

Intel® C++ Compiler for Linux* Systems User's Guide

Document Number: 253254-031

Table Of Contents

<i>Intel® C++ Compiler User's Guide</i>	<i>1</i>
Disclaimer and Legal Information	1
Welcome to the Intel® C++ Compiler	1
What's New in This Release	2
Features and Benefits	2
Product Web Site and Support	3
System Requirements	3
FLEXlm* Electronic Licensing	3
Related Publications	4
How to Use This Document	5
<i>Compiler Options Quick Reference</i>	<i>6</i>
New Options	6
Options Quick Reference Guide	9
Compiler Options Cross Reference	30
Default Compiler Options	37
Deprecated and Unsupported Compiler Options	38
<i>Volume I: Building Applications</i>	<i>39</i>
Building Applications from the Command Line	41
Building Applications in Eclipse*	43
Compilation Options	67
Linking	82
Debugging	83
Creating and Using Libraries	87
gcc* Compatibility	95
Language Conformance	104
<i>Volume II: Optimizing Applications</i>	<i>108</i>
Optimization Levels	108
Floating-point Optimizations	110
Optimizing for Specific Processors	114
Interprocedural Optimizations	121
Profile-guided Optimizations	131
High-level Language Optimizations (HLO)	150
Parallel Programming	155
Optimization Support Features	192
<i>Reference</i>	<i>202</i>
Compiler Limits	202
Key Files	203
Diagnostics and Messages	207
Intel Math Library	211
Intel® C++ Intrinsics Reference	243
Intel® C++ Class Libraries	382

Intel® C++ Compiler User's Guide

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 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 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 1996 - 2004.

Welcome to the Intel® C++ Compiler

Welcome to the Intel® C++ Compiler. Before you use the compiler, see System Requirements.

Most Linux* distributions include the GNU* C library, assembler, linker, and others. The Intel C++ Compiler includes the Dinkumware* C++ library. See Libraries Overview.

Please look at the individual sections within each main section of this User's Guide to gain an overview of the topics presented. For the latest information, visit the Intel Web site:

<http://www.intel.com/software/products/compilers/clin/>

See Getting Started for basic information on running the compiler.

What's New in This Release

New features for this version of the Intel® C++ Compiler include:

- New Eclipse IDE integration
- New compiler options
- New predefined macros
- Support for exported templates
- Support for template Instantiation
- Invoking the compiler with `icc` and `icpc`
- New defaults for gcc interoperability options
- Support for thread-local storage
- Support for high-level optimization for C on IA-32
- Support for additional debug information
- Deprecated compiler options

For further information on New Features, see the Release Notes.

Features and Benefits

The Intel® C++ Compiler allows your software to perform best on computers based on the Intel architecture. Using new compiler optimizations, such as profile-guided optimization, prefetch instruction and support for Streaming SIMD Extensions (SSE) and Streaming SIMD Extensions 2 (SSE2), the Intel C++ Compiler provides high performance.

Feature	Benefit
High Performance	Achieve a significant performance gain by using optimizations
Support for Streaming SIMD Extensions	Advantage of Intel microarchitecture
Automatic vectorizer	Advantage of SIMD parallelism in your code achieved automatically
OpenMP* Support	Shared memory parallel programming
Floating-point optimizations	Improved floating-point performance
Data prefetching	Improved performance due to the accelerated data delivery
Interprocedural optimizations	Larger application modules perform better
Profile-guided optimization	Improved performance based on profiling frequently-used functions
Processor dispatch	Taking advantage of the latest Intel architecture features while maintaining object code compatibility with previous generations of Intel® Pentium® processors (for IA-32-based systems only).

Product Web Site and Support

For the latest information about Intel® C++ Compiler, visit:

<http://www.intel.com/software/products/compilers/clin/>

For specific details on the Itanium® architecture, visit the web site at

<http://www.intel.com/software/products/browse/itanium.htm>

System Requirements

IA-32 Processor System Requirements

- A computer based on a Pentium® processor or subsequent IA-32 based processor (Pentium 4 processor recommended).
- 128 MB of RAM (256 MB recommended).
- 100 MB of disk space.

Itanium® Processor System Requirements

- A computer with an Itanium processor.
- 256 MB of RAM.
- 100 MB of disk space.

Software Requirements

See the Release Notes for a complete list of system requirements.

FLEXlm* Electronic Licensing

The Intel® C++ Compiler uses Macrovision's FLEXlm* licensing technology. The compiler requires a valid license file in the `/licenses` directory in the installation path. The default directory is `/opt/intel_cc_80/licenses`. The license files have a `.lic` file extension.

If you require a counted license, see *Using the Intel® License Manager for FLEXlm** (`flex_ug.pdf`).

Related Publications

The following documents provide additional information relevant to the Intel® C++ Compiler:

- ISO/IEC 9989:1990, Programming Languages--C
- ISO/IEC 14882:1998, Programming Languages--C++.
- *The Annotated C++ Reference Manual*, Special Edition, Ellis, Margaret; Stroustrup, Bjarne, Addison Wesley, 1991. Provides information on the C++ programming language.
- *The C++ Programming Language*, 3rd edition, 1997: Addison-Wesley Publishing Company, One Jacob Way, Reading, MA 01867.
- *The C Programming Language*, 2nd edition, Kernighan, Brian W.; Ritchie, Dennis W., Prentice Hall, 1988. Provides information on the K & R definition of the C language.
- *C: A Reference Manual*, 3rd edition, Harbison, Samuel P.; Steele, Guy L., Prentice Hall, 1991. Provides information on the ANSI standard and extensions of the C language.
- *Intel Architecture Software Developer's Manual, Volume 1: Basic Architecture*, Intel Corporation, doc. number 243190.
- *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference Manual*, Intel Corporation, doc. number 243191.
- *Intel Architecture Software Developer's Manual, Volume 3: System Programming*, Intel Corporation, doc. number 243192.
- *Intel® Itanium® Assembler User's Guide*.
- *Intel® Itanium®-based Assembly Language Reference Manual*.
- *Itanium® Architecture Software Developer's Manual Vol. 1: Application Architecture*, Intel Corporation, doc. number 245317-001.
- *Itanium® Architecture Software Developer's Manual Vol. 2: System Architecture*, Intel Corporation, doc. number 245318-001.
- *Itanium® Architecture Software Developer's Manual Vol. 3: Instruction Set Reference*, Intel Corporation, doc. number 245319-001.
- *Itanium® Architecture Software Developer's Manual Vol. 4: Itanium® Processor Programmer's Guide*, Intel Corporation, doc. number 245319-001.
- *Intel Architecture Optimization Manual*, Intel Corporation, doc. number 245127.
- *Intel Processor Identification with the CPUID Instruction*, Intel Corporation, doc. number 241618.
- *Intel Architecture MMX(TM) Technology Programmer's Reference Manual*, Intel Corporation, doc. number 241618.
- *Pentium® Pro Processor Developer's Manual (3-volume Set)*, Intel Corporation, doc. number 242693.
- *Pentium® II Processor Developer's Manual*, Intel Corporation, doc. number 243502-001.
- *Pentium® Processor Specification Update*, Intel Corporation, doc. number 242480.
- *Pentium® Processor Family Developer's Manual*, Intel Corporation, doc. numbers 241428-005.

Most Intel documents are also available from the Intel Corporation Web site at <http://developer.intel.com/software/products/>

How to Use This Document

This User's Guide explains how to use the Intel® C++ Compiler. It provides information on how to get started with the Intel C++ Compiler, how this compiler operates, and what capabilities it offers for high performance. You learn how to use the standard and advanced compiler optimizations to gain maximum performance of your application. This documentation assumes that you are familiar with the C++ standard programming language and with the Intel processor architecture. You should also be familiar with the host computer's operating system.



Note

This document explains how information and instructions apply differently to each targeted architecture. If there is no specific indication to either architecture, the description is applicable to both architectures.

Notation Conventions

Style	Definition
<i>This type style</i>	indicates an element of syntax, a reserved word, a keyword, a file name, or part of a program example (text appears in lowercase unless UPPERCASE is required)
This type style	indicates what you type as input
<i>This type style</i>	indicates an argument on a command line or an option's argument
[<i>items</i>]	indicates that the items enclosed in brackets are optional
{ <i>item</i> <i>item</i> }	indicates a set of choices from which you must select one
... (ellipses)	indicates that an argument can be repeated several times

Compiler Options Quick Reference

Conventions Used in the Options Quick Reference Tables

Convention	Definition
[-]	If an option includes "[-]" as part of the definition, then the option can be used to enable or disable the feature. For example, the <code>-c99 [-]</code> option can be used as <code>-c99</code> (enable c99 support) or <code>-c99-</code> (disable c99 support).
[n]	Indicates that the value <code>n</code> in [] can be omitted or have various values.
Values in { } with vertical bars	Used for option's version; for example, option <code>-x{K W N B P}</code> has these versions: <code>-xK</code> , <code>-xW</code> , <code>-xN</code> , <code>-xB</code> , and <code>-xP</code> .
{n}	Indicates that option must include one of the fixed values for <code>n</code> .
Words in <i>this style</i> following an option	Indicate option's required argument(s). Arguments are separated by comma if more than one are required.

