed, they stop waiting and sleep. Sleeping allows the threads to be used, until more parallel work becomes available, by non-OpenMP threaded code that may execute between parallel regions, or by other applications. The amount of time to wait before sleeping is set either by the `KMP_BLOCKTIME` environment variable or by the `kmp_set_blocktime()` function. A small `KMP_BLOCKTIME` value may offer better overall performance if your application contains non-OpenMP threaded code that executes between parallel regions. A larger `KMP_BLOCKTIME` value may be more appropriate if threads are to be reserved solely for use for OpenMP execution, but may penalize other concurrently-running OpenMP or threaded applications.

#### Throughput

In a dedicated (batch or single user) parallel environment where all processors are exclusively allocated to the program for its entire run, it is most important to effectively utilize all of the processors all of the time. The turnaround mode is designed to keep active all of the processors involved in the parallel computation

in order to minimize the execution time of a single job. In this mode, the worker threads actively wait for more parallel work, without yielding to other threads.

### Note

Avoid over-allocating system resources. This occurs if either too many threads have been specified, or if too few processors are available at run time. If system resources are over-allocated, this mode will cause poor performance. The throughput mode should be used instead if this occurs.

## OpenMP Environment Variables

This topic describes the standard OpenMP environment variables (with the `OMP_` prefix) and Intel-specific environment variables (with the `KMP_` prefix) that are Intel extensions to the standard Fortran Compiler .

### Standard Environment Variables

Variable	Description	Default
<code>OMP_SCHEDULE</code>	Sets the run-time schedule type and chunk size.	<code>static</code> , no chunk size specified
<code>OMP_NUM_THREADS</code>	Sets the number of threads to use during execution.	Number of processors
<code>OMP_DYNAMIC</code>	Enables ( <code>true</code> ) or disables ( <code>false</code> ) the dynamic adjustment of the number of threads.	<code>false</code>
<code>OMP_NESTED</code>	Enables ( <code>true</code> ) or disables ( <code>false</code> ) nested parallelism.	<code>false</code>

### Intel Extension Environment Variables

Environment Variable	Description	Default
<code>KMP_ALL_THREADS</code>	Sets the maximum number of threads that can be used by any parallel region.	$\max(32, 4 * \text{OMP\_NUM\_THREADS}, 4 * \text{number of processors})$

KMP_BLOCKTIME	<p>Sets the time, in milliseconds, that a thread should wait, after completing the execution of a parallel region, before sleeping.</p> <p>See also the throughput execution mode and the KMP_LIBRARY environment variable. Use the optional character suffix s, m, h, or d, to specify seconds, minutes, hours, or days.</p>	200 milliseconds
KMP_LIBRARY	<p>Selects the OpenMP run-time library throughput. The options for the variable value are: serial, turnaround, or throughput indicating the execution mode. The default value of throughput is used if this variable is not specified.</p>	throughput (execution mode)

KMP_MONITOR_STACKSIZE	Sets the number of bytes to allocate for the monitor thread, which is used for book-keeping during program execution. Use the optional suffix b, k, m, g, or t, to specify bytes, kilobytes, megabytes, gigabytes, or terabytes.	max(32k, system minimum thread stack size)
KMP_STACKSIZE	Sets the number of bytes to allocate for each parallel thread to use as its private stack. Use the optional suffix b, k, m, g, or t, to specify bytes, kilobytes, megabytes, gigabytes, or terabytes.	IA-32: 2m Itanium compiler: 4m
KMP_VERSION	Enables (set) or disables (unset) the printing of OpenMP run-time library version information during program execution.	Disabled

## OpenMP Run-time Library Routines

OpenMP provides several run-time library routines to assist you in managing your program in parallel mode. Many of these run-time library routines have corresponding environment variables that can be set as defaults. The run-time library routines enable you to dynamically change these factors to assist in controlling your program. In all cases, a call to a run-time library routine overrides any corresponding environment variable.



The following table specifies the interface to these routines. The names for the routines are in user name space. The `omp_lib.f`, `omp_lib.h` and `omp_lib.mod` header files are provided in the `include` directory of your compiler installation. The `omp_lib.h` header file is provided in the `include` directory of your compiler installation for use with the Fortran `INCLUDE` statement. The `omp_lib.mod` file is provided in the `Include` directory for use with the Fortran `USE` statement.

There are definitions for two different locks, `omp_lock_t` and `omp_nest_lock_t`, which are used by the functions in the table that follows.

This topic provides a summary of the OpenMP run-time library routines. For detailed descriptions, see the OpenMP Fortran version 2.0 specifications.

Function	Description
<b>Execution Environment Routines</b>	
subroutine <code>omp_set_num_threads(num_threads)</code> integer <code>num_threads</code>	Sets the number of threads to use for subsequent parallel regions.
integer function <code>omp_get_num_threads()</code>	Returns the number of threads that are being used in the current parallel region.
integer function <code>omp_get_max_threads()</code>	Returns the maximum number of threads that are available for parallel execution.
integer function <code>omp_get_thread_num()</code>	Determines the unique thread number of the thread currently executing this section of code.
integer function <code>omp_get_num_procs()</code>	Determines the number of processors available to the program.
logical function <code>omp_in_parallel()</code>	Returns <code>.true.</code> if called within the dynamic extent of a parallel region executing in parallel; otherwise returns <code>.false.</code>
subroutine <code>omp_set_dynamic(dynamic_threads)</code> logical <code>dynamic_threads</code>	Enables or disables dynamic adjustment of the number of threads

	used to execute a parallel region. If <i>dynamic_threads</i> is <i>.true.</i> , dynamic threads are enabled. If <i>dynamic_threads</i> is <i>.false.</i> , dynamic threads are disabled. Dynamics threads are disabled by default.
logical function <code>omp_get_dynamic()</code>	Returns <i>.true.</i> if dynamic thread adjustment is enabled, otherwise returns <i>.false.</i> .
subroutine <code>omp_set_nested(<i>nested</i>)</code> integer <i>nested</i>	Enables or disables nested parallelism. If <i>nested</i> is <i>.true.</i> , nested parallelism is enabled. If <i>nested</i> is <i>.false.</i> , nested parallelism is disabled. Nested parallelism is disabled by default.
logical function <code>omp_get_nested()</code>	Returns <i>.true.</i> if nested parallelism is enabled, otherwise returns <i>.false.</i> .
<b>Lock Routines</b>	
subroutine <code>omp_init_lock(<i>lock</i>)</code> integer (kind= <code>omp_lock_kind</code> ):: <i>lock</i>	Initializes the lock associated with <i>lock</i> for use in subsequent calls.
subroutine <code>omp_destroy_lock(<i>lock</i>)</code> integer (kind= <code>omp_lock_kind</code> ):: <i>lock</i>	Causes the lock associated with <i>lock</i> to become undefined.
subroutine <code>omp_set_lock(<i>lock</i>)</code> integer (kind= <code>omp_lock_kind</code> ):: <i>lock</i>	Forces the executing thread to wait until the lock associated with <i>lock</i> is available. The thread is granted ownership of the lock when it becomes available.
subroutine <code>omp_unset_lock(<i>lock</i>)</code> integer (kind= <code>omp_lock_kind</code> ):: <i>lock</i>	Releases the executing thread from ownership of

	the lock associated with <i>lock</i> . The behavior is undefined if the executing thread does not own the lock associated with <i>lock</i> .
<pre>logical omp_test_lock(lock) integer (kind=omp_lock_kind)::lock</pre>	Attempts to set the lock associated with <i>lock</i> . If successful, returns <code>.true.</code> , otherwise returns <code>.false.</code>
<pre>subroutine omp_init_nest_lock(lock) integer(kind=omp_nest_lock_kind)::lock</pre>	Initializes the nested lock associated with <i>lock</i> for use in the subsequent calls.
<pre>subroutine omp_destroy_nest_lock(lock) integer(kind=omp_nest_lock_kind)::lock</pre>	Causes the nested lock associated with <i>lock</i> to become undefined.
<pre>subroutine omp_set_nest_lock(lock) integer(kind=omp_nest_lock_kind)::lock</pre>	Forces the executing thread to wait until the nested lock associated with <i>lock</i> is available. The thread is granted ownership of the nested lock when it becomes available.
<pre>subroutine omp_unset_nest_lock(lock) integer(kind=omp_nest_lock_kind)::lock</pre>	Releases the executing thread from ownership of the nested lock associated with <i>lock</i> if the nesting count is zero. Behavior is undefined if the executing thread does not own the nested lock associated with <i>lock</i> .
<pre>integer omp_test_nest_lock(lock) integer(kind=omp_nest_lock_kind)::lock</pre>	Attempts to set the nested lock associated with <i>lock</i> . If successful, returns the nesting count, otherwise returns zero.
<b>Timing Routines</b>	
<pre>double-precision function omp_get_wtime()</pre>	Returns a double-precision value equal to

	the elapsed wallclock time (in seconds) relative to an arbitrary reference time. The reference time does not change during program execution.
double-precision function <code>omp_get_wtick()</code>	Returns a double-precision value equal to the number of seconds between successive clock ticks.

## Intel Extension Routines

The Intel® Fortran Compiler implements the following group of routines as an extension to the OpenMP run-time library: getting and setting stack size for parallel threads and memory allocation.

The Intel extension routines described in this section can be used for low-level debugging to verify that the library code and application are functioning as intended. It is recommended to use these routines with caution because using them requires the use of the `-openmp_stubs` command-line option to execute the program sequentially. These routines are also generally not recognized by other vendor's OpenMP-compliant compilers, which may cause the link stage to fail for these other compilers.

### Stack Size

In most cases, environment variables can be used in place of the extension library routines. For example, the stack size of the parallel threads may be set using the `KMP_STACKSIZE` environment variable rather than the `kmp_set_stacksize()` library routine.

#### Note

A run-time call to an Intel extension routine takes precedence over the corresponding environment variable setting.

The routines `kmp_set_stacksize()` and `kmp_get_stacksize()` take a 32-bit argument only. The routines `kmp_set_stacksize_s()` and `kmp_get_stacksize_s()` take a `size_t` argument, which can hold 64-bit integers.

On Itanium-based systems, it is recommended to always use `kmp_set_stacksize_s()` and `kmp_get_stacksize_s()`. These `_s()` variants must be used if you need to set a stack size  $2^{**}32$  bytes (4 gigabytes).

See the definitions of stack size routines in the table that follows.

## Memory Allocation

The Intel® Fortran Compiler implements a group of memory allocation routines as an extension to the OpenMP\* run-time library to enable threads to allocate memory from a heap local to each thread. These routines are: `kmp_malloc`, `kmp_calloc`, and `kmp_realloc`.

The memory allocated by these routines must also be freed by the `kmp_free` routine. While it is legal for the memory to be allocated by one thread and `kmp_free`'d by a different thread, this mode of operation has a slight performance penalty.

See the definitions of these routines in the table that follows.

Function/Routine	Description
<b>Stack Size</b>	
<pre>function kmp_get_stacksize_s() integer(kind=kmp_size_t_kind)kmp_ge t _stacksize_s</pre>	Returns the number of bytes that will be allocated for each parallel thread to use as its private stack. This value can be changed via the <code>kmp_get_stacksize_s</code> routine, prior to the first parallel region or via the <code>KMP_STACKSIZE</code> environment variable.
<pre>function kmp_get_stacksize() integer kmp_get_stacksize</pre>	This routine is provided for backwards compatibility only; use <code>kmp_get_stacksize_s</code> routine for compatibility across different families of Intel processors.
<pre>subroutine kmp_set_stacksize_s(size) integer (kind=kmp_size_t_kind) size</pre>	Sets to <i>size</i> the number of bytes that will be allocated for each parallel thread to use as its private stack. This value can also be set via the <code>KMP_STACKSIZE</code> environment variable. In order for <code>kmp_set_stacksize_s</code> to have an effect, it must be called before the beginning of

	the first (dynamically executed) parallel region in the program.
<pre>subroutine kmp_set_stacksize(<i>size</i>) integer <i>size</i></pre>	This routine is provided for backward compatibility only; use <code>kmp_set_stacksize_s(<i>size</i>)</code> for compatibility across different families of Intel processors.
<b>Memory Allocation</b>	
<pre>function kmp_malloc(<i>size</i>) integer(kind=kmp_pointer_kind)kmp_m alloc integer(kind=kmp_size_t_kind)<i>size</i></pre>	Allocate memory block of <i>size</i> bytes from thread-local heap.
<pre>function kmp_calloc(<i>nelem</i>,<i>elsize</i>) integer(kind=kmp_pointer_kind)kmp_c alloc integer(kind=kmp_size_t_kind)<i>nelem</i> integer(kind=kmp_size_t_kind)<i>elsize</i></pre>	Allocate array of <i>nelem</i> elements of size <i>elsize</i> from thread-local heap.
<pre>function kmp_realloc(<i>ptr</i>, <i>size</i>) integer(kind=kmp_pointer_kind)kmp_r ealloc integer(kind=kmp_pointer_kind)<i>ptr</i> integer(kind=kmp_size_t_kind)<i>size</i></pre>	Reallocate memory block at address <i>ptr</i> and <i>size</i> bytes from thread-local heap.
<pre>subroutine kmp_free(<i>ptr</i>) integer (kind=kmp_pointer_kind) <i>ptr</i></pre>	Free memory block at address <i>ptr</i> from thread-local heap. Memory must have been previously allocated with <code>kmp_malloc</code> , <code>kmp_calloc</code> , or <code>kmp_realloc</code> .

## Examples of OpenMP Usage

The following examples show how to use the OpenMP feature. See more examples in the OpenMP Fortran version 2.0 specifications.

### do: A Simple Difference Operator

This example shows a simple parallel loop where each iteration contains a different number of instructions. To get good load balancing, dynamic scheduling is used. The `end do` has a `nowait` because there is an implicit barrier at the end of the parallel region.

```
subroutine do_1 (a,b,n)
real a(n,n), b(n,n)
c$omp parallel
```

```

c$omp&   shared(a,b,n)
c$omp&   private(i,j)
c$omp do schedule(dynamic,1)
do i = 2, n
do j = 1, i
b(j,i) = ( a(j,i) + a(j,i-1) ) / 2
enddo
enddo
c$omp end do nowait
c$omp end parallel
end

```

### do: Two Difference Operators

This example shows two parallel regions fused to reduce `fork/join` overhead. The first `end do` has a `nowait` because all the data used in the second loop is different than all the data used in the first loop.

```

      subroutine do_2
      (a,b,c,d,m,n)
      real a(n,n), b(n,n), c(m,m),
      d(m,m)
      c$omp parallel
      c$omp&   shared(a,b,c,d,m,n)
      c$omp&   private(i,j)
      c$omp do schedule(dynamic,1)
      do i = 2, n
      do j = 1, i
      b(j,i) = ( a(j,i) + a(j,i-1) ) / 2
      enddo
      enddo
      c$omp end do nowait
      c$omp do schedule(dynamic,1)
      do i = 2, m
      do j = 1, i
      d(j,i) = ( c(j,i) + c(j,i-1) ) / 2
      enddo
      enddo
      c$omp end do nowait
      c$omp end parallel
      end

```

### sections: Two Difference Operators

This example demonstrates the use of the `sections` directive. The logic is identical to the preceding `do` example, but uses `sections` instead of `do`. Here the speedup is limited to 2 because there are only two units of

work whereas in `do: Two Difference Operators` above there are  $n-1 + m-1$  units of work.

```

subroutine sections_1
(a,b,c,d,m,n)
real a(n,n), b(n,n), c(m,m),
d(m,m)
!$omp parallel
!$omp& shared(a,b,c,d,m,n)
!$omp& private(i,j)
!$omp sections
!$omp section
do i = 2, n
do j = 1, i
b(j,i)=( a(j,i) + a(j,i-1) ) /
2
enddo
enddo

!$omp section
do i = 2, m
do j = 1, i
d(j,i)=( c(j,i) + c(j,i-1) ) /
2
enddo
enddo
!$omp end sections nowait
!$omp end parallel
end

```

### single: Updating a Shared Scalar

This example demonstrates how to use a `single` construct to update an element of the shared array `a`. The optional `nowait` after the first loop is omitted because it is necessary to wait at the end of the loop before proceeding into the `single` construct.

```

subroutine sp_1a
(a,b,n)
real a(n), b(n)
!$omp parallel
!$omp& shared(a,b,n)
!$omp& private(i)
!$omp do
do i = 1, n
a(i) = 1.0 / a(i)
enddo
!$omp single
a(1) = min( a(1), 1.0
)
!$omp end single
!$omp do
do i = 1, n
b(i) = b(i) / a(i)
enddo

```



```
!$omp end do nowait  
!$omp end parallel  
end
```

# Debugging Multithreaded Programs

## Debugging Multithread Programs Overview

The debugging of multithreaded program discussed in this section applies to both the OpenMP Fortran API and the Intel Fortran parallel compiler directives. When a program uses parallel decomposition directives, you must take into consideration that the bug might be caused either by an incorrect program statement or it might be caused by an incorrect parallel decomposition directive. In either case, the program to be debugged can be executed by multiple threads simultaneously.

To debug the multithreaded programs, you can use:

- Intel Debugger for IA-32 and Intel Debugger for Itanium-based applications (ldb)
- Intel Fortran Compiler debugging options and methods; in particular, Compiling Source Lines with Debugging Statements.
- Intel parallelization extension routines for low-level debugging.
- VTune(TM) Performance Analyzer to define the problematic areas.

Other best known debugging methods and tips include:

- Correct the program in single-threaded, uni-processor environment
- Statically analyze locks
- Use trace statement (such as `print` statement)
- Think in parallel, make very few assumptions
- Step through your code
- Make sense of threads and callstack information
- Identify the primary thread
- Know what thread you are debugging
  - Single stepping in one thread does not mean single stepping in others
- Watch out for context switch

## Debugger Limitations for Multithread Programs

Debuggers such as Intel Debugger for IA-32 and Intel Debugger for Itanium-based applications support the debugging of programs that are executed by multiple threads. However, the currently available versions of such debuggers do not directly support the debugging of parallel decomposition directives, and therefore, there are limitations on the debugging features.

Some of the new features used in OpenMP are not yet fully supported by the debuggers, so it is important to understand how these features work to know how to debug them. The two problem areas are:

- Multiple entry points
- Shared variables

You can use routine names (for example, `padd`) and entry names (for example, `__PADD`, `__PADD_6__par_loop0`). Fortran Compiler, by default, first mangles lower/mixed case routine names to upper case. For example, `pAdd( )` becomes `PADD( )`, and this becomes entry name by adding one underscore. The secondary entry name mangling happens after that. That's why `"__par_loop"` part of the entry name stays as lower case. Debugger for some reason didn't take the upper case routine name `"PADD"` to set the breakpoint. Instead, it accepted the lower case routine name `"padd"`.

## Debugging Parallel Regions

### Debugging Parallel Regions

The compiler implements a parallel region by enabling the code in the region and putting it into a separate, compiler-created entry point. Although this is different from outlining – the technique employed by other compilers, that is, creating a subroutine, – the same debugging technique can be applied.

### Constructing an Entry-point Name

The compiler-generated parallel region entry point name is constructed with a concatenation of the following strings:

- `"__"` character
- entry point name for the original routine (for example, `__parallel`)
- `"_"` character
- line number of the parallel region
- `__par_region` for OpenMP parallel regions (`!$OMP PARALLEL`)  
`__par_loop` for OpenMP parallel loops (`!$OMP PARALLEL DO`),  
`__par_section` for OpenMP parallel sections (`!$OMP PARALLEL SECTIONS`)
- sequence number of the parallel region (for each source file, sequence number starts from zero.)

### Debugging Code with Parallel Region

Example 1 illustrates the debugging of the code with parallel region. Example 1 is produced by this command:

```
ifort -openmp -g -O0 -S file.f90
```

Let us consider the code of subroutine `parallel` in Example 1.

#### Subroutine `PARALLEL ( )` source listing

```

1  subroutine parallel
2  integer id,OMP_GET_THREAD_NUM
3  !$OMP PARALLEL PRIVATE(id)
4  id = OMP_GET_THREAD_NUM()
5  !$OMP END PARALLEL
6  end

```

The parallel region is at line 3. The compiler created two entry points: `parallel_and__parallel_3__par_region0`. The first entry point corresponds to the subroutine `parallel()`, while the second entry point corresponds to the OpenMP parallel region at line 3.

#### Example 1 Debugging Code with Parallel Region

#### Machine Code Listing of the Subroutine `parallel()`

```

    .globl parallel_
parallel_:
..B1.1:                # Preds ..B1.0
..LN1:
pushl    %ebp                #1.0
movl    %esp, %ebp          #1.0
subl    $44, %esp           #1.0
pushl    %edi                #1.0
... ..
..B1.13:               # Preds ..B1.9
addl    $-12, %esp          #6.0
movl    $.2.1_2_kmpc_loc_struct_pack.2, (%esp) #6.0
movl    $0, 4(%esp)         #6.0
movl    $_parallel__6__par_region1, 8(%esp)   #6.0
call    __kmpc_fork_call    #6.0
        # LOE
..B1.31:               # Preds ..B1.13
addl    $12, %esp           #6.0
        # LOE
..B1.14:               # Preds ..B1.31 ..B1.30
..LN4:
leave                    #9.0
ret                                #9.0
        # LOE
.type parallel_,@function
.size parallel_,.-parallel_
.globl _parallel__3__par_region0
_parallel__3__par_region0:
# parameter 1: 8 + %ebp
# parameter 2: 12 + %ebp
..B1.15:               # Preds ..B1.0
pushl    %ebp                #9.0

```

```

movl    %esp, %ebp                #9.0
subl    $44, %esp                 #9.0
..LN5:
call    omp_get_thread_num_      #4.0
        # LOE eax
..B1.32:                                # Preds ..B1.15
movl    %eax, -32(%ebp)          #4.0
        # LOE
..B1.16:                                # Preds ..B1.32
movl    -32(%ebp), %eax          #4.0
movl    %eax, -20(%ebp)          #4.0
..LN6:
leave                                       #9.0
ret                                         #9.0
        # LOE
.type   _parallel__3__par_region0,@function
.size   _parallel__3__par_region0,.,_parallel__3__par_region0
.globl  _parallel__6__par_region1
_parallel__6__par_region1:
# parameter 1: 8 + %ebp
# parameter 2: 12 + %ebp
..B1.17:                                # Preds ..B1.0
pushl   %ebp                          #9.0
movl    %esp, %ebp                  #9.0
subl    $44, %esp                   #9.0
..LN7:
call    omp_get_thread_num_      #7.0
        # LOE eax
..B1.33:                                # Preds ..B1.17
movl    %eax, -28(%ebp)           #7.0
        # LOE
..B1.18:                                # Preds ..B1.33
movl    -28(%ebp), %eax           #7.0
movl    %eax, -16(%ebp)           #7.0
..LN8:
leave                                       #9.0
ret                                         #9.0
.align  4,0x90
# mark_end;

```

Debugging the program at this level is just like debugging a program that uses POSIX threads directly. Breakpoints can be set in the threaded code just like any other routine. With GNU debugger, breakpoints can be set to source-level routine names (such as `parallel`). Breakpoints can also be set to entry point names (such as `parallel_` and `_parallel__3__par_region0`). Note that Intel Fortran Compiler for Linux converted the upper case Fortran subroutine name to the lower case one.

## Debugging Multiple Threads

When in a debugger, you can switch from one thread to another. Each thread has its own program counter so each thread can be in a different place in the code. Example 2 shows a Fortran subroutine `PADD()`. A breakpoint can be set at the entry point of OpenMP parallel region.

### Source listing of the Subroutine `PADD()`

```

12.      SUBROUTINE PADD(A, B, C, N)
13.      INTEGER N
14.      INTEGER A(N), B(N), C(N)
15.      INTEGER I, ID, OMP_GET_THREAD_NUM
16.      !$OMP PARALLEL DO SHARED (A, B, C, N)
PRIVATE (ID)
17.      DO I = 1, N
18.          ID = OMP_GET_THREAD_NUM()
19.          C(I) = A(I) + B(I) + ID
20.      ENDDO
21.      !$OMP END PARALLEL DO
22.      END

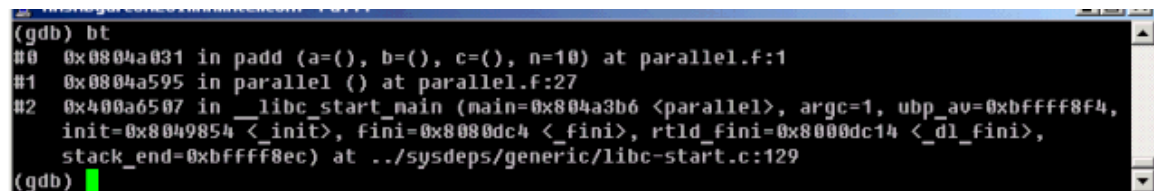
```

## The Call Stack Dumps

The first call stack below is obtained by breaking at the entry to subroutine `PADD` using GNU debugger. At this point, the program has not executed any OpenMP regions, and therefore has only one thread. The call stack shows a system runtime `__libc_start_main` function calling the Fortran main program `parallel()`, and `parallel()` calls subroutine `padd()`. When the program is executed by more than one thread, you can switch from one thread to another. The second and the third call stacks are obtained by breaking at the entry to the parallel region. The call stack of master contains the complete call sequence. At the top of the call stack is `_padd__6__par_loop0()`. Invocation of a threaded entry point involves a layer of Intel OpenMP library function calls (that is, functions with `__kmp` prefix). The call stack of the worker thread contains a partial call sequence that begins with a layer of Intel OpenMP library function calls.

ERRATA: GNU debugger sometimes fails to properly unwind the call stack of the immediate caller of Intel OpenMP library function `__kmpc_fork_call()`.

### Call Stack Dump of Master Thread upon Entry to Subroutine `PADD`



```

(gdb) bt
#0  0x0804a031 in padd (a=(), b=(), c=(), n=10) at parallel.f:1
#1  0x0804a595 in parallel () at parallel.f:27
#2  0x400a6507 in __libc_start_main (main=0x804a3b6 <parallel>, argc=1, ubp_av=0xbffff8f4,
init=0x8049854 <_init>, fini=0x8080dc4 <_fini>, rtdl_fini=0x8000dc14 <_dl_fini>,
stack_end=0xbffff8ec) at ../sysdeps/generic/libc-start.c:129
(gdb)

```

## Switching from One Thread to Another

```

(gdb) info threads
* 4 Thread 2051 (LWP 17512)  0x0804a38a in __padd_6__par_loop0 () at parallel.f:13
  3 Thread 1026 (LWP 17511)  0x40144a31 in __libc_nanosleep () from /lib/i686/libc.so.6
  2 Thread 2049 (LWP 17510)  0x4016f9f7 in __poll (fds=0x80abd5c, nfds=1, timeout=2000)
    at ../sysdeps/unix/sysv/linux/poll.c:63
  1 Thread 1024 (LWP 17493)  0x0804a38a in __padd_6__par_loop0 () at parallel.f:13
(gdb)

```

## Call Stack Dump of Master Thread upon Entry to Parallel Region

```

(gdb) bt
#0  0x0804a38a in __padd_6__par_loop0 () at parallel.f:13
#1  0x080763d9 in __invoke_3 () at proton/libi/getstat.c:241
#2  0x0807b26c in __kmpc_invoke_task_func () at proton/libi/getstat.c:241
(gdb)

```

## Call Stack Dump of Worker Thread upon Entry to Parallel Region

```

(gdb) bt
#0  0x400b8aa5 in __sigsuspend (set=0x40d9e958)
    at ../sysdeps/unix/sysv/linux/sigsuspend.c:45
#1  0x4007e079 in __pthread_wait_for_restart_signal (self=0x40d9ebe0) at pthread.c:967
#2  0x4007abdc in pthread_cond_wait (cond=0x80971b8, mutex=0x8096068) at restart.h:34
#3  0x08075cf2 in __kmp_suspend () at proton/libi/getstat.c:241
#4  0x4007bc7f in pthread_start_thread_event (arg=0x40d9ebe0) at manager.c:298
(gdb)

```

## Example 2 Debugging Code Using Multiple Threads with Shared Variables

### Subroutine PADD() Machine Code Listing

```

        .globl padd_
padd_:
# parameter 1: 8 + %ebp
# parameter 2: 12 + %ebp
# parameter 3: 16 + %ebp
# parameter 4(n): 20 + %ebp
..B1.1:                                # Preds ..B1.0
..LN1:
pushl   %ebp                                #1.0
... ..

..B1.19:                                # Preds ..B1.15
addl   $-28, %esp                            #6.0
movl   $.2.1_2_kmpc_loc_struct_pack.1, (%esp) #6.0
movl   $4, 4(%esp)                            #6.0
movl   $_padd_6__par_loop0, 8(%esp)           #6.0
movl   -196(%ebp), %eax                       #6.0
movl   %eax, 12(%esp)                         #6.0
movl   -152(%ebp), %eax                       #6.0
movl   %eax, 16(%esp)                         #6.0