New Options

Some compiler options are only available on certain systems. In the following table, these options are indicated with labels as follows:

Label	Meaning
i32	Option available on IA-32-based systems
i32em	Option available on Intel® Extended Memory 64 Technology (Intel® EM64T) systems
i64	Option available on Itanium®-based systems

- If no label is present, the option is available on all supported systems
- If "only" appears in the label, that option is only available on the identified system

The compiler options listed in the following table are new to this release.

Option	Description	Default
<code>-cxxlib-gcc=GCC-root-dir</code>	Specifies the top-level location of the gcc binaries and libraries.	OFF
<code>-debug [no]inline_debug_info</code>	Produces enhanced source position information for inlined code.	OFF
<code>-debug [no]variable_locations</code>	Produces additional debug information for scalar local variables using a feature of the DWARF object module format known as location lists.	OFF
<code>-debug extended</code>	Turns on the three <code>-debug</code> options: <ul style="list-style-type: none"> • <code>-debug inline_info</code> • <code>-debug variable_locations</code> 	OFF
<code>-export</code>	Enable recognition of exported templates. Supported in C++ mode only.	OFF
<code>-export_dir dir</code>	Specifies a directory name for the exported template search path.	OFF
<code>-fabi-version</code>	Directs the compiler to select a specific ABI implementation.	OFF
<code>-finline-functions</code>	Inline any function at the compiler's discretion. Same as <code>-ip</code> .	OFF
<code>-fno-exceptions</code>	The <code>-fno-exceptions</code> option turns off exception handling table generation, resulting in smaller code. Any use of exception handling constructs - try blocks, throw statements will produce an error. Exception specifications are parsed but ignored. A preprocessor symbol <code>__EXCEPTIONS</code> is defined when this option is not used. It is undefined when this option is present.	OFF
<code>-fno-implicit-inline-templates</code>	Do not emit code for implicit instantiations of inline templates. For C++ only.	OFF

Option	Description	Default
<code>-fno-implicit-templates</code>	Never emit code for non-inline templates which are instantiated implicitly (i.e. by use); only emit code for explicit instantiations. For C++ only.	OFF
<code>-ftls-model=<i>model</i></code>	Change thread-local storage model, where <i>model</i> can be the following: <ul style="list-style-type: none"> • <code>global-dynamic</code> • <code>local-dynamic</code> • <code>initial-exec</code> • <code>local-exec</code> 	OFF
<code>-g0</code>	Disable generation of symbolic debug information.	OFF
<code>-[no-]global-hoist</code>	Enables [disables] hoisting and speculative loads of global variables.	OFF
<code>-ipo[<i>value</i>]</code>	Enables interprocedural optimizations across files. The optional <i>value</i> argument controls the maximum number of link-time compilations (or number of object files) that are spawned. The default for <i>value</i> is 1 when <i>value</i> is not specified for small applications. It will generate two or more object files for large applications.	OFF
<code>-ipo_separate</code>	Creates one object file for every source file. This option overrides <code>-ipo[<i>value</i>]</code> .	OFF
<code>-kernel</code> (i64 only)	Generates code for inclusion in the kernel. Prevents generation of speculation as support may not be available when code runs. Suppresses software pipelining.	OFF
<code>-MP</code>	Add a phony target for each dependency.	OFF
<code>-MQ<i>target</i></code>	Same as <code>-MT</code> , but quotes special Make characters.	OFF
<code>-MT<i>target</i></code>	Change the default target rule for dependency generation.	OFF
<code>-Os</code>	Enable speed optimizations, but disable some optimizations which increase code size for small speed benefit.	OFF

Option	Description	Default
<code>-Qlocation,gas,path</code>	Specifies the GNU assembler.	OFF
<code>-Qlocation,gld,path</code>	Specifies the GNU linker.	OFF
<code>-reserve-kernel-regs</code> (i64 only)	Reserves registers f12-f15 and f32-f127 for use by the kernel. These will not be used by the compiler.	OFF
<code>-std=gnu89</code>	ISO C90 plus GNU extensions. Includes some C99 features.	ON
<code>-std=gnu++98</code>	Same as <code>-std=gnu89</code> .	OFF
<code>-[no]traceback</code>	Generate [do not generate] extra information in the object file that allows the display of source file traceback information at run time when a severe error occurs.	OFF

Options Quick Reference Guide

Some compiler options are only available on certain systems. In the following table, these options are indicated with labels as follows:

Label	Meaning
i32	Option available on IA-32-based systems
i32em	Option available on Intel® Extended Memory 64 Technology (Intel® EM64T) systems
i64	Option available on Itanium®-based systems

- If no label is present, the option is available on all supported systems
- If "only" appears in the label, that option is only available on the identified system

Option	Description	Default
<code>-A-</code>	Disables all predefined macros.	OFF
<code>-[no]align</code> (i32 only)	Analyze and reorder memory layout for variables and arrays.	OFF
<code>-Aname[(value)]</code>	Associates a symbol <i>name</i> with the specified sequence of <i>value</i> . Equivalent to an <code>#assert</code> preprocessing directive.	OFF
<code>-alias_args[-]</code>	This option implies arguments may be aliased [not aliased].	<code>-alias_args-</code>

Option	Description	Default
-ansi	Equivalent to GNU* ANSI.	OFF
-ansi_alias[-]	-ansi_alias directs the compiler to assume that the program adheres to the rules defined in the ISO C Standard. If your program adheres to these rules, then this option will allow the compiler to optimize more aggressively. If it doesn't adhere to these rules, then it can cause the compiler to generate incorrect code.	-ansi_alias-
-auto_ilp32	Specifies that the application cannot exceed a 32-bit address space, which allows the compiler to use 32-bit pointers whenever possible. To use this option, you must also specify -ip0[<i>value</i>]. Using the -auto_ilp32 option on programs that can exceed 32-bit address space (2**32) may cause unpredictable results during program execution. This option has no effect on Intel® EM64T systems unless the -axP or -xP option is also used.	OFF
-ax{K W N B P} (i32, i32em)	<p>Generates specialized code for processor-specific codes K, W, N, B, and P while also generating generic IA-32 code.</p> <ul style="list-style-type: none"> • K = Intel® Pentium® III and compatible Intel processors • W = Intel Pentium 4 and compatible Intel processors • N = Intel Pentium 4 and compatible Intel processors • B = Intel Pentium M and compatible Intel processors • P = Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3) <p>Only the -axW and -axP options are available on Intel® EM64T.</p>	OFF

Option	Description	Default
<code>-C</code>	Places comments in preprocessed source output.	OFF
<code>-c</code>	Stops the compilation process after an object file has been generated. The compiler generates an object file for each C or C++ source file or preprocessed source file. Also takes an assembler file and invokes the assembler to generate an object file.	OFF
<code>-c99[-]</code>	Enables [disables] C99 support for C programs.	OFF
<code>-complex_limited_range[-]</code>	Enables the use of "delete basic algebraic expansions" of some arithmetic operations involving data of type <code>_Complex</code> . This can cause some performance improvements in programs that use <code>_Complex</code> arithmetic, but values at the extremes of the exponent range may not compute correctly. Default is <code>-complex_limited_range-</code> .	OFF
<code>-create_pch filename</code>	Manual creation of precompiled header (<code>filename.pchi</code>).	OFF
<code>-cxxlib-gcc[=GCC-root-dir]</code>	Link using C++ run-time libraries provided with <code>gcc</code> . This option is ON by default if your <code>gcc</code> version is 3.2, 3.3, or 3.4. Use the optional argument, <code>=GCC-root-dir</code> , to specify the top-level location of the <code>gcc</code> binaries and libraries.	OFF
<code>-cxxlib-icc</code>	Link using C++ run-time libraries provided by Intel. This option is ON by default if your <code>gcc</code> version is less than 3.2.	OFF
<code>-debug [no]inline_debug_info</code>	Produces enhanced source position information for inlined code.	OFF
<code>-debug [no]variable_locations</code>	Produces additional debug information for scalar local variables using a feature of the DWARF object module format known as location lists.	OFF

Option	Description	Default
-debug extended	Turns on the three -debug options: <ul style="list-style-type: none"> -debug inline_info -debug variable_locations 	OFF
-dM	Output macro definitions in effect after preprocessing (use with -E).	OFF
-Dname [=value]	Defines a macro <i>name</i> and associates it with the specified <i>value</i> . Equivalent to a #define preprocessor directive.	OFF
-dryrun	Show driver tool commands, but do not execute tools.	OFF
-dynamic-linkerfilename	Selects a dynamic linker (<i>filename</i>) other than the default.	OFF
-E	Stops the compilation process after the C or C++ source files have been preprocessed, and writes the results to stdout.	OFF
-EP	Preprocess to stdout omitting #line directives.	OFF
-export	Enable recognition of exported templates. Supported in C++ mode only.	OFF
-export_dir dir	Specifies a directory name for the exported template search path.	OFF
-falias	Assume aliasing in program.	ON
-fabi-version=n	Directs the compiler to select a specific ABI implementation.	OFF

Option	Description	Default
<code>-fast</code>	The <code>-fast</code> option maximizes speed across the entire program. For Itanium-based systems, <code>-fast</code> sets <code>-O3</code> , <code>-ipo</code> , and <code>-static</code> . For IA-32 and Intel® EM64T systems, <code>-fast</code> sets <code>-O3</code> , <code>-ipo</code> , <code>-static</code> , and <code>-xP</code> . Note that on IA-32 and Intel® EM64T systems, programs compiled with the <code>-xP</code> option will detect non-compatible processors and generate an error message during execution.	OFF
<code>-fcode-asm</code>	Produce assembly file with optional code annotations. Requires <code>-S</code> .	OFF
<code>-ffnalias</code>	Assume aliasing within functions.	ON
<code>-finline-functions</code>	Inline any function at the compiler's discretion. Same as <code>-ip</code> .	OFF
<code>-fminshared</code>	Compilation is for the main executable. Absolute addressing can be used and non-position independent code generated for symbols that are at least protected.	OFF
<code>-fno-alias</code>	Assume no aliasing in program.	OFF
<code>-fno-common</code>	Enables the compiler to treat common variables as if they were defined, allowing the use of <code>gp_{rel}</code> addressing of common data variables.	OFF
<code>-fno-exceptions</code>	The <code>-fno-exceptions</code> option turns off exception handling table generation, resulting in smaller code. Any use of exception handling constructs - try blocks, throw statements will produce an error. Exception specifications are parsed but ignored. A preprocessor symbol <code>__EXCEPTIONS</code> is defined when this option is not used. It is undefined when this option is present.	OFF

Option	Description	Default
<code>-fno-fnalias</code>	Assume no aliasing within functions, but assume aliasing across calls.	OFF
<code>-fno-implicit-inline-templates</code>	Do not emit code for implicit instantiations of inline templates. For C++ only.	OFF
<code>-fno-implicit-templates</code>	Never emit code for non-inline templates which are instantiated implicitly (i.e. by use); only emit code for explicit instantiations. For C++ only.	OFF
<code>-f[no-]rtti</code> (i32 and i64)	Enable [disable] RTTI support.	<code>-frtti</code>
<code>-fnsplit[-]</code>	Enables [disables] function splitting. Default is ON with <code>-prof_use</code> . To disable function splitting when you use <code>-prof_use</code> , also specify <code>-fnsplit-</code> .	OFF
<code>-fp</code> (i32, i32em)	Disable using the EBP register as general purpose register.	OFF
<code>-fpic, -fPIC</code>	For IA-32, this option generates position independent code. For Itanium-based systems, this option generates code allowing full symbol preemption.	OFF
<code>-fp_port</code> (i32 only)	Round fp results at assignments and casts. Some speed impact.	OFF
<code>-fpstkchk</code> (i32 only)	Generates extra code after every function call to assure the FP stack is in the expected state.	OFF
<code>-fr32</code> (i64 only)	Use only lower 32 floating-point registers.	OFF
<code>-fshort-enums</code>	Allocate as many bytes as needed for enumerated types.	OFF
<code>-fsource-asm</code>	Produce assemblable file with optional code annotations. Requires <code>-S</code> .	OFF
<code>-fsyntax-only</code>	Same as <code>-syntax</code> .	OFF

Option	Description	Default
<code>-ftls-model=model</code>	Change thread-local storage model, where <i>model</i> can be the following: <ul style="list-style-type: none"> • <code>global-dynamic</code> • <code>local-dynamic</code> • <code>initial-exec</code> • <code>local-exec</code> 	OFF
<code>-ftz[-]</code> (i32em, i64)	Flushes denormal results to zero. The option is turned ON with <code>-O3</code> .	OFF
<code>-funsigned-bitfields</code>	Change default bitfield type to unsigned.	OFF
<code>-funsigned-char</code>	Change default char type to unsigned.	OFF
<code>-f[no]verbose-asm</code>	Produce assemblable file with compiler comments. Default: <code>-fverbose-asm</code>	ON
<code>-fvisibility-default=file</code>	Space separated symbols listed in the <i>file</i> argument will get visibility set to default.	OFF
<code>-fvisibility-extern=file</code>	Space separated symbols listed in the <i>file</i> argument will get visibility set to extern.	OFF
<code>-fvisibility-hidden=file</code>	Space separated symbols listed in the <i>file</i> argument will get visibility set to hidden.	OFF
<code>-fvisibility-internal=file</code>	Space separated symbols listed in the <i>file</i> argument will get visibility set to internal.	OFF
<code>-fvisibility-protected=file</code>	Space separated symbols listed in the <i>file</i> argument will get visibility set to protected.	OFF
<code>-fvisibility=[extern default protected hidden internal]</code>	Global symbols (common and defined data and functions) will get the visibility attribute given by default. Symbol visibility attributes explicitly set in the source code or using the symbol visibility attribute file options will override the <code>-fvisibility</code> setting.	OFF

Option	Description	Default
<code>-fwritable-strings</code> (i32 only)	Ensure that string literals are placed in a writable data section.	OFF
<code>-g</code>	Generates symbolic debugging information in the object code for use by source-level debuggers. The <code>-g</code> option changes the default optimization from <code>-O2</code> to <code>-O0</code> .	OFF
<code>-g0</code> (i32 only)	Disable generation of symbolic debug information.	OFF
<code>-gcc-name=name</code>	Use this option to specify the location of g++ when compiler cannot locate gcc C++ libraries. For use with <code>-cxxlib-gcc</code> configuration. Use this option when referencing a non-standard gcc installation.	OFF
<code>-gcc-version=nnn</code>	This option provides compatible behavior with gcc, where <i>nnn</i> indicates the gcc version.	OFF
<code>-[no-]global-hoist</code>	Enables [disables] hoisting and speculative loads of global variables.	OFF
<code>-H</code>	Print "include" file order and continue compilation.	OFF
<code>-help</code>	Prints compiler options summary.	OFF
<code>-idirafterdir</code>	Add directory (<i>dir</i>) to the second include file search path (after <code>-I</code>).	OFF
<code>-Idirectory</code>	Specifies an additional <i>directory</i> to search for include files.	OFF
<code>-i_dynamic</code>	Link Intel provided libraries dynamically.	OFF
<code>-inline_debug_info</code>	Produces enhanced source position information for inlined code. It also provides enhanced debug information useful for function call traceback. To use this option for debugging, you must also specify <code>-g</code> .	OFF

Option	Description	Default
<code>-ip</code>	Enables interprocedural optimizations for single file compilation.	OFF
<code>-IPF_fma[-]</code> (i64 only)	Enable [disable] the combining of floating-point multiplies and add/subtract operations.	OFF
<code>-IPF_fltacc[-]</code> (i64 only)	Enable [disable] optimizations that affect floating-point accuracy.	OFF
<code>-IPF_flt_eval_method0</code> (i64 only)	Floating-point operands evaluated to the precision indicated by the program.	OFF
<code>-IPF_fp_relaxed[-]</code> (i64 only)	Enable [disable] use of faster but slightly less accurate code sequences for math functions, such as divide and square root.	OFF
<code>-IPF_fp_speculationmode</code> (i64 only)	Enable floating-point speculations with the following <i>mode</i> conditions: <ul style="list-style-type: none"> • <code>fast</code> - speculate floating-point operations • <code>safe</code> - speculate only when safe • <code>strict</code> - same as off • <code>off</code> - disables speculation of floating-point operations 	OFF
<code>-ip_no_inlining</code>	Disables inlining that would result from the <code>-ip</code> interprocedural optimization, but has no effect on other interprocedural optimizations.	OFF
<code>-ip_no_pinlining</code> (i32, i32em)	Disable partial inlining. Requires <code>-ip</code> or <code>-ipo[<i>value</i>]</code> .	OFF
<code>-ipo[<i>n</i>]</code>	Enables interprocedural optimizations across files. The optional <i>n</i> argument controls the maximum number of link-time compilations (or number of object files) that are spawned. The default for <i>value</i> is 1 when <i>value</i> is not specified.	OFF