```

```

movl    -112(%ebp), %eax                #6.0
movl    %eax, 20(%esp)                  #6.0
lea     20(%ebp), %eax                  #6.0
movl    %eax, 24(%esp)                  #6.0
call    __kmpc_fork_call                 #6.0
        # LOE
..B1.39:                                # Preds ..B1.19
addl    $28, %esp                       #6.0
jmp     ..B1.31                          # Prob 100% #6.0
        # LOE
..B1.20:                                # Preds ..B1.30
... ..
call    __kmpc_for_static_init_4        #6.0
        # LOE
..B1.40:                                # Preds ..B1.20
addl    $36, %esp                       #6.0
        # LOE
... ..
..B1.26:                                # Preds ..B1.28 ..B1.21
addl    $-8, %esp                       #6.0
movl    $.2.1_2_kmpc_loc_struct_pack.1, (%esp) #6.0
movl    -8(%ebp), %eax                  #6.0
movl    %eax, 4(%esp)                   #6.0
call    __kmpc_for_static_fini          #6.0
        # LOE
..B1.41:                                # Preds ..B1.26
addl    $8, %esp                        #6.0
jmp     ..B1.31                          # Prob 100% #6.0
        # LOE
..B1.27:                                # Preds ..B1.28 ..B1.25
..LN7:
call    omp_get_thread_num_             #8.0
        # LOE eax
..B1.42:                                # Preds ..B1.27
... ..
cmpl    %edx, %eax                      #10.0
jle     ..B1.27                          # Prob 50% #10.0
jmp     ..B1.26                          # Prob 100% #10.0
        # LOE
.type   padd_,@function
.size   padd_,.-padd_
.globl  _padd__6__par_loop0
_padd__6__par_loop0:
# parameter 1: 8 + %ebp
# parameter 2: 12 + %ebp
# parameter 3: 16 + %ebp
# parameter 4: 20 + %ebp
# parameter 5: 24 + %ebp
# parameter 6: 28 + %ebp
..B1.30:                                # Preds ..B1.0

```



```

..LN16:
pushl    %ebp                                #13.0
movl     %esp, %ebp                          #13.0
subl     $208, %esp                          #13.0
movl     %ebx, -4(%ebp)                      #13.0
..LN17:
movl     8(%ebp), %eax                        #6.0
movl     (%eax), %eax                        #6.0
movl     %eax, -8(%ebp)                      #6.0
movl     28(%ebp), %eax                      #6.0
..LN18:
movl     (%eax), %eax                        #7.0
movl     (%eax), %eax                        #7.0
movl     %eax, -80(%ebp)                    #7.0
movl     $1, -76(%ebp)                      #7.0
movl     -80(%ebp), %eax                    #7.0
testl   %eax, %eax                          #7.0
jg       ..B1.20                            # Prob 50%  #7.0
                                     # LOE
..B1.31:                                # Preds ..B1.41 ..B1.39
..B1.38 ..B1.30
..LN19:
movl     -4(%ebp), %ebx                      #13.0
leave                                       #13.0
ret                                           #13.0
.align   4,0x90
# mark_end;

```

## Debugging Shared Variables

When a variable appears in a `PRIVATE`, `FIRSTPRIVATE`, `LASTPRIVATE`, or `REDUCTION` clause on some block, the variable is made private to the parallel region by redeclaring it in the block. `SHARED` data, however, is not declared in the threaded code. Instead, it gets its declaration at the routine level. At the machine code level, these shared variables become incoming subroutine call arguments to the threaded entry points (such as `___PADD_6__par_loop0`).

In Example 2, the entry point `___PADD_6__par_loop0` has six incoming parameters. The corresponding OpenMP parallel region has four shared variables. First two parameters (parameters 1 and 2) are reserved for the compiler's use, and each of the remaining four parameters corresponds to one shared variable. These four parameters exactly match the last four parameters to `__kmpc_fork_call()` in the machine code of `PADD`.

### Note

The `FIRSTPRIVATE`, `LASTPRIVATE`, and `REDUCTION` variables also require shared variables to get the values into or out of the parallel region.

Due to the lack of support in debuggers, the correspondence between the shared variables (in their original names) and their contents cannot be seen in the debugger at the threaded entry point level. However, you can still move to the call stack of one of the subroutines and examine the contents of the variables at that level. This technique can be used to examine the contents of shared variables. In Example 2, contents of the shared variables A, B, C, and N can be examined if you move to the call stack of `PARALLEL( )`.



# Optimization Support Features

## Optimization Support Features Overview

This section describes the Intel® Fortran features such as directives, intrinsics, run-time library routines and various utilities which enhance your application performance in support of compiler optimizations. These features are Intel Fortran language extensions that enable you optimize your source code directly. This section includes examples of optimizations supported by Intel extended directives and intrinsics or library routines that enhance and/or help analyze performance.

For complete detail of the Intel® Fortran Compiler directives and examples of their use, see Chapter 14, "Directive Enhanced Compilation," in the *Intel® Fortran Language Reference*. For intrinsic procedures, see Chapter 9, "Intrinsic Procedures," in the *Intel® Fortran Language Reference*.

A special topic describes options that enable you to generate optimization reports for major compiler phases and major optimizations. The optimization report capability is used for Itanium®-based applications only.

## Compiler Directives

### Compiler Directives Overview

This section discusses the Intel® Fortran language extended directives that enhance optimizations of application code, such as software pipelining, loop unrolling, prefetching and vectorization. For complete list, descriptions and code examples of the Intel® Fortran Compiler directives, see "Directive Enhanced Compilation" in the *Intel® Fortran Language Reference*.

### Pipelining for Itanium®-based Applications

The `SWP` | `NOSWP` directives indicate preference for a loop to get software-pipelined or not. The `SWP` directive does not help data dependence, but overrides heuristics based on profile counts or lop-sided control flow.

The syntax for this directive is:

```
!DEC$ SWP or CDEC$ SWP
```

```
!DEC$ NOSWP or CDEC$ NOSWP
```

The software pipelining optimization triggered by the `SWP` directive applies instruction scheduling to certain innermost loops, allowing instructions within a loop to be split into different stages, allowing increased instruction level parallelism. This can reduce the impact of long-latency operations, resulting in faster loop execution. Loops chosen for software pipelining are always innermost loops that do not contain procedure calls that are not inlined. Because the optimizer no longer considers fully unrolled loops as innermost loops, fully unrolling loops can allow an additional loop to become the innermost loop (see `-unroll[n]`). You can request and view the optimization report to see whether software pipelining was applied (see Optimizer Report Generation).

<b>SWP</b>
<pre>!DEC\$ SWP do i = 1, m if (a(i) .eq. 0) then b(i) = a(i) + 1 else b(i) = a(i)/c(i) endif enddo</pre>

## Loop Count and Loop Distribution

### LOOP COUNT (N) Directive

The `LOOP COUNT (n)` directive indicates the loop count is likely to be `n`.

The syntax for this directive is:

```
!DEC$ LOOP COUNT(n) or CDEC$ LOOP COUNT(n)
```

where `n` is an integer constant.

The value of loop count affects heuristics used in software pipelining, vectorization and loop-transformations.

<b>LOOP COUNT (N)</b>
<pre>!DEC\$ LOOP COUNT (10000) do i =1,m b(i) = a(i) +1 ! This is likely to enable ! the loop to get software- ! pipelined enddo</pre>

## Loop Distribution Directive

The `DISTRIBUTE POINT` directive indicates to compiler a preference of performing loop distribution.

The syntax for this directive is:

```
!DEC$ DISTRIBUTE POINT or CDEC$ DISTRIBUTE POINT
```

Loop distribution may cause large loops be distributed into smaller ones. This may enable more loops to get software-pipelined. If the directive is placed inside a loop, the distribution is performed after the directive and any loop-carried dependency is ignored. If the directive is placed before a loop, the compiler will determine where to distribute and data dependency is observed. Currently only one distribute directive is supported if it is placed inside the loop.

### DISTRIBUTE POINT

```
!DEC$ DISTRIBUTE POINT
do i =1, m
b(i) = a(i) +1
....
c(i) = a(i) + b(i) ! Compiler will decide
where
                ! to distribute.
                ! Data dependency is observed
....
d(i) = c(i) + 1
enddo

do i =1, m
b(i) = a(i) +1
....
!DEC$ DISTRIBUTE POINT
call sub(a, n)      ! Distribution will start
here,
                ! ignoring all loop-carried
                ! dependency
c(i) = a(i) + b(i)
....
d(i) = c(i) + 1
enddo
```

## Loop Unrolling Support

The `UNROLL` directive tells the compiler how many times to unroll a counted loop.

The syntax for this directive is:

```
CDEC$ UNROLL or !DEC$ UNROLL
```

```
CDEC$ UNROLL [n] or !DEC$ UNROLL [n]
```

```
CDEC$ NOUNROLL or !DEC$ NOUNROLL
```

where *n* is an integer constant. The range of *n* is 0 through 255.

The UNROLL directive must precede the `do` statement for each `do` loop it affects.

If *n* is specified, the optimizer unrolls the loop *n* times. If *n* is omitted or if it is outside the allowed range, the optimizer assigns the number of times to unroll the loop.

The UNROLL directive overrides any setting of loop unrolling from the command line.

Currently, the directive can be applied only for the innermost loop nest. If applied to the outer loop nests, it is ignored. The compiler generates correct code by comparing *n* and the loop count.

UNROLL
<pre>CDEC\$ UNROLL(4) do i = 1, m b(i) = a(i) + 1 d(i) = c(i) + 1 enddo</pre>

## Prefetching Support

The PREFETCH and NOPREFETCH directives assert that the data prefetches be generated or not generated for some memory references. This affects the heuristics used in the compiler.

The syntax for this directive is:

```
CDEC$ PREFETCH or !DEC$ PREFETCH
```

```
CDEC$ NOPREFETCH or !DEC$ NOPREFETCH
```

```
CDEC$ PREFETCH a,b or !DEC$ PREFETCH a,b
```

```
CDEC$ NOPREFETCH a,b or !DEC$ NOPREFETCH a,b
```

If loop includes expression  $a(j)$ , placing `PREFETCH a` in front of the loop, instructs the compiler to insert prefetches for  $a(j + d)$  within the loop.  $d$  is determined by the compiler. This directive is supported when option `-O3` is on.

### PREFETCH

```
CDEC$ NOPREFETCH c
CDEC$ PREFETCH a
do i = 1, m
b(i) = a(c(i)) + 1
enddo
```

## Vectorization Support

The directives discussed in this topic support vectorization.

### IVDEP Directive

#### Syntax:

```
CDEC$ IVDEP
!DEC$ IVDEP
```

The `IVDEP` directive instructs the compiler to ignore assumed vector dependences. To ensure correct code, the compiler treats an assumed dependence as a proven dependence, which prevents vectorization. This directive overrides that decision. Use `IVDEP` only when you know that the assumed loop dependences are safe to ignore.

For example, if the expression  $j \geq 0$  is always true in the code fragment bellow, the `IVDEP` directive can communicate this information to the compiler. This directive informs the compiler that the conservatively assumed loop-carried flow dependences for values  $j < 0$  can be safely ignored:

```
!DEC$ IVDEP
do i = 1,
100
a(i) =
a(i+j)
enddo
```

#### Note

The proven dependences that prevent vectorization are not ignored, only assumed dependences are ignored.

The usage of the directive differs depending on the loop form, see examples below.

<b>Loop 1</b>
<pre>Do i   = a(*) + 1   a(*) = enddo</pre>
<b>Loop 2</b>
<pre>Do i   a(*) =     = a(*) + 1 enddo</pre>

For loops of the form 1, use old values of *a*, and assume that there is no loop-carried flow dependencies from *DEF* to *USE*.

For loops of the form 2, use new values of *a*, and assume that there is no loop-carried anti-dependencies from *USE* to *DEF*.

In both cases, it is valid to distribute the loop, and there is no loop-carried output dependency.

<b>Example 1</b>
<pre>CDEC\$ IVDEP do j=1,n   a(j) = a(j+m) + 1 enddo</pre>
<b>Example 2</b>
<pre>CDEC\$ IVDEP do j=1,n   a(j) = b(j) + 1   b(j) = a(j+m) + 1 enddo</pre>

Example 1 ignores the possible backward dependencies and enables the loop to get software pipelined.

Example 2 shows possible forward and backward dependencies involving array *a* in this loop and creating a dependency cycle. With *IVDEP*, the backward dependencies are ignored.

*IVDEP* has options: *IVDEP:LOOP* and *IVDEP:BACK*. The *IVDEP:LOOP* option implies no loop-carried dependencies. The *IVDEP:BACK* option implies no backward dependencies.

The *IVDEP* directive is also used with the `-ivdep_parallel` option for Itanium®-based applications.



For more details on the `IVDEP` directive, see "Directive Enhanced Compilation," in the *Intel® Fortran Language Reference*.

## Overriding Vectorizer's Efficiency Heuristics

In addition to `IVDEP` directive, there are more directives that can be used to override the efficiency heuristics of the vectorizer:

```
VECTOR ALWAYS
NOVECTOR
VECTOR ALIGNED
VECTOR UNALIGNED
VECTOR NONTEMPORAL
```

The `VECTOR` directives control the vectorization of the subsequent loop in the program, but the compiler does not apply them to nested loops. Each nested loop needs its own directive preceding it. You must place the vector directive before the loop control statement.

### The `VECTOR ALWAYS` and `NOVECTOR` Directives

The `VECTOR ALWAYS` directive overrides the efficiency heuristics of the vectorizer, but it only works if the loop can actually be vectorized, that is: use `IVDEP` to ignore assumed dependences.

#### Syntax:

```
!DEC$ VECTOR ALWAYS
!DEC$ NOVECTOR
```

The `VECTOR ALWAYS` directive can be used to override the default behavior of the compiler in the following situation. Vectorization of non-unit stride references usually does not exhibit any speedup, so the compiler defaults to not vectorizing loops that have a large number of non-unit stride references (compared to the number of unit stride references). The following loop has two references with `stride 2`. Vectorization would be disabled by default, but the directive overrides this behavior.

#### Vector Aligned

```
!DEC$ VECTOR ALWAYS
do i = 1, 100, 2
a(i) = b(i)
enddo
```

If, on the other hand, avoiding vectorization of a loop is desirable (if vectorization results in a performance regression rather than improvement), the `NOVECTOR` directive can be used in the source text to disable vectorization of a loop. For instance, the Intel® Compiler vectorizes the following example loop by default. If

this behavior is not appropriate, the `NOVECTOR` directive can be used, as shown below.

<b>NOVECTOR</b>
<pre>!DEC\$ NOVECTOR do i = 1, 100 a(i) = b(i) + c(i) enddo</pre>

## The `VECTOR ALIGNED` and `UNALIGNED` Directives

### Syntax:

```
!DEC$ VECTOR ALIGNED
!DEC$ VECTOR UNALIGNED
```

Like `VECTOR ALWAYS`, these directives also override the efficiency heuristics. The difference is that the qualifiers `UNALIGNED` and `ALIGNED` instruct the compiler to use, respectively, unaligned and aligned data movement instructions for all array references. This disables all the advanced alignment optimizations of the compiler, such as determining alignment properties from the program context or using dynamic loop peeling to make references aligned.

### Note

The directives `VECTOR [ALWAYS, UNALIGNED, ALIGNED]` should be used with care. Overriding the efficiency heuristics of the compiler should only be done if the programmer is absolutely sure the vectorization will improve performance. Furthermore, instructing the compiler to implement all array references with aligned data movement instructions will cause a run-time exception in case some of the access patterns are actually unaligned.

## The `VECTOR NONTEMPORAL` Directive

**Syntax:** `!DEC$ VECTOR NONTEMPORAL`

The `VECTOR NONTEMPORAL` directive results in streaming stores on Pentium® 4 based systems. A floating-point type loop together with the generated assembly are shown in the example below. For large `n`, significant performance improvements result on a Pentium 4 systems over a non-streaming implementation.

The following example illustrates the use of the `NONTEMPORAL` directive:

<b>NONTEMPORAL</b>
<pre>subroutine set(a,n) integer i,n real a(n)</pre>

```

!DEC$ VECTOR NONTEMPORAL
!DEC$ VECTOR ALIGNED
  do i = 1, n
    a(i) = 1
  enddo
end
program setit
parameter(n=1024*1204)
real a(n)
integer i
do i = 1, n
  a(i) = 0
enddo
call set(a,n)
do i = 1, n
  if (a(i).ne.1) then
    print *, 'failed
nontemp.f', a(i), i
    stop
  endif
enddo
print *, 'passed nontemp.f'
end

```

## Optimizations and Debugging

This topic describes the command-line options that you can use to debug your compilation and to display and check compilation errors.

The options that enable you to get debug information while optimizing are as follows:

-O0	Disables optimizations. Enables <code>-fp</code> option.
-g	Generates symbolic debugging information and line numbers in the object code for use by the source-level debuggers. Turns off <code>-O2</code> and makes <code>-O0</code> the default unless <code>-O2</code> (or <code>-O1</code> or <code>-O3</code> ) is explicitly specified in the command line together with <code>-g</code> .
-fp IA-32 only	Disables the use of the <code>ebp</code> register in optimizations. Directs to use the <code>ebp</code> -based stack frame for all functions.

### Support for Symbolic Debugging, `-g`

Use the `-g` option to direct the compiler to generate code to provide symbolic debugging information and line numbers in the object code that will be used by your source-level debugger. For example:

```
ifort -g prog1.f
```

Turns off `-O2` and makes `-O0` the default unless `-O2` (or `-O1` or `-O3`) is explicitly specified in the command line together with `-g`.

## The Use of `ebp` Register

### `-fp` (IA-32 only)

Most debuggers use the `ebp` register as a stack frame pointer to produce a stack backtrace. The `-fp` option disables the use of the `ebp` register in optimizations and directs the compiler to generate code that maintains and uses `ebp` as a stack frame pointer for all functions so that a debugger can still produce a stack backtrace without turning off `-O1`, `-O2`, or `-O3` optimizations.

Note that using this option reduces the number of available general-purpose registers by one, and results in slightly less efficient code.

### `-fp` Summary

Default	OFF
<code>-O1</code> , <code>-O2</code> , or <code>-O3</code>	Disable <code>-fp</code>
<code>-O0</code>	Enables <code>-fp</code>

## The `-traceback` Option

The `-traceback` option also forces the compiler to use `ebp` as the stack frame pointer. In addition, the `-traceback` option causes the compiler to generate extra information into the object file, which allows a symbolic stack traceback to be produced if a run-time failure occurs.

## Combining Optimization and Debugging

The `-O0` option turns off all optimizations so you can debug your program before any optimization is attempted. To get the debug information, use the `-g` option.

The compiler lets you generate code to support symbolic debugging while one of the `-O1`, `-O2`, or `-O3` optimization options is specified on the command line along with `-g`, which produces symbolic debug information in the object file.

Note that if you specify an `-O1`, `-O2`, or `-O3` option with the `-g` option, some of the debugging information returned may be inaccurate as a side-effect of optimization.

It is best to make your optimization and/or debugging choices explicit:

- If you need to debug your program excluding any optimization effect, use the `-O0` option, which turns off all the optimizations.
- If you need to debug your program with optimization enabled, then you can specify the `-O1`, `-O2`, or `-O3` option on the command line along with `-g`.

### Note

The `-g` option slows down the program when no optimization level (`-On`) is specified. In this case `-g` turns on `-O0`, which is what slows the program down. However, if, for example, both `-O2` and `-g` are specified, the code should run very nearly at the same speed as if `-g` were not specified.

Refer to the table below for the summary of the effects of using the `-g` option with the optimization options.

These options	Produce these results
<code>-g</code>	Debugging information produced, <code>-O0</code> enabled (optimizations disabled), <code>-fp</code> enabled for IA-32-targeted compilations
<code>-g -O1</code>	Debugging information produced, <code>-O1</code> optimizations enabled.
<code>-g -O2</code>	Debugging information produced, <code>-O2</code> optimizations enabled.
<code>-g -O3 -fp</code>	Debugging information produced, <code>-O3</code> optimizations enabled, <code>-fp</code> enabled for IA-32-targeted compilations.

## Debugging and Assembling

The assembly listing file is generated without debugging information, but if you produce an object file, it will contain debugging information. If you link the object file and then use the GDB debugger on it, you will get full symbolic representation.

## Optimizer Report Generation

The Intel® Fortran Compiler provides options to generate and manage optimization reports.

- `-opt_report` generates optimizations report and places it in a file specified in `-opt_report_filename`. If `-opt_report_file` is not specified,

`-opt_report` directs the report to `stderr`. The default is OFF: no reports are generated.

- `-opt_report_filefilename` generates optimizations report and directs it to a file specified in `filename`.
- `-opt_report_level{min/med/max}` specifies the detail level of the optimizations report. The `min` argument provides the minimal summary and the `max` the full report. The default is `-opt_report_levelmin`.
- `-opt_report_routine [substring]` generates reports from all routines with names containing the `substring` as part of their name. If `[substring]` is not specified, reports from all routines are generated. The default is to generate reports for all routines being compiled.

## Specifying Optimizations to Generate Reports

The compiler can generate reports for an optimizer you specify in the `phase` argument of the

`-opt_report_phasephase` option.

The option can be used multiple times on the same command line to generate reports for multiple optimizers.

Currently, the reports for the following optimizers are supported:

Optimizer Logical Name	Optimizer Full Name
ipo	Interprocedural Optimizer
hlo	High-level Language Optimizer
ilo	Intermediate Language Scalar Optimizer
ecg	Itanium Compiler Code Generator
all	All optimizers

When one of the above logical names for optimizers are specified all reports from that optimizer will be generated. For example, `-opt_report_phaseipo` and `-opt_report_phaseecg` generate reports from the interprocedural optimizer and the code generator.

Each of the optimizers can potentially have specific optimizations within them. Each of these optimizations are prefixed with the optimizer's logical name. For example:

Optimizer_optimization	Full Name
ipo_inl	Interprocedural Optimizer, inline expansion of functions
ipo_cp	Interprocedural Optimizer, copy propagation
hlo_unroll	High-level Language Optimizer, loop unrolling
hlo_prefetch	High-level Language Optimizer, prefetching
ilo_copy_propagation	Intermediate Language Scalar Optimizer, copy propagation
ecg_swp	Itanium Compiler Code Generator, software pipelining

### Command Syntax Example

The following command generates a report for the Itanium Compiler Code Generator (*ecg*):

```
ifort -c -opt_report -opt_report_phase ecg myfile.f
```

where:

- `-c` tells the compiler to stop at generating the object code, not linking
- `-opt_report` invokes the report generator
- `-opt_report_phase ecg` indicates the phase (*ecg*) for which to generate the report; the space between the option and the phase is optional.

The entire name for a particular optimization within an optimizer need not be specified in full, just a few characters is sufficient. All optimization reports that have a matching prefix with the specified optimizer are generated. For example, if `-opt_report_phase ilo_co` is specified, a report from both the constant propagation and the copy propagation are generated.

### The Availability of Report Generation

The `-opt_report_help` option lists the logical names of optimizers and optimizations that are currently available for report generation.

**For IA-32 systems**, the reports can be generated for:

- `ilo`
- `hlo` if `-O3` is on
- `ipo` if interprocedural optimizer is invoked with `-ip` or `-ipo`
- all the above optimizers if `-O3` and `-ip` or `-ipo` options are on

**For Itanium-based systems**, the reports can be generated for:

- ilo
- ecg
- hlo if -O3 is on
- ipo if interprocedural optimizer is invoked with -ip or -ipo
- all the above optimizers if -O3 and -ip or -ipo options are on

 **Note**

If hlo or ipo report is requested, but the controlling option (-O3 or -ip--ipo, respectively) is not on, the compiler generates an empty report.



# Index

.il files .....	2	<b>8</b>	80-bit.....	6
[			8-bit.....	6, 25, 128
-[no]altparam compiler option.....	43		8-byte .....	6, 25, 50
-[no]logo compiler option .....	93	<b>A</b>	ABI.....	53
—			ABS .....	130
_s() variants .....	180		absence	
<b>1</b>			loop-carried memory dependency	
128-bit .....	6, 60		.....	116
128-bit Streaming SIMD Extensions			accessing	
.....	128		16-bit.....	25
16-bit			accuracy	
accessing .....	25		controlling .....	65
16-bit .....	6, 128		add.....	4, 66
16-byte .....	28, 43, 131		test2 .....	103
16-byte-aligned .....	123		add/subtract	
1-byte .....	20		operations.....	43, 65,66
<b>2</b>			added performance.....	85
24-bit significand			adheres	
pc32.....	43		ANSI .....	47
<b>3</b>			advanced PGO options.....	90
32-bit			affected aspect of the program.....	73
exceed .....	73		affects	
pointers.....	73		inlining.....	80
single-precision .....	128		SSE.....	60
32-bit .....	6, 25, 89, 180		after	
32k.....	174		FORTRAN 77 .....	4
3-byte .....	6		vectorization .....	128, 131
<b>4</b>			ALIAS .....	93
4-byte .....	6, 20, 50		aliased .....	47
<b>5</b>			-align compiler option.....	6
5000 interval			ALIGNED.....	197
set.....	112		aligning	
5000 interval .....	112		data.....	131
53-bit significand			alignment.....	6
pc64.....	43		options .....	50
<b>6</b>			strategy .....	131
64-bit .....	6, 25, 89, 180		all	
64-bit double-precision .....	128		input/output.....	39
64-bit MMX(TM).....	128		optimizers .....	204
64-bit significand			alloca procedure .....	82
pc80.....	43		ALLOCATABLE.....	47

- allowing
  - optimizer ..... 35
- alternate ..... 43
- ALWAYS vector ..... 197
- analyzer..... 141
- analyzing ..... 6
  - effects of multiple IPO ..... 79
- AND..... 161, 169
- another possible environment
  - variable setting ..... 39
- ANSI ..... 43, 66
  - adheres..... 47
  - conforms..... 60
- ANSI Fortran
  - R818 ..... 151
- ansi\_alias ..... 43, 47
- ANTI ..... 135
- anti-dependencies ..... 197
- API..... 117, 140
- applications
  - features contributes..... 117
- application's ..... 103,141
  - code coverage..... 95
  - tests ..... 95
  - visual presentation ..... 95
- applies
  - ATOMIC..... 161
- ar library manager..... 79
- argument
  - aliasing ..... 131
  - using efficiently..... 15
- arranging
  - data items ..... 6
- array ..... 28, 35, 43, 115, 128, 183
  - accessing ..... 15
  - assumed-shape..... 15
  - compiler creates ..... 15
  - derived-part ..... 6
  - natural storage order..... 20
  - operations..... 130
  - output argument array types..... 15
  - requirements ..... 15
  - using efficiently..... 15
- assembling ..... 201
- assembly files ..... 32, 75, 78, 79, 197
  - specifying ..... 201
- assume compiler option.....20, 35
- assumed-shape arrays ..... 15
- ATOMIC directive..... 141, 161
- ATTRIBUTES C ..... 93
- auto compiler option ..... 47
- automatic allocation
  - variables ..... 47
- automatic compiler option..... 47
- automatic processor-specific
  - optimization ..... 71
- AUTOMATIC statement..... 47
- auto-parallelization
  - data flow ..... 117, 135, 138
  - diagnostic ..... 138
  - enabling ..... 136
  - environment variables ..... 136
  - overview..... 134
  - processing ..... 135
  - programming with..... 135
  - threshold control ..... 138
  - threshold needed..... 136
- auto-parallelized
  - loops ..... 138
- auto-parallelized..... 43, 138
- auto-parallelizer's
  - control ..... 117, 134 -138
  - enabling ..... 117, 136
  - threshold ..... 138
- auto-vectorization.....2, 28, 117
- auto-vectorizer ..... 124
- availability
  - report generation ..... 204
- avoid
  - EQUIVALENCE ..... 6
  - mixed data type arithmetic
    - expressions..... 25
    - small integer items ..... 25
    - small logical items ..... 25
  - unaligned data ..... 6
  - vectorization ..... 197
- ax{K|W|N|B|P} compiler option .... 28, 57, 121
- B**
- BACK option of IVDEP ..... 197
- BACKSPACE ..... 20
- backtrace compiler option..... 201

- BARRIER directive
  - executes ..... 141
  - use ..... 161
- basic PGO options
  - profile-guided optimization ..... 86
- bcolor option of code-coverage tool
  - ..... 95
- before
  - inserting ..... 141
  - vectorization ..... 128
- begin
  - parallel construct ..... 147
  - serial execution ..... 147
  - worksharing construct ..... 147
- best performance
  - function ..... 72
- big-endian ..... 35
  - little-endian ..... 39
- binding ..... 141
- bitwise AND ..... 130
- block size ..... 128
- BLOCKSIZE
  - increasing ..... 20
  - omitting ..... 20
  - values ..... 20
- bound
  - denormalized single-precision ... 60
- Bourne shell ..... 32
- browsing
  - frames ..... 95
- BUFFERCOUNT
  - buffered\_io option ..... 20
  - default ..... 20
  - increase ..... 20
- buffers
  - UBC ..... 20
- byterecl keyword ..... 20
- C**
- c compiler option ..... 35, 75
- c\$omp ..... 149, 183
- c\$OMP BARRIER ..... 161
- c\$OMP DO PRIVATE ..... 161
- c\$OMP END PARALLEL ..... 161
- c\$OMP PARALLEL ..... 161
- cache size intrinsic ..... 28
- cachsize ..... 28
- call stack dumps
  - master thread ..... 189
  - worker thread ..... 189
- call WORK ..... 156, 158, 161, 169
- callee ..... 82
- calls
  - malloc ..... 53
  - OMP\_SET\_NUM\_THREADS .. 156
- callstack ..... 185
- causing
  - unaligned data ..... 6
- cc\_omp keyword ..... 43
- ccdefault compiler option ..... 43
- ccolor option of code coverage tool
  - ..... 95
- CDEC\$ prefix of directives.. 136, 194, 195, 196, 197
- CEIL rounding mode ..... 28
- changing
  - number ..... 156
- character data ..... 6, 50
- characteristics ..... 57, 73, 141
- checking
  - floating-point stack state ..... 47
  - inefficient unaligned data ..... 6
- choosing
  - inline ..... 25
- chunk size ..... 172
- clause containing reduction ..... 169
- clauses ..... 141, 147, 149, 151 - 171, 193
  - comma-separated list ..... 156, 158
  - cross-reference ..... 151
  - list ..... 166
  - summary ..... 151
- cleanup ..... 128
- code
  - assembly ..... 32, 197
  - preparing ..... 141
- Code DO ..... 25
- codecov\_option ..... 95
- code-coverage tool ..... 95
- coding ..... 4, 25
  - Intel® architectures ..... 28
- coloring scheme
  - setting ..... 95