Option	Description	Default
-ipo_c	Generates a multifile object file (<code>ipo_out.o</code>) that can be used in further link steps.	OFF
-ipo_obj	Forces the compiler to create real object files when used with <code>-ipo[<i>value</i>]</code> .	OFF
-ipo_S	Generates a multifile assemblable file named <code>ipo_out.s</code> that can be used in further link steps.	OFF
-ipo_separate	Creates one object file for every source file. This option overrides <code>-ipo[<i>value</i>]</code> .	OFF
-isystemdir	Add directory <i>dir</i> to the start of the system include path.	OFF
-ivdep_parallel (i64 only)	This option indicates there is absolutely no loop-carried memory dependency in the loop where the IVDEP directive is specified.	OFF
-Kc++	Compile all source or unrecognized file types as C++ source files.	OFF
-kernel (i64 only)	Generates code for inclusion in the kernel. Prevents generation of speculation as support may not be available when code runs. Suppresses software pipelining.	OFF
-Knopic, -KNOPIC (i64 only)	Use <code>fpic</code> instead of this option.	ON for Itanium-based systems OFF for IA-32
-KPIC, -Kpic	Use <code>fpic</code> instead of this option.	OFF
-Ldirectory	Instruct linker to search <i>directory</i> for libraries.	OFF
-M	Generates makefile dependency lines for each source file, based on the <code>#include</code> lines found in the source file.	OFF

Option	Description	Default
<p><code>-march=<i>cpu</i></code> (i32 only)</p>	<p>Generate code exclusively for a given <i>cpu</i>. Values for <i>cpu</i> are:</p> <ul style="list-style-type: none"> • <code>pentiumpro</code> - Intel Pentium Pro processors • <code>pentiumii</code> - Intel Pentium II processors. • <code>pentiumiii</code> - Intel Pentium III processors. • <code>pentium4</code> - Intel Pentium 4 processors. 	<p>OFF</p>
<p><code>-mcpu=<i>cpu</i></code></p>	<p>Optimize for a specific <i>cpu</i>. For IA-32, <i>cpu</i> values are:</p> <ul style="list-style-type: none"> • <code>pentium</code> - Optimize for Pentium processor. • <code>pentiumpro</code> - Optimize for Pentium Pro, Pentium II and Pentium III processors. • <code>pentium4</code> - Optimize for Pentium 4 processor (Default). <p>The only option available on Intel® EM64T systems is <code>-mcpu=pentium4</code>. For Itanium-based Systems, <i>cpu</i> values are:</p> <ul style="list-style-type: none"> • <code>itanium</code> - Optimize for Itanium processor. • <code>itanium2</code> - Optimize for Itanium 2 processor (Default). 	<p>ON <code>pentium4</code> on IA-32</p> <p><code>itanium2</code> on Itanium-based Systems</p>
<p><code>-MD</code></p>	<p>Preprocess and compile. Generate output file (<code>.d</code> extension) containing dependency information.</p>	<p>OFF</p>
<p><code>-MF<i>file</i></code></p>	<p>Generate makefile dependency information in <i>file</i>. Must specify <code>-M</code> or <code>-MM</code>.</p>	<p>OFF</p>
<p><code>-MG</code></p>	<p>Similar to <code>-M</code>, but treats missing header files as generated files.</p>	<p>OFF</p>
<p><code>-MM</code></p>	<p>Similar to <code>-M</code>, but does not include system header files.</p>	<p>OFF</p>

Option	Description	Default
-MMD	Similar to -MD, but does not include system header files.	OFF
-mp	Favors conformance to the ANSI C and IEEE 754 standards for floating-point arithmetic.	OFF
-mp1	Improve floating-point precision (speed impact is less than -mp).	OFF
-MP	Add a phony target for each dependency.	OFF
-mrelax (i64 only)	Pass -relax to the linker.	ON
-mno-relax (i64 only)	Do not pass -relax to the linker.	OFF
-MQtarget	Same as -MT, but quotes special Make characters.	OFF
-mserialize-volatile (i64 only)	Impose strict memory access ordering for volatile data object references.	OFF
-mno-serialize-volatile (i64 only)	The compiler may suppress both run-time and compile-time memory access ordering for volatile data object references. Specifically, the .rel/.acq completers will not be issued on referencing loads and stores.	OFF
-MTtarget	Change the default target rule for dependency generation.	OFF
-nobss_init	Places variables that are initialized with zeroes in the DATA section. Disables placement of zero-initialized variables in BSS (use DATA).	OFF
-no_cpprt	Do not link in C++ run-time libraries.	OFF
-nodefaultlibs	Do not use standard libraries when linking.	
-no-gcc	Do not predefine the <code>__GNUC__</code> , <code>__GNUC_MINOR__</code> , and <code>__GNUC_PATCHLEVEL__</code> macros.	OFF

Option	Description	Default
<code>-nolib_inline</code>	Disables inline expansion of standard library functions.	OFF
<code>-nostartfiles</code>	Do not use standard startup files when linking.	OFF
<code>-nostdinc</code>	Same as <code>-X</code> .	OFF
<code>-nostdlib</code>	Do not use standard libraries and startup files when linking.	OFF
<code>-O</code>	Same as <code>-O1</code> on IA-32. Same as <code>-O2</code> on Itanium-based systems.	OFF
<code>-O0</code>	Disables optimizations.	OFF
<code>-O1</code>	Enable optimizations. Optimizes for speed. For Itanium compiler, <code>-O1</code> turns off software pipelining to reduce code size.	ON (i32)
<code>-O2</code>	Same as <code>-O1</code> on IA-32. Same as <code>-O</code> on Itanium-based systems.	ON (i64)
<code>-O3</code>	Enable <code>-O2</code> plus more aggressive optimizations that may increase the compilation time. Impact on performance is application dependent, some applications may not see a performance improvement.	OFF
<code>-Obn</code>	Controls the compiler's inline expansion. The amount of inline expansion performed varies with the value of <i>n</i> as follows: <ul style="list-style-type: none"> • 0: Disables inlining. • 1: Enables (default) inlining of functions declared with the <code>__inline</code> keyword. Also enables inlining according to the C++ language. • 2: Enables inlining of any function. However, the compiler decides which functions to inline. Enables interprocedural optimizations and has the same effect as <code>-ip</code>. 	OFF
<code>-ofile</code>	Name output <i>file</i> .	OFF

Option	Description	Default
<code>-openmp</code>	Enables the parallelizer to generate multi-threaded code based on the OpenMP* directives. The <code>-openmp</code> option works with both <code>-O0</code> and any optimization level of <code>-O1</code> , <code>-O2</code> , and <code>-O3</code> .	OFF
<code>-openmp_profile</code>	The <code>-openmp_profile</code> option enables analysis of OpenMP* applications with Thread Profiler, which is required to use this option.	OFF
<code>-openmp_report{0 1 2}</code>	Controls the OpenMP parallelizer's diagnostic levels.	OFF
<code>-openmp_stubs</code>	Enables OpenMP programs to compile in sequential mode. The OpenMP directives are ignored and a stub OpenMP library is linked sequentially.	OFF
<code>-opt_report</code>	Generates an optimization report directed to <code>stderr</code> , unless <code>-opt_report_file</code> is specified.	OFF
<code>-opt_report_filefilename</code>	Specifies the <i>filename</i> for the optimization report. It is not necessary to invoke <code>-opt_report</code> when this option is specified.	OFF
<code>-opt_report_levellevel</code>	Specifies the verbosity <i>level</i> of the output. Valid <i>level</i> arguments: <ul style="list-style-type: none"> • <code>min</code> • <code>med</code> • <code>max</code> <p>If a <i>level</i> is not specified, <code>min</code> is used by default.</p>	OFF

Option	Description	Default
<code>-opt_report_phasename</code>	Specifies the compilation <i>name</i> for which reports are generated. The option can be used multiple times in the same compilation to get output from multiple phases. Valid <i>name</i> arguments: <ul style="list-style-type: none"> • <code>ipo</code>: Interprocedural Optimizer • <code>hlo</code>: High Level Optimizer • <code>ilo</code>: Intermediate Language Scalar Optimizer • <code>ecg</code>: Code Generator • <code>omp</code>: OpenMP* • <code>all</code>: All phases 	OFF
<code>-opt_report_routinesubstring</code>	Specifies a routine <i>substring</i> . Reports from all routines with names that include <i>substring</i> as part of the name are generated. By default, reports for all routines are generated.	OFF
<code>-opt_report_help</code>	Displays all possible settings for <code>-opt_report_phase</code> . No compilation is performed.	OFF
<code>-Os</code>	Enable speed optimizations, but disable some optimizations which increase code size for small speed benefit.	OFF
<code>-p</code>	Same as <code>-qp</code> .	OFF
<code>-P, -F</code>	Stops the compilation process after C or C++ source files have been preprocessed and writes the results to files named according to the compiler's default file-naming conventions.	OFF
<code>-parallel</code>	Detects parallel loops capable of being executed safely in parallel and automatically generates multithreaded code for these loops.	OFF

Option	Description	Default
<code>-par_report{0 1 2 3}</code>	<p>Controls the auto-parallelizer's diagnostic levels 0, 1, 2, or 3 as follows:</p> <ul style="list-style-type: none"> • <code>-par_report0</code>: no diagnostic information is displayed. • <code>-par_report1</code>: indicates loops successfully auto-parallelized (default). • <code>-par_report2</code>: loops successfully and unsuccessfully auto-parallelized. • <code>-par_report3</code>: same as 2 plus additional information about any proven or assumed dependences inhibiting auto-parallelization. 	OFF
<code>-par_threshold[n]</code>	<p>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. This option is used for loops whose computation work volume cannot be determined at compile time. Default: $n=100$.</p>	OFF
<code>-pc32</code> (i32, i32em)	Set internal FPU precision to 24-bit significand.	OFF
<code>-pc64</code> (i32, i32em)	Set internal FPU precision to 53-bit significand.	OFF
<code>-pc80</code> (i32, i32em)	Set internal FPU precision to 64-bit significand.	ON
<code>-pch</code>	Automatic processing for precompiled headers.	OFF
<code>-pch_dir <i>dirname</i></code>	Directs the compiler to find and/or create a file for precompiled headers in <i>dirname</i> .	OFF
<code>-prec_div</code> (i32, i32em)	Disables the floating point division-to-multiplication optimization. Improves precision of floating-point divides.	OFF

Option	Description	Default
-prefetch[-] (i32 only)	Enables [disables] the insertion of software prefetching by the compiler. Default: -prefetch.	ON
-prof_dir <i>dirname</i>	Specify the directory (<i>dirname</i>) to hold profile information (*.dyn, *.dpi).	OFF
-prof_file <i>filename</i>	Specify the <i>filename</i> for profiling summary file.	OFF
-prof_format_32	By default, the Intel compiler creates 64-bit profiling counters (.dyn and .dpi). This option creates 32-bit counters for compatibility with the Intel C++ Compiler 7.0.	OFF
-prof_gen[x]	Instruments the program to prepare for instrumented execution and also creates a new static profile information file (.spi). With the x qualifier, extra source position is collected which enables code coverage tools.	OFF
-prof_use	Uses dynamic feedback information.	OFF
-Qinstall <i>dir</i>	Sets <i>dir</i> as root of compiler installation.	OFF
-Qlocation, <i>tool,path</i>	Sets <i>path</i> as the location of the tool specified by <i>tool</i> .	OFF
-Qoption, <i>tool,list</i>	Passes an argument <i>list</i> to another <i>tool</i> in the compilation sequence, such as the assembler or linker.	OFF
-qp	Compile and link for function profiling with UNIX* <i>prof tool</i>	OFF
-rcd (i32 only)	Disables changing of the FPU rounding control. Enables fast float-to-int conversions.	OFF
-reserve-kernel-regs (i64 only)	Reserves registers f12-f15 and f32-f127 for use by the kernel. These will not be used by the compiler.	OFF

Option	Description	Default
-[no]restrict	Enables/disables pointer disambiguation with the <code>restrict</code> qualifier.	OFF
-S	Generates assemblable files with <code>.s</code> suffix, then stops the compilation.	OFF
-scalar_rep[-]	The <code>-scalar_rep[-]</code> compiler option enables [disables] scalar replacement performed during loop transformations.	OFF
-shared	Produce a shared object.	OFF
-shared-libcxa	Link Intel <code>libcxa</code> C++ library dynamically.	OFF
-sox[-] (i32, i32em)	Enables [disables] the saving of compiler options and version information in the executable file.	-sox-
-static	Prevents linking with shared libraries.	OFF
-static-libcxa	Link Intel <code>libcxa</code> C++ library statically.	OFF
-std=gnu89	ISO C90 plus GNU extensions. Includes some C99 features.	ON
-std=gnu++98	Same as <code>-std=gnu89</code> .	OFF
-strict_ansi	Strict ANSI conformance dialect.	OFF
-syntax	Checks the syntax of a program and stops the compilation process after the C or C++ source files and preprocessed source files have been parsed. Generates no code and produces no output files. Warnings and messages appear on <code>stderr</code> .	OFF
-T <i>file</i>	Direct linker to read link commands from <i>file</i> .	OFF
-tcheck	The <code>-tcheck</code> compiler option enables analysis of threaded applications with Intel® Thread Checker, which is required to use this option.	OFF

Option	Description	Default
-tpp1 (i64 only)	Targets optimization for the Itanium processor.	OFF
-tpp2 (i64 only)	Targets optimization for the Itanium® 2 processor. Generated code is compatible with the Itanium processor.	ON
-tpp5 (i32 only)	Targets the optimizations for the Pentium processor.	OFF
-tpp6 (i32 only)	Targets the optimizations for the Pentium Pro, Pentium II and Pentium III processors.	OFF
-tpp7 (i32, i32em)	Targets optimizations for the Intel Pentium 4 processors.	ON
-[no]traceback	Tells the compiler to generate [not generate] extra information in the object file to allow the display of source file traceback information at run time when a severe error occurs.	OFF
-Uname	Suppresses any definition of a macro <i>name</i> . Equivalent to a #undef preprocessing directive.	OFF
-unrolln	Disable loop unrolling for n=0.	OFF
-unroll n	Disable loop unrolling for n=0.	OFF
-use_asm	Produce objects through assembler.	OFF
-use_msasm (i32 only)	Accept the Microsoft* MASM-style inlined assembly format instead of GNU-style.	OFF
-use_pch <i>filename</i>	Manual use of precompiled header (<i>filename.pchi</i>).	OFF
-u <i>symbol</i>	Pretend the <i>symbol</i> is undefined.	OFF
-V	Display compiler version information.	OFF
-v	Show driver tool commands and execute tools.	

Option	Description	Default
-vec_report[n] (i32, i32em)	Controls the amount of vectorizer diagnostic information. <ul style="list-style-type: none"> • n = 0 no diagnostic information • n = 1 indicates vectorized loops (DEFAULT) • n = 2 indicates vectorized/non-vectorized loops • n = 3 indicates vectorized/non-vectorized loops and prohibiting data dependence information • n = 4 indicates non-vectorized loops • n = 5 indicates non-vectorized loops and prohibiting data 	OFF
-w	Disable all warnings.	OFF
-Wall	Enable all warnings.	OFF
-Wbrief	Enable a mode in which a shorter form of the diagnostic output is used. When enabled, the original source line is not displayed and the error message text is not wrapped when too long to fit on a single line.	OFF
-Wcheck	Performs compile-time code checking for code that exhibits non-portable behavior, represents a possible unintended code sequence, or possibly affects operation of the program because of a quiet change in the ANSI C Standard.	OFF
-wn	Control diagnostics. <ul style="list-style-type: none"> • n = 0 displays errors (same as -w) • n = 1 displays warnings and errors (DEFAULT) • n = 2 displays remarks, warnings, and errors 	-w1

Option	Description	Default
<code>-wdL1[,L2, ...]</code>	Disables diagnostics <i>L1</i> through <i>LN</i> .	OFF
<code>-weL1[,L2, ...]</code>	Changes severity of diagnostics <i>L1</i> through <i>LN</i> to error.	OFF
<code>-Werror</code>	Force warnings to be reported as errors.	OFF
<code>-wnn</code>	Limits the number of errors displayed prior to aborting compilation to <i>n</i> .	OFF
<code>-wrL1[,L2, ...]</code>	Changes the severity of diagnostics <i>L1</i> through <i>LN</i> to remark.	OFF
<code>-wwL1[,L2, ...]</code>	Changes severity of diagnostics <i>L1</i> through <i>LN</i> to warning.	OFF
<code>-Wl,o1[,o2, ...]</code>	Pass options <i>o1</i> , <i>o2</i> , etc. to the linker for processing.	OFF
<code>-Wp,o1[,o2, ...]</code>	Pass options <i>o1</i> , <i>o2</i> , etc. to the preprocessor.	OFF
<code>-Wp64 (i32em, i64)</code>	Print diagnostics for 64-bit porting.	OFF
<code>-x type</code>	All source files found subsequent to <code>-x type</code> will be recognized as one of the following <i>types</i> : <ul style="list-style-type: none"> • <code>c</code> - C source file • <code>c++</code> - C++ source file • <code>c-header</code> - C header file • <code>cpp-output</code> - C preprocessed file • <code>assembler</code> - assemblable file • <code>assembler-with-cpp</code> - Assemblable file that needs to be preprocessed. • <code>none</code> - Disable recognition and revert to file extension. 	OFF
<code>-X</code>	Removes the standard directories from the list of directories to be searched for include files.	OFF

Option	Description	Default
<code>-x{K W N B P}</code> (i32, i32em)	<p>Generates specialized code for processor-specific codes K, W, N, B, and P.</p> <ul style="list-style-type: none"> • K = Intel® Pentium® III and compatible Intel processors • W = Intel Pentium 4 and compatible Intel processors • N = Intel Pentium 4 and compatible Intel processors • B = Intel Pentium M and compatible Intel processors • P = Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3) <p>Only the <code>-xW</code> and <code>-xP</code> options are available on Intel® EM64T.</p>	OFF
<code>-xlinker val</code>	Pass <i>val</i> directly to the linker for processing.	OFF
<code>-zP{1 2 4 8 16}</code>	Packs structures on 1, 2, 4, 8, or 16 byte boundaries.	OFF