combined parallel/worksharing	
constructs .....	141, 160
command line	
options .....	50
syntax .....	95, 103, 149, 193
comma-separated .....	161
comma-separated list ..	151, 161, 166
clauses .....	156, 158
variables .....	151
COMMON	
FIELDS .....	166
block ...	6, 28, 35, 50, 53, 141, 156, 165-167
statement.....	6, 28, 47, 50
Compaq* Visual Fortran .....	2
compilation	
controlling .....	50
customizing process of .....	35
efficient .....	35
optimizing .....	35
options .....	43, 47, 50, 53
phase.....	74
techniques .....	35
compiler	
applying heuristic .....	138
commands.....	35
compiler's IL .....	78
compiler-created .....	186
compiler-generated .....	28, 186
compiler-supplied library .....	84
compiling with OpenMP* .....	149
creating array descriptor .....	15
creating temporary array .....	15
default optimizations .....	43
defining the size of the array	
elements .....	6
directives .....	194
efficient compilation.....	35
Intel(R) extension routines .....	180, 186
IPO benefits.....	73
issuing warnings.....	78, 89, 90, 91, 121, 151
merging the data from all .dyn files .....	91
optimization levels.....	57
producing	
pgopti.dpi file.....	92
producing .....	78
producing .....	92
programming with OpenMP* ....	141
relocating the source files .....	94
report generation .....	204
selecting routines for inlining.....	82
treating assumed dependence	197
vectorization	
support .....	197
vectorization .....	121
compiling source lines .....	185
COMPLEX.....	6, 25, 43, 47
conditional parallel region execution	
setting .....	156
conditional parallel region execution .....	156
conforming	
ANSI .....	60
IEEE 754 .....	66
constructing	
entry-point name .....	186
constructing.....	186
containing	
8-byte.....	6
IR .....	75, 76
substring .....	204
CONTINUE .....	126
controlling	
accuracy .....	65
advanced PGO optimizations ....	90
alignment with options.....	50
auto-parallelizer's diagnostic levels .....	117, 138
compilation process .....	35
complex flow.....	117
computation of stacks and	
variables.....	47
data scope attributes .....	158
floating-point accuracy .....	66
floating-point computations .....	65
generation of profile information .....	110
inline expansion.....	84
loop vectorization .....	197

number of threads ..... 156  
 OpenMP\* diagnostics ..... 149  
 rounding..... 64  
 speculation ..... 57  
 your program with OpenMP\* ... 177  
 conventions ..... 4  
 converting  
   little-endian ..... 35  
 COPYIN clause... 147, 151, 156, 166  
 copyprivate..... 151  
 correct usage  
   countable ..... 127  
 COS..... 128, 130  
 COUNT ..... 127  
 countable  
   correct usage..... 127  
 coverage analysis  
   modules subset..... 95  
 coverage analysis ..... 95  
 CPU  
   CPU\_TIME ..... 32  
   CPU-intensive ..... 20, 32  
   use ..... 15  
 CRAY\* pointer  
   preventing aliasing ..... 47  
 creating  
   DPI list ..... 103  
   multifile IPO executable using xild  
     ..... 76  
   multifile IPO executable with  
     command line ..... 75  
   multithreaded applications ..... 28  
 criteria  
   inline function expansion..... 82  
 criteria ..... 82  
 criteria for inline function expansion  
   ..... 82  
 CRITICAL directive ..... 141, 151, 161  
   use ..... 161  
 critical/ordered ..... 141  
 cross-iteration ..... 141  
 cross-platform, -ansi\_alias..... 47  
 Csh ..... 39  
 customizing  
   compilation process ..... 35  
 CVF options ..... 2

**D**

data  
   alignment ..... 6, 50  
   alignment example ..... 131  
   data-dependence ..... 115  
   declarations  
     ordering ..... 6  
   declarations ..... 6  
   dependence.... 114, 124, 135, 138,  
     194  
   dependence analysis ..... 124  
   dependence vectorization patterns  
     ..... 124  
   environment directive ..... 147  
   flow  
     program's loops ..... 134  
   flow ..... 117  
   flow ..... 134  
   items  
     arranging..... 6  
   items ..... 6  
   options ..... 47  
   settings ..... 43  
   sharing ..... 141  
   type ..... 4, 6, 25, 60, 117, 121, 169  
 data scope attribute clauses ..... 166  
 data type arithmetic expressions ... 25  
 DATE\_AND\_TIME ..... 32  
 DAZ flag ..... 2, 28, 72  
 DCLOCK ..... 32  
 dcommons keyword ..... 6, 35, 50  
 DCU (data cache unit) ..... 131  
 debugger limitations  
   multithread programs ..... 185  
 debugging  
   code ..... 186  
   multiple threads ..... 189  
   multithread programs overview 185  
   parallel regions ..... 186  
   shared variables ..... 193  
   statements ..... 185  
   code using multiple threads .... 189  
 DEC ..... 93, 136, 194, 195, 196, 197  
 DEF  
   USE ..... 197  
 DEFAULT

BUFFERCOUNT .....	20	enhanced compilation .....	193, 194, 197
disabling options .....	43	format.....	136, 149
Itanium®-based applications .....	82	IVDEP	
default behavior		informs compiler .....	197
compiler options .....	43	IVDEP .....	116
DEFAULT Clause		name.....	136, 149
specify .....	167	overview.....	194
use .....	167	preceding.....	197
deferred-shape.....	15	relieve .....	117
demangle option of the code		usage rules .....	141
coverage tool.....	95	use .....	149
denormal		VECTOR.....	197
exceptions .....	28	directory	
flushing .....	60, 65	specifying.....	90
values .....	28, 60, 65	directory.....	90
denormalized .....	60	disable	
denormalized single-precision		-fp.....	201
bound.....	60	function splitting.....	89
denormalized single-precision.....	60	inlining.....	43
denormalized values.....	60	intrinsic inlining .....	57
denormals .....	28	IPO.....	73
denormals-are-zero .....	28	-On optimizations .....	57
dependence .....	135	disclaimer .....	1
DEQUEUE .....	161	disk I/O .....	20
dequeuing .....	161	dispatch options .....	68
derived-type data .....	6	DISTRIBUTE POINT directive.....	195
describes		division-to-multiplication optimization	
characteristics .....	57	.....	64
Profile IGS .....	110	DO directive .....	141, 158, 161, 168
developing		DO loop .....	15, 20, 25, 158, 168, 172
multithreaded application .....	28	DO WHILE.....	126, 127, 158, 161
device-specific .....	20	document number .....	1
diagnostic reports .....	138, 149	DO-ENDDO.....	126
diagnostics . 117, 121, 123, 136, 138, 149		DOUBLE.....	28
difference operators.....	183	DOUBLE PRECISION	
different applications		returns.....	177
optimizing .....	57	types .....	130
differential coverage		variables	
running.....	95	KIND.....	43
source.....	95	variables .....	43
DIMENSION.....	126, 127	-double_size {n} compiler option ...	43
dimension-by-dimension.....	124	double_size 64 .....	43
directive		dpi	
controls .....	147	customer.dpi .....	95
		file .....	86, 89, 93, 95, 103

- DPI list
  - Create ..... 103
  - dpi\_list file ..... 103
  - dpi\_list tests\_list ..... 103
  - line ..... 103
  - options ..... 95
  - pgopti.dpi ..... 95
- dps ..... 43
- dps compiler option ..... 43
- dummy argument ..... 15, 20, 35
- dummy\_aliases ..... 35, 47
- dumping
  - profile data ..... 93
  - profile information ..... 111, 112
- during
  - instrumentation ..... 93
  - interprocedural optimizations ..... 84
- dyn (dynamic-information files)
  - files .... 86, 90, 92, 93, 95, 103-112
- dynamic
  - counters ..... 95
  - DYNAMIC ..... 161, 172
  - dynamic\_threads ..... 177
  - dynamic-information
    - files ..... 86, 89
  - profile counters
    - resetting ..... 112
- E**
- eax ..... 186, 189
- ebp register ..... 43, 186, 189
  - use ..... 201
- ebx ..... 189
- ecg ..... 204
- ecg\_swp ..... 204
- EDB ..... 6, 32
- edi ..... 186
- edx ..... 189
- effective auto-parallelization usage
  - ..... 135
- effects
  - analyzing ..... 79
  - multiple IPO ..... 79
- efficiency ..... 4, 25
- efficient
  - code ..... 25
  - compilation ..... 35
  - use of arrays ..... 15
    - record buffers ..... 20
- efficient compilation ..... 35
- elapsed time ..... 103
- elsize ..... 180
- email ..... 95
- enable
  - auto-parallelizer ..... 117, 136
  - DEC ..... 43
  - denormals-as-zero ..... 28
  - fp option ..... 60, 201
  - implied-DO loop collapsing ..... 20
  - inlining ..... 84
  - O2 optimizations ..... 57
  - parallelizer ..... 117
  - SIMD-encodings ..... 128
  - test-prioritization ..... 103
- encounters
  - SINGLE ..... 158
- end ..... 126, 127, 131
  - DO ..... 158
  - parallel construct ..... 147
  - REDUCTION ..... 169
  - worksharing
    - construct ..... 147
    - worksharing ..... 141
    - worksharing ..... 151
- END CRITICAL directive ..... 161
- END DO directive ..... 141, 158
- END INTERFACE ..... 93
- END MASTER directive ..... 151, 161
- END ORDERED directive ..... 151, 161
- END PARALLEL
  - directive ..... 147, 156
- END PARALLEL ..... 141
- END PARALLEL ..... 151
- END PARALLEL DO
  - directive ..... 160
- END PARALLEL DO ..... 141, 151
- END PARALLEL DO ..... 160
- END PARALLEL SECTIONS
  - directive ..... 160
- END PARALLEL SECTIONS ..... 141, 151
- END PARALLEL SECTIONS ..... 160
- END SECTION directive ..... 158

END SECTIONS	
directive .....	141, 158
END SECTIONS .....	151
END SECTIONS .....	158
END SINGLE	
directive .....	141, 158
END SINGLE .....	151
END SINGLE .....	158
END SUBROUTINE .....	93, 131
endian .....	39
Enhanced Debugger .....	32
ensuring natural alignment .....	6
entry	
parallel region .....	189
subroutine PADD .....	189
entry/exit .....	135
entry-point name	
constructing .....	186
environment	
data environment directive .....	147
OpenMP environments routines	
.....	177
uniprocessor .....	185
variables .....	20, 39, 91, 110, 136, 149, 156, 174, 177, 180
environment .....	147
EQUIVALENCE statement ..	6, 20, 47
avoid .....	6
EQV .....	161, 169
ERRATA .....	189
errno variable	
setting .....	84
error_limit 30 .....	43
-error_limit n compiler option .....	43
esp .....	186, 189
examples	
OpenMP .....	183
PGO .....	92
vectorization .....	131
examples 4, 6, 15, 20, 25, 28, 32, 35,	
39, 43, 47, 50, 53, 57, 60, 64, 65,	
66, 68, 69, 71, 72, 75, 76, 78, 79,	
80, 84, 85, 86	
exceed	
32-bit .....	73
EXCEPTION list .....	39
executable files .....	35
executing	
BARRIER .....	141
SINGLE .....	158
test-prioritization .....	103
execution	
environment routines .....	177
flow .....	147
existing	
pgopti.dpi .....	91
exit	
worksharing .....	141
explicit symbol visibility specification	
.....	53
explicit-shape arrays .....	15
EXTENDED PRECISION .....	66
extended-precision .....	25
extensions support .....	68
EXTERN symbol visibility attribute	
value .....	53
<b>F</b>	
F_UFMTENDIAN variable	
setting .....	39
value .....	39
-fast compiler option .....	57
fcolor .....	95
feature .....	1, 2, 4, 6, 28, 35
application .....	117
display .....	95
enable .....	39
OpenMP contains .....	141
overview .....	193
work .....	185
feedback compilation .....	92
FIELDS .....	165, 166
file	
.dpi .....	89, 95, 103
.dyn files .....	93
assembly .....	76, 78, 79, 201
containing	
intermediate representation (IR)	
.....	75
list .....	103
containing .....	35, 103
default output .....	35
dynamic-information .....	89



- executable .. 32, 35, 71, 75, 76, 78, 86, 89, 92, 121, 136, 140, 141
- input.....20
- multiple IPO.....74
- multiple source files.....35
- name
  - pgopti.dpi .....93
- name.....4, 90
- object 2, 35, 43, 74, 75, 76, 78, 79, 89
- pathname .....53
- real object files .....78
- relocating the source files .....94
- required .....75, 95, 103
- specifying symbol files .....53
- FIRSTPRIVATE clause .....135, 141, 147, 151, 156, 158, 166, 167, 168, 193
- floating-point
  - applications
    - optimizing.....28
  - applications .....28
  - arithmetic precision
    - IA-32 systems.....64
    - Itanium-based systems .....65
    - mp option.....60
    - mp1 option.....60
    - options .....60
    - overview.....60
  - exceptions .....28
  - handling .....60
  - floating-point-to-integer .....64
  - multiply and add (FA).....66
  - stack state checking.....47
  - type .....60
- FLOW .....135
- FLUSH directive
  - use .....161
- FLUSH directive.....141, 151
- FLUSH directive.....161
- flushing
  - denormal.....60, 65
  - zero denormal .....60, 65
- FMA.....66
- fnsplit- compiler option .....89
- FOR\_SET\_FPE intrinsic
  - FOR\_M\_ABRUPT\_UND .....72
  - FOR\_SET\_FPE intrinsic .....60
  - FOR\_SET\_FPE intrinsic .....72
- fork/join.....183
- format
  - auto-parallelization directives...136
  - big-endian .....39
  - expressions .....20
  - floating-point applications .....28
  - OpenMP directives .....149
- formatted files
  - unformatted files.....20
- FORT\_BUFFERED
  - run-time environment variable ...20
- FORT\_BUFFERED.....20
- Fortran
  - API .....141, 147, 185
  - FORTTRAN 77
    - dummy aliases .....35
    - FORTTRAN 77 .....4, 6, 15
    - FORTTRAN 77 .....35
    - Fortran standard .....4
    - Fortran uninitialized .....53
    - Fortran USE statement .....177
  - INCLUDE statement.....177
- Fourier-Motzkin .....124
- FP .....43,60, 201
  - multiply.....65
  - operations evaluation .....65
  - options .....60
  - results .....65
- fp compiler option
  - fp summary.....201
- fpstkchk compiler option .....2
- frames
  - browsing .....95
- ftz compiler option.....28
- FTZ flag .....2, 28
  - Itanium®-based systems .....60
  - setting .....72
- FTZ flag .....2, 28, 65
- full name.....204
- function
  - best performance .....72
  - function splitting
    - disabling.....89

function/routine.....	180	Hyper-Threading technology	28, 117, 140
function/subroutine.....	47		
<b>G</b>		<b>I</b>	
-g compiler option .....	201	I/O	
GCC .....	76	list.....	20
ld .....	75	parsing .....	20
GCD .....	124	performance	
GDB		improving .....	4, 20
use .....	201	performance .....	4
general-purpose registers.....	201	performance .....	20
generating		IA-32	
instrumented code.....	89	applications.....	113
non-SSE .....	28	floating-point arithmetic .....	64
processor-specific function version		Hyper-Threading Technology-	
.....	71	enabled .....	117
profile-optimized executable .....	89	Intel® Debugger .....	185
reports .....	204	Intel® Enhanced Debugger.....	32
vectorization reports.....	121	IA-32 only .....	4, 60, 115
gigabytes.....	174, 180	IA-32 systems .....	2, 64, 68, 72
global symbols .....	53	IA-32-based	
GNU .....	186, 189	little-endian .....	39
GOT (global offset table) .....	53	processors .....	39, 141
GP-relative .....	53	IA-32-specific feature.....	121
GUIDED (schedule type) .....	172	IA-32-targeted compilations.....	201
guidelines .....	6, 25	IAND .....	151, 161, 169
advanced PGO.....	90	identifying	
auto-parallelization .....	135	synchronization .....	161
coding .....	28	identifying .....	161
vectorization .....	123	IEEE .....	60,72
<b>H</b>		IEEE 754 .....	28, 66
help		conform .....	66
od utility .....	39	IEOR.....	151, 161, 169
HIDDEN visibility attribute .....	53	IF	
high performance		generated .....	95
programming .....	6	statement.....	95
high-level.....	4, 35	IF clause.....	156
optimizer .....	204	-iface compiler option.....	43
parallelization .....	135	ifort..	4, 35, 43, 50, 53, 60, 65, 68, 69, 71, 72, 74, 75, 79, 80, 84, 86, 92, 103, 115, 121, 136, 138, 149, 186, 201, 204
HLO		IL	
hlo_prefetch.....	204	compiler reads .....	78
hlo_unroll.....	204	files .....	2, 78
overview .....	113	produced.....	78
prefetching.....	117	ilo .....	204
unrolling .....	115		
HTML files .....	95		