Compiler Options Cross Reference

Linux*	Windows*	Description	Linux Default
<code>-A-</code>	<code>/QA-</code>	Remove all predefined macros.	OFF
<code>-Aname [val]</code>	<code>/QAname [val]</code>	Create an assertion name having value <i>val</i> .	OFF
<code>-ansi</code>	<code>/Za</code>	Enable/disable assumption of ANSI conformance.	ON

Linux*	Windows*	Description	Linux Default
-ax{K W N B P}	/Qax{K W N B P}	Generates specialized code for processor-specific codes K, W, N, B, and P while also generating generic IA-32 code. <ul style="list-style-type: none"> • K = Intel® Pentium® III and compatible Intel processors • W = Intel Pentium 4 and compatible Intel processors • N = Intel Pentium 4 and compatible Intel processors • B = Intel Pentium M and compatible Intel processors • P = Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3) 	OFF
-C	/C	Don't strip comments.	OFF
-c	/c	Compile to object (.o) only, do not link.	OFF
-Dname [=value]	/Dname [=value]	Define macro.	OFF
-E	/E	Preprocess to stdout.	OFF
-fp	/Oy-	Use EBP-based stack frame for all functions.	OFF
-g	/Zi	Produce symbolic debug information in object file. The -g option changes the default optimization from -O2 to -O0.	OFF
-H	/QH	Print include file order.	OFF
-help	/help	Print help message listing.	OFF

Linux*	Windows*	Description	Linux Default
<code>-Idirectory</code>	<code>/Idirectory</code>	Add directory to include file search path.	OFF
<code>-inline_debug_info</code>	<code>/Qinline_debug_info</code>	Preserve the source position of inlined code instead of assigning the call-site source position to inlined code.	OFF
<code>-ip</code>	<code>/Qip</code>	Enable single-file IP optimizations (within files).	OFF
<code>-ip_no_inlining</code>	<code>/Qip_no_inlining</code>	Optimize the behavior of IP: disable full and partial inlining (requires <code>-ip</code> or <code>-ipo[<i>value</i>]</code>).	OFF
<code>-ipo[<i>value</i>]</code>	<code>/Qipo[<i>value</i>]</code>	Enable multifile IP optimizations (between files).	OFF
<code>-ipo_obj</code>	<code>/Qipo_obj</code>	Optimize the behavior of IP: force generation of real object files (requires <code>-ipo[<i>value</i>]</code>).	OFF
<code>-KPIC</code>	NA	Generate position independent code (same as <code>-Kpic</code>).	OFF
<code>-Kpic</code>	NA	Generate position independent code (same as <code>-KPIC</code>).	OFF
<code>-m</code>	NA	Instruct linker to produce map file.	OFF
<code>-M</code>	<code>/QM</code>	Generate makefile dependency information.	OFF
<code>-mp</code>	<code>/Op[-]</code>	Maintain floating-point precision (disables some optimizations).	OFF
<code>-mp1</code>	<code>/Qprec</code>	Improve floating-point precision (speed impact is less than <code>-mp</code>).	OFF

Linux*	Windows*	Description	Linux Default
-nobss_init	/Qnobss_init	Disable placement of zero-initialized variables in BSS (use DATA).	OFF
-nolib_inline	/Oi[-]	Disable inline expansion of intrinsic functions.	OFF
-O	/O2		OFF
-ofile	/Fefile or /Fofile	Name output file.	OFF
-O0	/Od	Disable optimizations.	OFF
-O1	/O1	Optimizes for speed.	OFF
-O2	/O2		ON
-P	/EP	Preprocess to file.	OFF
-pc32	/Qpc 32	Set internal FPU precision to 24-bit significand.	OFF
-pc64	/Qpc 64	Set internal FPU precision to 53-bit significand.	OFF
-pc80	/Qpc 80	Set internal FPU precision to 64-bit significand.	ON
-prec_div	/Qprec_div	Improve precision of floating-point divides (some speed impact).	OFF
-prof_dirdirectory	/Qprof_dirdirectory	Specify directory for profiling output files (*.dyn and *.dpi).	OFF
-prof_filefilename	/Qprof_filefilename	Specify file name for profiling summary file.	OFF
-prof_gen[x]	/Qprof_genx	Instrument program for profiling; with the x qualifier, extra information is gathered.	OFF
-prof_use	/Qprof_use	Enable use of profiling information during optimization.	OFF

Linux*	Windows*	Description	Linux Default
<code>-Qinstall dir</code>	NA	Set <i>dir</i> as root of compiler installation.	OFF
<code>-Qlocation, str, dir</code>	<code>/Qlocation, tool, path</code>	Set <i>dir</i> as the location of tool specified by <i>str</i> .	OFF
<code>-Qoption, str, opts</code>	<code>/Qoption, tool, list</code>	Pass options <i>opts</i> to tool specified by <i>str</i> .	OFF
<code>-qp, -p</code>	NA	Compile and link for function profiling with UNIX* gprof tool.	OFF
<code>-rcd</code>	<code>/Qrcd</code>	Enable fast floating-point-to-integer conversions.	OFF
<code>-restrict</code>	<code>/Qrestrict</code>	Enable the restrict keyword for disambiguating pointers.	OFF
<code>-S</code>	<code>/S</code>	Generates assemblable files with <code>.s</code> suffix, then stops the compilation.	OFF
<code>-sox[-]</code>	<code>/Qsox</code>	Enable [disable] saving of compiler options and version in the executable.	<code>-sox-</code>
<code>-syntax</code>	<code>/Zs</code>	Perform syntax check only.	OFF
<code>-tpp5</code>	<code>/G5</code>	Optimize for Pentium processor.	OFF
<code>-tpp6</code>	<code>/G6</code>	Optimize for Pentium Pro, Pentium II and Pentium III processors.	OFF
<code>-tpp7</code>	<code>/G7</code>	Optimize for Pentium 4 processor.	OFF

Linux*	Windows*	Description	Linux Default
-[no]traceback	/[no]traceback	Generate [do not generate] extra information in the object file that allows the display of source file traceback information at run time when a severe error occurs.	OFF
-Uname	/Uname	Remove predefined macro.	OFF
-unroll0	/Qunroll0	Disable loop unrolling.	OFF
-V	/QV	Display compiler version information.	OFF
-w	/w	Display errors.	OFF
-w2	/W4	Enable remarks, warnings and errors.	
-Wbrief	/WL	Produces less verbose diagnostics.	OFF
-wn	/Wn	Control diagnostics. Display errors (n=0). Display warnings and errors (n=1). Display remarks, warnings, and errors (n=2).	OFF
-wdL1[,L2,...]	/Qwd[tag]	Disable diagnostics L1 through LN.	OFF
-weL1[,L2,...]	/Qwe[tag]	Change severity of diagnostics L1 through LN to error.	OFF
-wnn	/Qwn[tag]	Print a maximum of n errors.	OFF
-Wp64	/Wp64	Print diagnostics for 64-bit porting.	OFF
-wrL1[,L2,...]	/Qwr[tag]	Change severity of diagnostics L1 through LN to remark.	OFF
-wwL1[,L2,...]	/Qww[tag]	Change severity of diagnostics L1 through LN to warning.	OFF

Linux*	Windows*	Description	Linux Default
-X	/X	Remove standard directories from include file search path.	OFF
-x{K W N B P}	/Qx{K W N B P}	Generates specialized code for processor-specific codes K, W, N, B, and P while also generating generic IA-32 code. <ul style="list-style-type: none"> • K = Intel® Pentium® III and compatible Intel processors • W = Intel Pentium 4 and compatible Intel processors • N = Intel Pentium 4 and compatible Intel processors • B = Intel Pentium M and compatible Intel processors • P = Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3) 	OFF
-Zp{1 2 4 8 16}	/Zp[n]	Packs structures on 1, 2, 4, 8, or 16 byte boundaries.	OFF

Default Compiler Options

Some compiler options are only available on certain systems. In the following table, these options are indicated with labels as follows:

Label	Meaning
i32	Option available on IA-32-based systems
i32em	Option available on Intel® Extended Memory 64 Technology (Intel® EM64T) systems
i64	Option available on Itanium®-based systems

- If no label is present, the option is available on all supported systems
- If "only" appears in the label, that option is only available on the identified system

Option	Description
-alias_args	Enable C/C++ rule that function arguments may be aliased.
-ansi_alias	Enable use of ANSI aliasing rules in optimizations; user asserts that the program adheres to these rules.
-complex_limited_range-	Disable the use of the basic algebraic expansions of some complex arithmetic operations.
-falias	Assume aliasing in program.
-ffnalias	Assume aliasing within functions.
-frtti	Support for RTTI.
-fverbose-asm	Produce assembly file with compiler comments (requires -S).
-mcpu=pentium4 (i32 only)	Optimizes for Intel® Pentium® 4 processor.
-mcpu=itanium2 (i64 only)	Optimizes for Intel® Itanium® 2 processor.
-O1	Same as -O2 on IA-32. Same as -O on Itanium-based systems.
-pc80 (i32, i32em)	Set internal floating-point precision to 64-bit significand.
-prefetch	Enables the insertion of software prefetching by the compiler.
-sox-	Disable saving of compiler options and version in the executable.

Option	Description
-std=gnu89	ISO C90 plus GNU extensions. Includes some C99 features.
-tpp2 (i64 only)	Target optimization to the Itanium 2 processor.
-tpp7 (i32 only)	Target optimization to the Pentium 4 2 processor.
-w1	Display warnings and errors.

Deprecated and Unsupported Compiler Options

Deprecated Options

Occasionally, compiler options are marked as "deprecated." Deprecated options are still supported in the current release, but are planned to be unsupported in future releases. The following options are deprecated in this release of Intel® C++ Compiler:

- -Qansi

Deprecated options are not limited to this list.

Unsupported Options

Some Intel C++ Compiler options are no longer supported. If you use an unsupported option, the compiler issues a warning, ignores the option, then proceeds with compilation. This version of the Intel C++ Compiler no longer supports the following compiler options:

- -axi
- -axM
- -xi
- -xM
- -of_check
- -fdiv_check

Unsupported options are not limited to this list.

Volume I: Building Applications

Getting Started

You can Invoke the Compiler from a system command prompt, or you can use the compiler with the Eclipse* Integrated Development Environment.

Getting Help

- Documentation conventions are described in How to Use This Document.
- If you are using the compiler from the command line, you can execute `icc -help` for a summary of command-line options.
- If you need additional help in using the Intel C++ Compiler, see Product Web Site and Support.

Default Behavior of the Compiler

If you do not specify any options when you invoke the Intel® C++ Compiler, the compiler uses the following default settings:

- Produces executable output with filename `a.out`.
- Invokes options specified in a configuration file first. See Configuration Files.
- The location of shared objects is specified by the `LD_LIBRARY_PATH` environment variable.
- Sets 8 bytes as the strictest alignment constraint for structures.
- Displays error and warning messages.
- Performs standard optimizations using the default `-O2` option. See Setting Optimization Levels.
- On operating systems that support characters in Unicode* (multi-byte) format, the compiler will process file names containing these characters.

If the compiler does not recognize a command-line option, that option is ignored and a warning is displayed. See Diagnostic Messages for detailed descriptions about system messages.

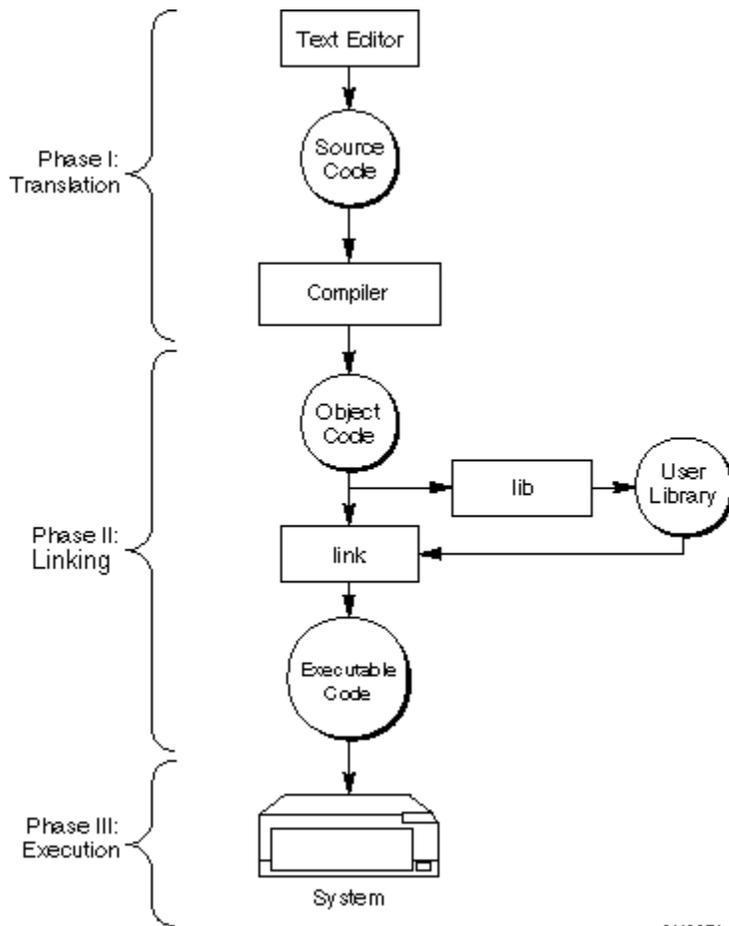
Compilation Phases

To produce an executable file, the compiler performs by default the compile and link phases. When invoked, the compiler driver determines which compilation phases to perform based on the file name extension and the compilation options specified in the command line.

The compiler passes object files and any unrecognized file name to the linker. The linker then determines whether the file is an object file (`.o`) or a library (`.a`). The compiler driver handles all types of input files correctly, thus it can be used to invoke any phase of compilation.

The relationship of the compiler to system-specific programming support tools is presented in this diagram:

Application Development Cycle



OM09714

Building Applications from the Command Line

Invoking the Compiler

The ways to invoke Intel® C++ Compiler are as follows:

- Invoke directly: Running Compiler from the Command Line
- Use system make file: Running from the Command Line with make

Invoking the Compiler from the Command Line

There are two necessary steps to invoke the Intel® C++ Compiler from the command line:

1. set the environment
2. invoke the compiler using `icc` or `icpc`

Set the Environment Variables

Before you can operate the compiler, you must set the environment variables to specify locations for the various components. The Intel C++ Compiler installation includes shell scripts that you can use to set environment variables. With the default compiler installation, these scripts are:

- `/opt/intel_cc_80/bin/iccvars.sh`
- `/opt/intel_cc_80/bin/iccvars.csh`

To run an environment script, enter one of the following on the command line:

```
prompt>>source /opt/intel_cc_80/bin/iccvars.sh
```

or

```
prompt>>source /opt/intel_cc_80/bin/iccvars.csh
```

If you want the script to run automatically when you start Linux*, add the same command to the end of your startup file.

Sample `.bash_profile` entry for `iccvars.sh`:

```
# set environment vars for Intel C++ compiler
source /opt/intel_cc_80/bin/iccvars.sh
```

Invoking the Compiler with `icc` or `icpc`

You can invoke the Intel C++ Compiler on the command line with either `icc` or `icpc`.

- When you invoke the compiler with `icc`, the compiler builds C source files using C libraries and C include files. If you use `icc` with a C++ source file, it will be compiled as a C++ file. Use `icc` to link C object files.
- When you invoke the compiler with `icpc` the compiler builds C++ source files using C++ libraries and C++ include files. If you use `icpc` with a C source file, it will be compiled as a C++ file. Use `icpc` to link C++ object files.

Command-line Syntax

When you invoke the Intel C++ Compiler with `icc` or `icpc`, use the following syntax:

```
prompt>{icc|icpc} [options] file1 [file2 . . .]
```

Argument	Description
options	Indicates one or more command-line options. The compiler recognizes one or more letters preceded by a hyphen (-). This includes linker options. See the Options Quick Reference
file1, file2 . . .	Indicates one or more files to be processed by the compilation system. You can specify more than one file. Use a space as a delimiter for multiple files.

Example:

```
prompt>icpc -prec_div -axP -Bstatic my_source1.cpp my_source2.cpp
```

Invoking the Compiler from the Command Line with `make`

To run `make` from the command line using Intel® C++ Compiler, make sure that `/usr/bin` is in your path. If you use a C shell, you can edit your `.cshrc` file and add:

```
setenv PATH /usr/bin:<full path to Intel compiler>
```



Note

To use the Intel compiler, your makefile must include the setting `CC=icc`. Use the same setting on the command line to instruct the makefile to use the Intel compiler. If your makefile is written for `gcc`, the GNU* C compiler, you will need to change those command line options not recognized by the Intel compiler.

Then you can compile:

```
prompt>make -f my_makefile
```

Compiler Input Files

The Intel® C++ Compiler recognizes the file name extensions listed in the following table:

Filename	Interpretation
filename.a	Object library
filename.i	When you invoke the compiler with <code>icc</code> , the <code>.i</code> files are treated as C source files. The <code>.i</code> files are treated as C++ sources if you compile with <code>icpc</code> .
filename.o	Compiled object module
filename.s	Assembly file
filename.so	Shared object file
filename.S	Assembly file that requires preprocessing
filename.c	C language source file
filename.C filename.cc filename.CC filename.cpp filename.cxx	C++ language source file

Building Applications in Eclipse*

The Intel® C++ Compiler for Linux (IA-32 only) includes a compiler integration with Eclipse* and the C/C++ Development Tools* (CDT). This functionality is an optional part of the compiler installation. For more information about CDT, see <http://www.eclipse.org/cdt/>.