- ILP ..... 117
- implied DO loop ..... 25
  - collapsing ..... 20
- improving
  - I/O performance ..... 20
  - run-time performance..... 25
- improving/restricting FP arithmetic
  - precision ..... 66
- include ..... 177
  - floating-point-to-integer ..... 64
  - Intel® Xeon(TM) ..... 2
- incorrect usage ..... 126
  - non-countable loop..... 127
- increase
  - BLOCKSIZE specifier ..... 20
  - BUFFERCOUNT specifier ..... 20
- individual module source view..... 95
- industry-standard ..... 140
- inefficient
  - code ..... 25
  - unaligned data..... 6
- infinity ..... 60
- init routine..... 82
- initialization ..... 169
- initializer ..... 53
- initiating
  - interval profile dumping ..... 112
- inlinable ..... 82
- inline
  - choose ..... 25
  - expansion ..... 60, 82, 204
    - controlling ..... 84
    - library functions ..... 84
  - function expansion
    - criteria ..... 82
  - function expansion ..... 82
- inline\_debug\_info compiler option 84
- inlined
  - library..... 84
  - source position ..... 84
- inlined ..... 25, 35, 53, 82
- inlining ..... 43, 57, 79, 85
  - affect..... 80
  - intrinsic..... 57
  - prevents..... 35
- INPUT ..... 158, 161
- arguments..... 15
  - files ..... 20
  - input/output..... 39
  - test-prioritization ..... 103
- instruction-level ..... 117
- instrumentation..... 93, 110
  - compilation ..... 86, 92
  - compilation/execution..... 89
  - repeat..... 90
- instrumented
  - code generating..... 89
  - execution—run ..... 92
  - program ..... 85
- INTEGER ..... 6, 15, 124-127, 189
  - variables ..... 25
- integer\_size{n} compiler option
  - integer\_size 32 ..... 43
- integer\_size{n} compiler option ... 43
- Intel®
  - architecture-based..... 141
  - architecture-based processors .28, 32
  - architecture-specific ..... 32
  - Fortran Compiler for 32-bit
    - application..... 2
  - Fortran Compiler for Itanium®-
    - based applications ..... 2
- Intel® architectures
  - coding ..... 4, 28
- Intel® Compiler
  - adjust ..... 80
  - coding ..... 20, 25, 28
  - directives ..... 194
  - refer ..... 60, 117
  - run..... 149
  - use ..... 6, 32, 76, 79
  - utilize..... 6
  - vectorizes ..... 197
- Intel® Debugger ..... 32
  - IA-32 applications..... 185
  - Itanium®-based applications.... 185
- Intel® Enhanced Debugger
  - IA-32 ..... 32
- Intel® extensions
  - extended intrinsics..... 28
  - OpenMP\* routines ..... 180

Intel® Fortran language ..4, 6, 15, 20, 28, 50, 72, 185, 193	
record structures .....	6
RTL .....	20
Intel® Itanium® Compiler .....	28
Intel® Itanium® processor ...	2, 43, 68
Intel® Pentium® 4 processor .	69, 71, 72
Intel® Pentium® III processor	69, 71, 72
Intel® Pentium® M processor ..	2, 68, 69, 71, 72, 130
Intel® Pentium® processors.....	2, 43, 68, 69, 71, 72, 85, 115, 117
Intel® processors....	2, 28, 68, 69, 72, 180
depending.....	71
optimizing for.....	68, 69, 71, 72
Intel® Threading Toolset .....	28, 32
Intel® VTune(TM) Performance Analyzer .....	32
Intel® Xeon(TM) processors ....	2, 43, 68, 85, 115, 117, 130
Intel®-specific .....	28, 140
INTERFACE.....	93
intermediate language scalar optimizer .....	204
intermediate results use memory.....	20
internal subprograms .....	25
INTERNAL visibility attribute .....	53
interprocedural during.....	84
use .....	28
interprocedural optimizations (IPO) .....	4, 35, 57, 73
compilation with real object files	78
criteria for inline function expansion .....	82
disable .....	73
inline expansion of user functions .....	84
library of IPO objects.....	79
multiple IPO executable .....	76
objects .....	79
options .....	28, 73, 74, 76
-ipo_c .....	79
-ipo_obj .....	78
-ipo_S .....	79
-Option specifies.....	80
interprocedural optimizer .....	73, 204
interthread .....	141
interval profile dumping initiating.....	112
intrinsic .....	4, 15, 28, 35, 43, 57, 84, 193
cachesize.....	28
functions .....	161
inlining.....	57
procedures.....	193
invoking GCC ld .....	76
invoking .....	76
IOR .....	151, 161, 169
-ip compiler option	60, 73, 80, 82, 84, 92, 204
ip_ninl_max_total_stats .....	80
ip_ninl_min_stats .....	80, 82
-ip_no_inlining compiler option	43, 84
-ip_no_pinlining compiler option ...	84
ip_specifier .....	80
-IPF_ft_eval_method{0 2} compiler option .....	65
-IPF_ftacc compiler option .....	65
-IPF_fma compiler option.....	65
-IPF_fp_speculation compiler option .....	65
IPO .....	204
-ipo compiler option.....	73
-ipo_c compiler option.....	79
-ipo_obj compiler option....	43, 78, 82, 121
-ipo_S compiler option	79, 75, 76, 79
IR .....	74, containing .....
object file .....	74
ISYNC .....	161
Itanium® architectures .....	28
Itanium® compiler	28, 43, 53, 60, 65, 80, 89, 115, 174
-auto_ilp32 compiler option.....	73
code generator .....	204

- Itanium® processors.....4, 28, 43, 68
- Itanium®-based applications
  - pipelining ..... 194
- Itanium®-based applications ..... 113, 194
- Itanium®-based compilation..... 86
- Itanium®-based multiprocessor .. 117
- Itanium®-based processors ..... 60
- Itanium®-based systems
  - default..... 82
  - Intel® Debugger ..... 185
  - optimization reports..... 204
  - pipelining ..... 194
  - software pipelining ..... 117
  - using intrinsics..... 28
- IVDEP directive..... 113, 116, 197
- ivdep\_parallel..... 116
- ivdep\_parallel compiler option... 113, 116, 197
- K**
- K|W|N|B|P .....57, 68, 69, 71
- KIND parameter..... 6, 25
  - double-precision variables ..... 43
  - specifying ..... 6
- kmp..... 174, 189
- KMP\_ALL\_THREADS ..... 174
- KMP\_BLOCKTIME ..... 174
- KMP\_BLOCKTIME value ..... 173
- kmp\_calloc ..... 180
- kmp\_free ..... 180
- kmp\_get\_stacksize ..... 180
- kmp\_get\_stacksize\_s ..... 180
- KMP\_LIBRARY..... 174
- kmp\_malloc..... 180
- KMP\_MONITOR\_STACKSIZE... 174
- kmp\_pointer\_kind ..... 180
- kmp\_realloc..... 180
- kmp\_set\_stacksize ..... 180
- kmp\_set\_stacksize\_s
  - order ..... 180
- kmp\_size\_t\_kind ..... 180
- KMP\_STACKSIZE ..... 174, 180
- KMP\_VERSION ..... 174
- kmpc\_for\_static\_fini..... 189
- kmpc\_for\_static\_init\_4..... 189
- kmpc\_fork\_call..... 186, 189, 193
- L**
- LASTPRIVATE
  - clauses..... 168
  - use ..... 168
- LASTPRIVATE... 141, 147, 151, 158, 166, 167, 168, 193
- layer..... 189
- ld 76, 92
- legal information..... 1
- level coverage ..... 95
- libc.so ..... 53
- libc\_start\_main ..... 189
- libdir ..... 43
- libdir keyword compiler option ..... 43
- libguide.a ..... 173
- libirc.a library ..... 92
- libraries
  - functions ..... 84
  - inline expansion..... 84
  - libintrins.a ..... 28
  - library I/O ..... 20
  - OpenMP runtime routines ..... 177
  - routines ..... 177
- limitations
  - loop unrolling ..... 115
- line
  - DPI list ..... 103
  - dpi\_list..... 103
  - lines compiled..... 138
- LINK\_command line..... 76
- linkage phase ..... 74
- list
  - tool generates..... 95
  - tool provides ..... 95
- listing ..... 25, 103, 166
  - file containing..... 103
  - xild..... 76
- little-endian
  - big-endian..... 39
  - converting ..... 35
- little-endian-to-big-endian conversion
  - environment variable..... 39
- Lock routines..... 177
- LOGICAL ..... 6, 47
- loop
  - blocking..... 128

body .....	130	matrix multiplication .....	133
changing .....	15	MAX.....	128, 130, 161, 169
collapsing .....	20	maximum number .	43, 115, 174, 177
computing .....	66	memory	
constructs .....	126	access.....	28
count.....	195	allocation .....	180
diagnostics .....	121, 138	dependency .....	116
directives .....	195	layout .....	28
distribution .....	195	MIN .....	80, 128, 130, 151, 161, 169,
exit conditions.....	127	183, 204	
interchange.....	133	min med max.....	204
LOOP option of IVDEP directive		minimizing	
.....	197	execution time .....	103
parallelization .....	117, 123	number.....	103
parallelizer .....	68	mintime option.....	103
parallelizing .....	68, 141	misaligned	
peeling .....	131, 197	data crossing 16-byte boundary	
sectioning .....	128	.....	131
skewing.....	114	mispredicted.....	86
transformations .....	66, 114, 195	mixing	
types vectorized .....	128	vectorizable .....	123
unrolling ...	57, 113, 115, 123, 194,	MM_PREFETCH.....	117
204		MMX(TM) technology.....	117
limitations.....	115	MODE.....	39
support.....	196	modules subset	
variable assignment.....	168	coverage analysis.....	95
vectorization .....	123, 197	modules subset.....	95
vectorized types .....	128	more replicated code .....	147
loop-carried memory dependency		-mp compiler option .....	60
absence .....	116	-mp1 compiler option .....	60
loop-carried memory dependency		multidimensional arrays .....	15, 124
.....	116	multifile .....	74
lower/mixed .....	185	multifile IPO	
<b>M</b>		IPO executable .....	75, 76
machine code listing		overview.....	74
subroutine.....	186	phases .....	74
maddr option .....	95	stores .....	74
maintainability .....	25	xild.....	76
makefile.....	76	multifile optimization.....	73
malloc		multiple threads	
calls .....	53	debugging .....	189
MASTER directive .....	141, 161, 189	multithread programs	
master thread.....	141, 151, 161	debugger limitations .....	185
call stack dump .....	189	overview.....	185
use .....	161	multithreaded 28, 117, 134-136, 140,	
math libraries .....	84	149, 185	

- applications
  - creating ..... 28
  - developing ..... 28
  - produces ..... 140, 141
  - run ..... 149
- mutually-exclusive
  - part ..... 43
- mutually-exclusive ..... 43
- N**
- names
  - optimizers ..... 204
- NAN value ..... 47, 65
- natural storage order ..... 20
- naturally aligned
  - data ..... 6
  - records ..... 6
  - reordered data ..... 6
- new optimizations ..... 2
- noalign compiler option ..... 50
- noalignments keyword ..... 6
- noauto compiler option ..... 47
- noauto\_scalar compiler option .... 47
- noautomatic compiler option ..... 47
- nobuffered\_io keyword ..... 20
- nocommons keyword ..... 50
- nodcommons keyword ..... 50
- nolib\_inline compiler option .... 60, 84
- nologo compiler option ..... 93
- non-countable loop
  - incorrect usage ..... 127
- NONE ..... 167
- noniterative worksharing SECTIONS
  - use ..... 158
- non-OpenMP ..... 173
- non-preemptable ..... 53
- non-SSE
  - generating ..... 28
- NONTEMPORAL ..... 197
- non-varying values ..... 25
- non-vectorizable loop ..... 123
- non-vectorized loops ..... 121
- NOP ..... 115
- NOPARALLEL directive ..... 135, 136
- nopartial option ..... 95
- NOPREFTCH directives ..... 197
- nosave compiler option ..... 47
- nosequence keyword ..... 50
- NOSWP directives ..... 194
- nototal ..... 103
- NOUNROLL ..... 196
- NOVECTOR directives ..... 197
- NOWAIT option ..... 158
- nozero compiler option ..... 47
- NUM ..... 103, 117
- num\_threads ..... 151, 177
- number ..... 60
  - changing ..... 156
  - minimizing ..... 103
- O**
- O compiler option ..... 57
- o filename compiler option .... 75, 79
- O0 compiler option ..... 57, 60, 201
- O1 compiler option ..... 57
- O2 compiler option .... 25, 35, 43, 50, 57, 60, 65, 92, 113, 115, 121, 136, 149, 201
  - O2 optimizations ..... 57, 201
- O3 compiler option 28, 89, 113, 121, 201
  - optimizations ..... 57, 60, 65, 201
- Ob{n} compiler option
  - Ob0 ..... 84
  - Ob1 ..... 43, 84
  - Ob2 ..... 43, 84
- object files ..... 35, 43, 75, 76, 79, 201
  - IR ..... 74
- od utility
  - help ..... 39
- omitting
  - BLOCKSIZE ..... 20
  - SEQUENCE ..... 6
- OMP ... 117, 141, 147, 149, 167, 174, 183
  - OMP ATOMIC ..... 161
  - OMP BARRIER ..... 158, 161
  - OMP CRITICAL ..... 161
  - OMP DO ..... 147, 156
  - OMP DO LASTPRIVATE ..... 168
  - OMP DO ORDERED,SCHEDULE ..... 161
  - OMP DO REDUCTION ..... 169
  - OMP END CRITICAL ..... 161

OMP END DO .....	156	omp_lock_kind .....	177
OMP END DO directives .....	156	omp_lock_t .....	177
OMP END MASTER .....	161	omp_nest_lock_kind .....	177
OMP END ORDERED .....	161	omp_nest_lock_t .....	177
OMP END PARALLEL 156, 158, 161, 168, 171, 186		OMP_NESTED .....	174
OMP END PARALLEL DO . 156, 160, 169, 189		OMP_NUM_THREADS .....	136, 149, 156, 174
OMP END PARALLEL SECTIONS .....	160	OMP_SCHEDULE .....	136, 141, 172, 174
OMP END SECTIONS .....	158	omp_set_dynamic .....	177
OMP END SINGLE .....	158	omp_set_lock .....	177
OMP FLUSH .....	161	omp_set_nest_lock .....	177
OMP MASTER .....	161	omp_set_nested .....	177
OMP ORDERED .....	161	omp_set_num_threads .....	156, 177
OMP PARALLEL 156, 158, 168, 186		omp_test_lock .....	177
OMP PARALLEL DEFAULT .....	156, 158, 161, 166, 171	omp_test_nest_lock .....	177
OMP PARALLEL DO ... 156, 160, 186		omp_unset_lock .....	177
OMP PARALLEL DO DEFAULT 167, 169		omp_unset_nest_lock .....	177
OMP PARALLEL DO SHARED... 189		-On compiler option .....	57
OMP PARALLEL IF .....	156	one thread .....	189
OMP PARALLEL PRIVATE. 168, 186		open statement	
OMP PARALLEL SECTIONS..... 160, 186		OPEN statement BUFFERED....	20
OMP SECTION .....	158, 160	-openmp compiler option .....	117, 149
OMP SECTIONS .....	158	OpenMP* .. 2, 4, 43, 47, 68, 117, 123, 134, 135, 140	
OMP SINGLE .....	158	clauses .....	151
OMP THREADPRIVATE .....	165, 166	contains	
omp_destroy_lock .....	177	feature .....	141
omp_destroy_nest_lock .....	177	contains .....	141
OMP_DYNAMIC .....	174	directives .....	151
omp_get_dynamic .....	177	environment variables .....	174
omp_get_max_threads .....	177	examples .....	183
omp_get_nested .....	177	extension environment variables .....	174
omp_get_num_procs .....	156, 177	Intel® extensions .....	180
omp_get_num_threads .....	171, 177	par_loop .....	186
omp_get_thread_num 161, 171, 177, 186, 189		par_region .....	186
omp_get_wtick .....	177	par_section .....	186
omp_get_wtime .....	177	parallelizer's	
omp_in_parallel .....	177	option controls .....	149
omp_init_lock .....	177	parallelizer's .....	149
omp_init_nest_lock .....	177	processing .....	141
omp_lib.mod file .....	177	run-time library routines .....	177
		synchronization directives .....	141
		usage .....	183
		uses .....	141



- OpenMP\*-compliant compilers.... 180
- openmp\_report{n} compiler option
  - openmp\_report0..... 149
  - openmp\_report1 ..... 43, 149
  - openmp\_report2..... 149
- openmp\_report{n} compiler option
  - ..... 117, 149
- openmp\_stubs compiler option . 117, 180
- operator/intrinsic ..... 169
- operator|intrinsic ..... 151
- opt\_report{n} compiler option
  - opt\_report\_file ..... 204
  - opt\_report\_filefilename ..... 204
  - opt\_report\_help ..... 204
  - opt\_report\_level ..... 204
  - opt\_report\_levelmin ..... 43, 204
  - opt\_report\_phasephase option
    - ..... 204
  - opt\_report\_routine..... 204
- optima record
  - use ..... 20
- optimization-level
  - options ..... 57
  - restricting ..... 60
  - setting ..... 57
- optimizations
  - compilation process ..... 6
  - debugging and optimizations ... 201
  - different application types ..... 4
  - floating-point arithmetic precision
    - ..... 60
  - HLO ..... 113
  - IPO..... 73
  - optimizer report generation..... 204
  - optimizing for specific processors
    - ..... 68
  - overview ..... 35
  - PGO..... 85
  - reports .....43, 193, 194, 204
- optimizer
  - allowing..... 35
  - full name ..... 204
  - logical name ..... 204
  - report generation ..... 204
  - reports ..... 204
- optimizer..... 35
- optimizer..... 194
- optimizer..... 196
- optimizer..... 204
- optimizers
  - names ..... 204
  - your code ..... 71
- optimizers ..... 204
- optimizing (see also optimizations)
  - application types..... 57
  - floating-point applications ..... 28
  - for specific processors ..... 2, 4, 68
- optimizing (see also optimizations)
  - ..... 28, 57, 68, 130
- option..... 57
  - causes ..... 57
  - controls
    - auto-parallelizer's..... 138
    - OpenMP parallelizer's ..... 149
  - controls ..... 90
  - controls ..... 138
  - controls ..... 149
  - disables..... 89
  - forces ..... 47
  - initializes ..... 47
  - places ..... 47
  - reduces ..... 201
  - sets
    - threshold ..... 138
    - visibility..... 53
  - sets ..... 53
  - sets ..... 57
  - sets ..... 138
- options
  - correspond..... 53
  - debugging summary..... 201
  - direct
    - compiler ..... 47, 65, 71, 121
  - enable ..... 28
    - auto-parallelizer ..... 136
  - improve run-time performance... 35
  - instruct ..... 123
  - output summary..... 201
  - overviews..... 117, 201
- OR ..... 95, 130, 161, 169
- ORDERED ..... 141, 151, 161

- clause ..... 158
- directive ..... 141, 158, 161
- specify ..... 161
- use ..... 161
- ordering
  - data declarations ..... 6
  - kmp\_set\_stacksize\_s ..... 180
- original serial code ..... 134
- other
  - operations ..... 130
  - options ..... 121
  - READ/WRITE statements ..... 39
  - tools ..... 80
- output20, 35, 47, 103, 135, 140, 158, 161, 197
  - argument ..... 15
- overriding
  - vectorizer's efficiency heuristics ..... 197
- overview ..... 6, 35, 57, 68, 73, 74, 85, 110, 113, 117, 121, 140, 185, 193
- P**
- PADD
  - using GNU ..... 185, 189, 193
- par\_report{n} compiler option
  - par\_report Output ..... 138
  - par\_report0 ..... 138
  - par\_report1 ..... 43, 117, 138
  - par\_report2 ..... 138
  - par\_report3 ..... 138
- par\_report{n} compiler option ..... 117
- par\_report{n} compiler option ..... 136
- par\_report{n} compiler option ..... 138
- par\_threshold{n} compiler option
  - par\_threshold0 ..... 138
  - par\_threshold100 ..... 138
- par\_threshold{n} compiler option ..... 117, 136
- par\_threshold{n} compiler option 138
- PARALLEL. 135, 136, 141, 147, 151, 156, 160, 166, 167, 168, 169, 171, 193
- parallel construct
  - begin ..... 147
  - end ..... 147
- parallel construct ..... 147
- PARALLEL directive ... 136, 156, 161
- PARALLEL DO
  - directive ..... 135, 172
  - use ..... 160
- PARALLEL DO... 117, 141, 151, 160, 161, 166-169, 171
- parallel invocations with makefile .76, 89
- PARALLEL PRIVATE ..... 117
- parallel processing ..... 140
  - directive groups ..... 141
  - thread model
    - pseudo code ..... 147
    - thread model ..... 147
- parallel program development ..... 117
- parallel regions ..... 141, 149
  - debugging ..... 186
  - directives ..... 156
  - entry ..... 189
- PARALLEL SECTIONS ..... 141, 151, 161, 167-169, 171
  - use ..... 160
- parallel/worksharing ..... 141, 160
- parallelism ..... 117
- parallelization 123, 134-136, 138, 141
  - loops ..... 117
  - parallelized . 43, 135, 147, 149, 185
  - parallelizer ..... 149
    - enables ..... 117
    - relieves ..... 134
- parsing
  - I/O ..... 20
- part
  - mutually-exclusive ..... 43
- pathname ..... 53
- pc{n} compiler option ..... 43, 64
  - pc32 compiler option
    - 24-bit significand ..... 43
  - pc64 ..... 43, 64
    - 53-bit significand ..... 43
  - pc80 ..... 43, 64
    - 64-bit significand ..... 43
- pc{n} compiler option ..... 43, 64
- pcolor ..... 95
- Pentium® 4 processors ..... 68

- Pentium® III processors ..... 68
- Pentium® M processors ..... 68
- performance analysis..... 140
- performance analyzer ..... 32, 185
- performance-critical ..... 95, 173
- performance-related options ..... 35
- performing ..... 20
  - data flow ..... 117, 134
  - I/O ..... 20
- PGO ..... 82, 85, 89-93, 110
  - environment variables..... 91
  - methodology..... 86
  - PGO API..... 93
  - phases ..... 86
  - usage model..... 86
- PGO API support
  - dumping and resetting profile
    - information..... 112
  - dumping profile information .... 111
  - interval profile dumping ..... 112
  - overview ..... 110
  - resetting the dynamic profile
    - counters ..... 112
  - resetting the profile information 112
- pgopti.dpi file ..... 86, 89, 94, 95
  - compiler produces..... 92
  - existing ..... 91
  - remove..... 91
- pgopti.spi..... 86, 95, 103
- PGOPTI\_Prof\_Dump..... 93, 111
- PGOPTI\_Prof\_Dump\_And\_Reset
  - ..... 112
- PGOPTI\_Prof\_Reset ..... 111, 112
- PGOPTI\_Set\_Interval\_Prof\_Dump
  - ..... 112
- pgouser.h ..... 110
- phase1 ..... 141
- phase2 ..... 141
- pipelining ..... 57, 194, 197, 204
  - Itanium®-based applications ... 194
  - optimization ..... 194
- placing
  - PREFETCH ..... 197
- pointer aliasing..... 47
- pointers .... 15, 47, 73, 117, 123, 151, 201
- position-independent code..... 53
- POSIX ..... 186
- prec\_div compiler option..... 64
- preemption
  - preemptable..... 53
  - preempted ..... 53, 82
- PREFETCH ..... 57, 113, 117
  - placing ..... 197
- prefetching ..... 57, 113, 194, 204
  - optimizations..... 117
  - option ..... 117
  - support..... 197
- preparing
  - code ..... 141
- preventing
  - CRAY\* pointers ..... 47
  - inlining..... 35
- PRINT statement ..... 95, 168
- prioritization ..... 103
- PRIVATE clause 135, 147, 151, 156, 161, 166-169, 171, 189, 193
- private scoping
  - variable ..... 141
- procedure names ..... 151
- process
  - overview..... 35
- process\_data..... 111
- processor..... 28, 68, 69, 71
  - processor-based..... 68
  - processor-instruction ..... 68
  - processor-specific
    - generating ..... 71
    - optimization..... 69, 71, 72
    - runtime checks..... 72
  - processor-specific ..... 2, 53
  - processor-specific ..... 71
  - processor-specific ..... 72
  - targeting ..... 68
- produced ..... 78, 89, 140, 141
  - IL 78
    - multithreaded..... 140, 141
    - profile-optimized ..... 89
- prof\_dir dirname compiler option.. 90
- prof\_dpi file..... 103
- prof\_dpi Test1.dpi ..... 103
- prof\_dpi Test2.dpi ..... 103

prof_dpi Test3.dpi .....	103	pushl .....	186, 189
PROF_DUMP_INTERVAL ....	91, 110	<b>Q</b>	
-prof_file filename compiler option	90	-qipo_fa xild option .....	76
-prof_gen[x] compiler option		-qipo_fo xild option .....	76
-prof_gen compilations.....	89	-Qoption compiler option.....	80
PROF_NO_CLOBBER .....	91	<b>R</b>	
-prof_use compiler option .....	89	-rcd compiler option .....	64
profile data		READ .....	20, 39, 135
dumping .....	93	READ DATA .....	124
profile IGS		READ/WRITE statements.....	39
describe .....	110	REAL	
environment variable.....	110	REAL DATA.....	124
functions .....	110	REAL .....	6, 25, 43, 47, 66, 128, 130,
variable .....	110	131	
profile information		real object files .....	78
dumping .....	111	REAL*10 variables .....	28
generation support .....	110	REAL*16.....	25
profile-guided optimizations (see		REAL*4.....	25
also PGO).....	92, 93	REAL*8.....	25
instrumented program.....	85	-real_size {n} compiler option	
methodology .....	86	-real_size 64 .....	43
overview .....	85	reassociation .....	65, 66, 169
phases .....	86	rec8byte keyword .....	50
utilities.....	93	RECL	
profile-optimized		value .....	20
executable .....	89	recnbyte keyword .....	50
generating .....	89	recommendations .....	28
produce.....	89	controlling alignment .....	50
profiling summary		record buffers	
specifying .....	90	efficient use of .....	20
profmerge		RECORD statement	
tool .....	93, 103	use .....	6
use .....	94	-recursive compiler option.....	47
utility .....	93	redeclaring .....	193
program		redirected standard .....	20
affected aspect.....	73	REDUCTION.....	147, 151, 156, 166
program's loops		clause .....	169
dataflow .....	134	completed .....	169
programming		end .....	169
high performance .....	6	use .....	169
project makefile.....	76	variables .....	169, 193
PROTECTED .....	53	reduction/induction variable .....	57
providing		ref_dpi_file	
superset.....	168	respect .....	95
pseudo code		relieving	
parallel processing model .....	147	I/O .....	20

- relocating
  - source files ..... 94
  - using profmerge ..... 94
- removing
  - pgopti.dpi ..... 91
- reordering
  - transformations ..... 123
- repeating
  - instrumentation ..... 90
- replicated code ..... 147
- report
  - availability ..... 204
  - generation ..... 204
  - optimizer ..... 204
  - stderr ..... 204
- resetting
  - dynamic profile counters ..... 112
  - profile information ..... 112
- restricting
  - FP arithmetic precision ..... 66
  - optimizations ..... 60
- RESULT ..... 72
- results
  - IPO ..... 85
- RETURN ..... 126, 127, 131
  - double-precision ..... 177
  - return values ..... 47
- REVERSE ..... 168
- rm PROF\_DIR ..... 103
- rounding
  - control ..... 64
  - significand ..... 64
- round-to-nearest ..... 64
- routines ..... 82
  - selecting ..... 82, 180
  - timing ..... 177
- RTL ..... 20
- run
  - differential coverage ..... 95
  - multithreaded ..... 149
  - test prioritization ..... 103
- run-time
  - call ..... 180
  - library routines ..... 177
  - peeling ..... 131
  - performance ..... 35
  - processor-specific checks ..... 2, 72
  - scheduling ..... 136
- S**
  - S compiler option ..... 78
  - safe\_cray\_ptr compiler option ..... 47
  - SAVE statement ..... 47
  - scalar ..... 47, 57, 113, 114, 128, 141, 156, 161, 169, 172, 183, 204
    - clean-up iterations ..... 131
    - replacement ..... 115
    - scalar\_integer\_expression ..... 151
    - scalar\_logical\_expression ..... 151
    - scalar\_rep ..... 115
  - scalar\_rep[-] compiler option ..... 115
  - SCHEDULE ..... 151
    - clause ..... 172
    - specifying ..... 172
    - use ..... 158
  - scoping ..... 166
  - SCRATCH ..... 165, 166
  - screenshot ..... 95
  - SECNDS ..... 32
  - SECTION ..... 141, 151, 158
  - SECTIONS ..... 141, 151, 156
    - directive ..... 158, 160, 168, 169
    - use ..... 158
  - sections\_1 ..... 183
  - selecting
    - routines ..... 82
  - selecting ..... 82
  - SEQUENCE
    - omit ..... 6
    - specify ..... 6
    - statement ..... 6, 50
    - use ..... 6
  - setenv ..... 39
  - setting
    - arguments ..... 6
    - coloring scheme ..... 95
    - conditional parallel region
      - execution ..... 156
    - email ..... 95
    - errno ..... 84
    - F\_UFMTENDIAN variable ..... 39
    - FTZ ..... 72
    - html files ..... 95

integer and floating-point data .....	6
optimization level.....	57
units .....	156
Sh .....	39
SHARED ....	117, 135, 147, 156, 166-169, 193
clause .....	171
debugging.....	193
shared scoping .....	141
shared variables .....	189
updating.....	183
use .....	171
significantd .....	43
round .....	64
SIMD .....	28, 117, 121, 123, 128
SIMD SSE2	
streaming.....	28
SIMD-encodings	
enabling .....	128
simple difference operator .....	183
SIN .....	128, 130
SINGLE .....	141, 151, 156
directive .....	158, 161
encounters.....	158
executing .....	158
use .....	158
single-instruction .....	123
single-precision.....	25, 60
single-statement loops.....	123
single-threaded .....	185
small logical data items .....	25
small_bar.....	28
SMP.....	28, 134, 140
software pipelining .....	117, 194, 195
source ....	2, 4, 6, 25, 35, 53, 60, 193, 197
code .....	149
coding guidelines .....	25
files relocation .....	94
input.....	136, 149
listing .....	186, 189
source position	
inlined .....	84
source position .....	84
view .....	95
specialized code .....	69, 71, 117, 121
specific	
optimizing .....	2, 4, 68
specifying	
8-byte data.....	43
DEFAULT .....	167
directory .....	90
END DO.....	158
KIND .....	6
ORDERED.....	161
profiling summary .....	90
RECL .....	20
schedule .....	172
SEQUENCE .....	6
symbol visibility explicitly.....	53
vectorizer .....	131
visibility without symbol file .....	53
spi	
file .....	95, 103
option .....	103
pgopti.spi .....	95, 103
SQRT .....	156
SSE .....	28, 60, 121, 128
SSE2 .....	28, 121
stacks .....	47
size .....	180
standard	
OpenMP* clauses.....	151
OpenMP* directives.....	151
OpenMP* environment variables	
.....	174
statements 15, 25, 39, 47, 50, 80, 95,	
130, 147, 156, 161, 177, 185, 196	
accessing.....	6
BLOCKSIZE .....	20
BUFFERCOUNT .....	20
BUFFERED .....	20
functions .....	25
STATIC.....	172
STATUS .....	20
stderr	
report .....	204
Stream_LF .....	20
streaming	
SIMD SSE2 .....	28
Streaming SIMD Extensions...28, 32,	
123, 128	

- single-precision ..... 128
- stride-1 ..... 123, 133
  - example ..... 133
- strings ..... 20
- strip-mining ..... 128
- STRUCTURE statements ..... 6, 50
- SUBDOMAIN ..... 171
- subl ..... 186, 189
- subobjects ..... 168
- suboption ..... 35
- subroutine
  - machine code listing ..... 186
  - PADD
    - entry ..... 189
    - source listing ..... 189
  - PADD ..... 189
  - PARALLEL ..... 186
  - PGOPTI\_PROF\_DUMP ..... 93
  - VEC\_COPY ..... 131
  - WORK ..... 161
- subscripts ..... 15, 39, 124, 133
  - array ..... 15
  - loop ..... 133
  - varying ..... 20
- substring ..... 82
  - containing ..... 204
- superset ..... 168
- support ..... 2, 4, 6, 28
  - loop unrolling ..... 196
  - MMX(TM) ..... 28
  - OpenMP\* Libraries ..... 134, 173
  - prefetching ..... 197
  - symbolic debugging ..... 201
  - vectorization ..... 197
  - worksharing ..... 141
- SWP directive ..... 194
- symbol
  - file ..... 53
  - preemption ..... 53
  - visibility attribute options ..... 53
- symbolic debugging ..... 201
- synchronization
  - constructs ..... 161
  - identify ..... 161
  - with my neighbor ..... 161
  - worksharing construct directives ..... 158
- synchronization .. 117, 134, 135, 140, 141, 161, 169
- syntax ..... 136, 149
- SYSTEM\_CLOCK ..... 32
- systems ..... 60
- T**
- table operators/intrinsics ..... 169
- TAN ..... 128
- targeting a processor ..... 68
- terabytes ..... 174
- test prioritization tool
  - Test1
    - Test1.dpi ..... 103
    - Test1.dpi 00 ..... 103
    - Test2.dpi ..... 103
  - Test2
    - adding ..... 103
    - Test2.dpi 00 ..... 103
  - Test3
    - Test3.dpi ..... 103
    - Test3.dpi 00 ..... 103
  - tests\_list file ..... 103
  - tselect command ..... 103
- testl ..... 189
- this release ..... 2
- THREADPRIVATE ..... 147, 151, 156, 166
  - directive ..... 141, 165
  - variables ..... 167
- threads ..... 156
- threshold ..... 28
  - auto-parallelization ..... 136
  - control ..... 138
  - option sets ..... 138
- TIME intrinsic procedure ..... 32
- timeout ..... 82
- timing
  - routines ..... 177
  - your application ..... 6, 32
- tips
  - troubleshooting ..... 138
- TLP ..... 117
- tool ..... 6, 28, 32, 74, 80, 93
  - code coverage

list.....	95	unformatted files .....	20
code coverage.....	95	unformatted I/O .....	20
test prioritization .....	103	uninterruptable .....	141
-tpp{n} compiler option		uniprocessor.....	141, 149, 185
-tpp1 .....	68	units	
-tpp2 .....	43, 68	setting .....	156
-tpp5 .....	68	unpredicatable.....	47
-tpp6 .....	68, 79	unproven distinction	
-tpp7 .....	43, 68	unvectorizable copy .....	131
-traceback compiler option .....	201	unproven distinction .....	131
transformations .....	57, 115, 140	UNROLL directive .....	196
reordering .....	123	-unroll[n] compiler option	
transformed parallel code .....	134	-unroll0.....	43, 115
troubleshooting		-unrolln.....	115
tips .....	138	-unroll[n] compiler option .....	115
TRUNC.....	28	unrolling.....	115, 196
tselect command.....	103	loop .....	115
two-dimensional .....	128	unvectorizable .....	123
array .....	28	unvectorizable copy due to	
type		unproven distinction .....	131
aliasability .....	47	updating	
casting .....	117	shared.....	183
INTEGER .....	47	usage	
padd_,@function .....	189	model.....	86, 103
parallel_,@function .....	186	requirements.....	103
part_dt .....	6	rules .....	76, 158
REAL .....	66	user functions.....	84
TYPE statement .....	6	user@system .....	32
types ... 4, 6, 15, 20, 25, 32, 35, 39,		users' source.....	78
43, 47, 50, 57, 60, 65, 66, 71,		using	
82, 117, 123, 128, 130, 135,		32-bit counters.....	89
136, 151, 156, 158, 161, 168,		advanced PGO .....	90
169, 172, 174, 186, 189, 197		ATOMIC .....	161
<b>U</b>		auto-parallelization .....	135
UBC		BARRIER.....	161
buffers.....	20	COPYIN .....	166
ucolor code-coverage tool option ..	95	CPU .....	15
ULIST .....	39	CRITICAL .....	161
unaligned data .....	6	DEF .....	197
UNALIGNED directives .....	197	DEFAULT .....	167
unary		ebp register.....	201
SQRT.....	128	EDB .....	6
unary .....	130	efficient data types .....	25
unbuffered.....	20	EQUIVALENCE statements.....	25
underflow/overflow .....	47	FIRSTPRIVATE.....	168
undispatched.....	158	FLUSH .....	161



- GDB ..... 201
- GOTO ..... 158
- GP-relative ..... 53
- implied-DO loops ..... 20
- Intel® performance analysis tools  
..... 32
- interprocedural optimizations .... 28,  
73
- intrinsics
  - Itanium®-based systems ..... 28
- intrinsics ..... 28
- ip ..... 73
- IPF\_ftacc ..... 65
- IPO ..... 73, 85
- IVDEP ..... 197
- LASTPRIVATE ..... 168
- MASTER ..... 161
- memory
  - intermediate results ..... 20
- memory ..... 20
- mp ..... 60
- noniterative worksharing
  - SECTIONS ..... 158
- non-SSE instructions ..... 28
- NONTEMPORAL ..... 197
- O3 ..... 35
- optimal record ..... 20
- ORDERED ..... 161
- orphaned directives ..... 147
- par\_report3 ..... 138
- par\_threshold0 ..... 138
- PARALLEL DO ..... 160
- PARALLEL SECTIONS ..... 160
- prec\_div ..... 64
- PRIVATE ..... 168
- profile-guided optimization ..... 92
- profmerge ..... 94
- profmerge utility
  - source relocation ..... 94
- profmerge utility ..... 94
- REAL ..... 25
- REAL variables ..... 28
- RECORD ..... 6
- REDUCTION ..... 169
- SCHEDULE ..... 158
- SECTIONS ..... 158
- SEQUENCE ..... 6
- SHARED ..... 171
- SINGLE ..... 158
- slow arithmetic operators ..... 25
- SSE ..... 28
- this document ..... 4
- THREADPRIVATE directive .... 147
- unbuffered disk writes ..... 20
- unformatted files
  - formatted files ..... 20
- unformatted files ..... 20
- vectorization ..... 28
- VTune(TM) Performance Analyzer  
..... 140, 141
- worksharing ..... 156
- xiar ..... 78
- utilities for PGO ..... 93
- utilize ..... 6
- V**
  - value ..... 2, 4, 6, 20
    - 1E-40 ..... 60
    - infinity ..... 60
    - mixed data type ..... 25
    - NaN ..... 60
    - specified for -src\_old and -  
src\_new ..... 94
    - threshold control ..... 138
    - visibility attributes ..... 53
- variables
  - AUTOMATIC ..... 43
  - automatic allocation ..... 47
  - comma-separated list ..... 151
  - correspond ..... 15
  - existing ..... 151
  - ISYNC ..... 161
  - length ..... 20
  - loop ..... 168
  - PGO environment ..... 91
  - private scoping ..... 141
  - profile IGS ..... 110
  - renaming ..... 57
  - scalars ..... 47
  - setting ..... 6, 180
- VAX\* ..... 50
- vec\_report{n} compiler option
  - vec\_report0 ..... 121

- vec\_report1 ..... 43, 121
- vec\_report2 ..... 121
- vec\_report3 ..... 121
- vec\_report4 ..... 121
- vec\_report5 ..... 121
- vector copy ..... 131
- VECTOR directives
  - VECTOR ALIGNED ..... 197
  - VECTOR ALWAYS ..... 197
  - VECTOR NONTEMPORAL ..... 197
  - VECTOR UNALIGNED ..... 197
- vectorizable ..... 124, 130, 133
  - mixing ..... 123
- vectorization (see also Loop)
  - avoiding ..... 197
  - examples ..... 131
  - key programming guidelines .... 123
  - levels ..... 121
  - loop ..... 197
  - options ..... 113, 121
  - overview ..... 121
  - reports ..... 121
  - support ..... 197
- vectorization (see also Loop) . 28, 66, 85, 117, 121, 131, 195, 197
- vectorize ..... 66, 123, 131
  - loops ..... 85
- vectorized ..... 43, 121, 126, 128, 131, 197
- vectorizer ..... 117, 123, 128, 197
  - efficiency heuristics
    - overriding ..... 197
  - efficiency heuristics ..... 197
  - options ..... 121
- vectorizing compilers ..... 123
- vectorizing loops ..... 197
- version numbers ..... 78
- versioned .il files ..... 2
- view
  - XMM ..... 32
- violation
  - FORTTRAN-77 ..... 35
- visibility
  - specifying ..... 53
  - symbol ..... 53
- visual presentation
  - application's code coverage ..... 95
  - visual presentation ..... 95
- vms compiler option ..... 6, 35
- VMS-related ..... 35
- VOLATILE statement ..... 20
- VTune(TM) Performance Analyzer 1, 32, 185
  - use ..... 140
- W**
- W0 compiler option ..... 6
- wallclock ..... 177
- what's new ..... 2
- whitespace ..... 53
- work ..... 156
  - work/pgopti.dpi file ..... 94
  - work/sources ..... 94
- worker thread
  - call stack dump ..... 189
- WORKSHARE ..... 140
- worksharing 117, 134, 135, 156, 160, 169
  - construct ..... 141, 158
    - begin ..... 147
    - end ..... 147
  - construct directives ..... 158
  - end ..... 141, 151
  - exits ..... 141
  - use ..... 156
- WRITE ..... 20, 39, 135, 161
  - WRITE DATA ..... 124
- write whole arrays ..... 20
- X**
- X\_AXIS ..... 158, 160, 161
- x{K|W|N|B|P} compiler option ..... 69, 121
  - xB ..... 2, 69, 72, 121
  - xK ..... 69, 72
  - xK|W|P ..... 66
  - xP ..... 2, 69, 121
- x86 processors ..... 69
- XFIELD ..... 165, 166
- xiar ..... 78
- xild ..... 79
  - listing ..... 76
  - options
    - ipo\_[no]verbose-asm ..... 76

-ipo_fcode-asm.....	76	YFIELD.....	165, 166
-ipo_fsource-asm.....	76	<b>Z</b>	
-qipo_fa.....	76	Z_AXIS.....	158, 160
-qipo_fo.....	76	zero denormal	
options.....	76	flushing.....	60, 65
tool.....	74	ZFIELD.....	165, 166
<b>XMM</b>		-Zp{n} compiler option	
view.....	32	-Zp16.....	50
<b>XOR</b> .....	130	-Zp8.....	43, 50
<b>Y</b>		-Zp{n} compiler option.....	50
Y_AXIS.....	158, 160, 161		