The Intel C++ Compiler integration with the Eclipse/CDT integrated development environment lets you develop, build, and run your C/C++ projects in a visual, interactive environment.

This section includes the following topics:

- Starting Eclipse*
- Using Online Help in Eclipse*
- Creating a New Project
- Setting Properties
- Standard and Managed Make Files

Starting Eclipse*

After you have installed the following:

- Intel® C++ Compiler for 32-bit applications
- Eclipse* integrated development environment
- Java* Runtime Environment (JRE)
- C/C++ Development Tools (CDT)

you can execute the `iccec` shell script to start Eclipse from a directory where you have write permission. With the default compiler installation, execute `iccec` as follows:

```
prompt>/opt/intel_cc_80/bin/iccec
```

You can also use `iccec` to pass Eclipse-specific parameters, such as:

- `-data <path>` - sets the location for the Eclipse workspace
- `-showlocation` - shows the location of the workspace in the Eclipse window title bar.

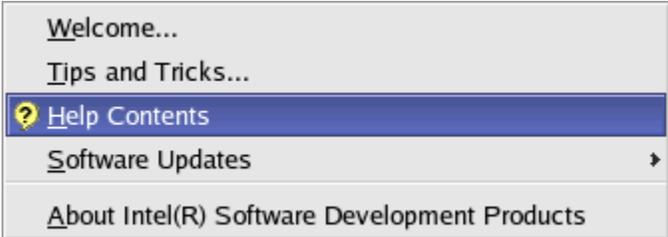
For example:

```
prompt>/opt/intel_cc_80/bin/iccec -data /cpp/eclipse -  
showlocation
```

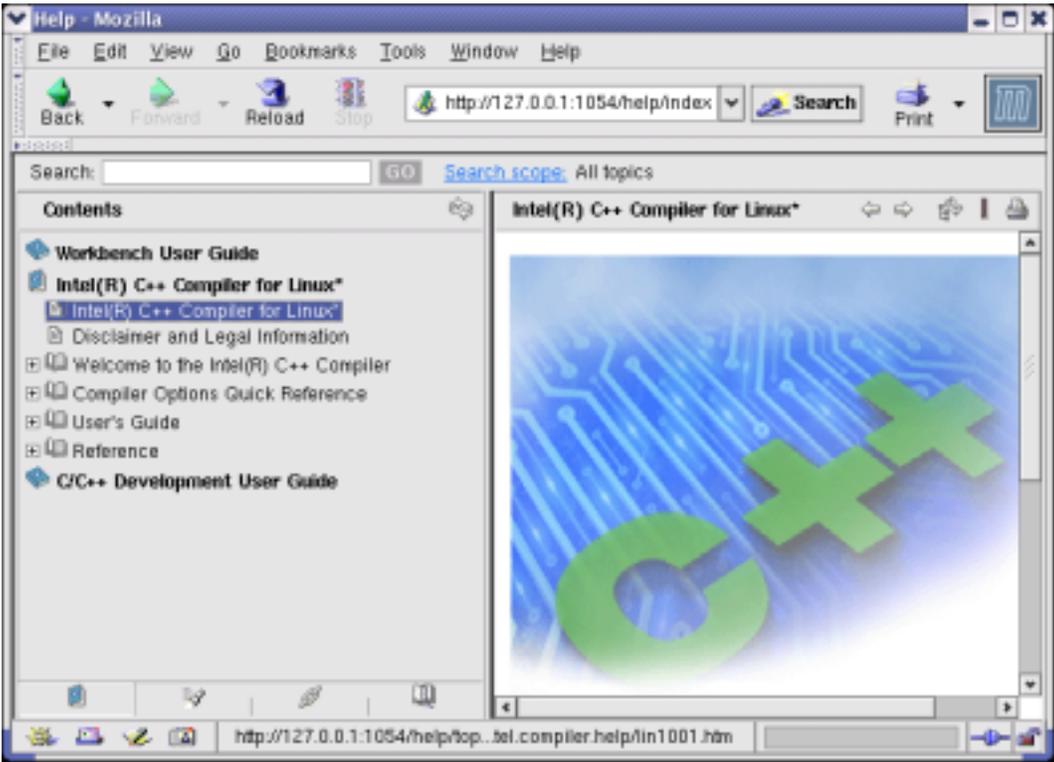
From the Eclipse **Help** menu, select **Help Contents > Workbench User's Guide > Tasks > Running Eclipse** for the complete list of Eclipse startup parameters.

Using Online Help in Eclipse*

The Intel® C++ Compiler integration with Eclipse*/CDT* includes online help. From the Eclipse toolbar, select **Help > Help Contents**.

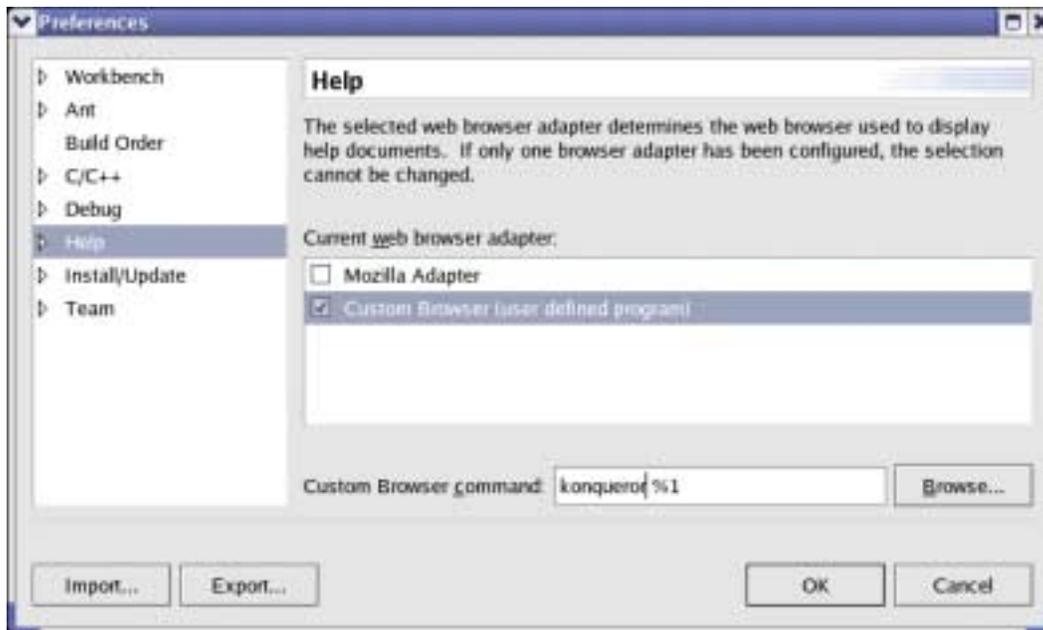


The **Help Contents** option lets you narrow your search for help information by presenting all the help modules registered with Eclipse. Select **Intel(R) C++ Compiler for Linux*** to open the Compiler User's Guide (this document). The **Help Contents** may also include links to the Eclipse **Workbench User Guide**, the **C/C++ Development User Guide**, and other similar documents.



Selecting a Different Browser

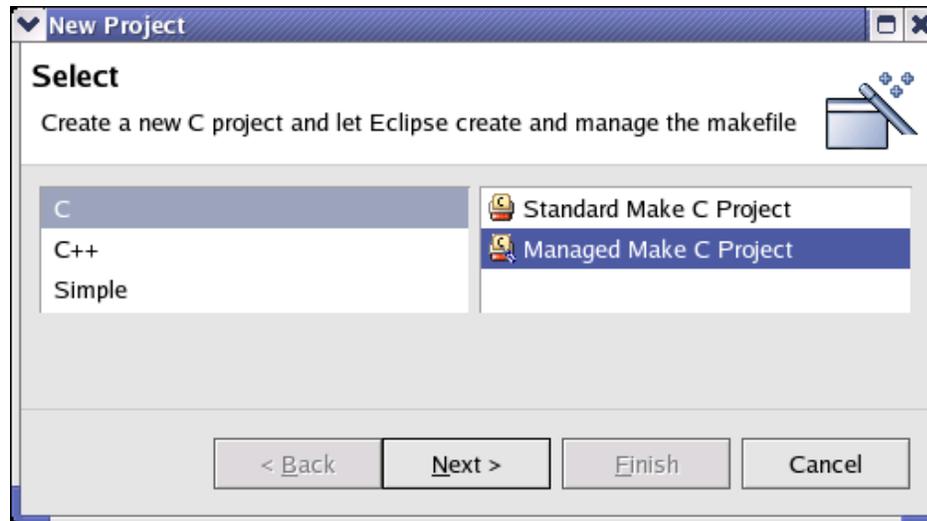
If you want to select a different browser to view the **Help Contents**, open **Windows > Preferences > Help** from the Eclipse toolbar. Check **Custom Browser (user defined program)**, then complete the necessary information in the **Custom Browser command** text box. Click **OK** to complete your browser selection.



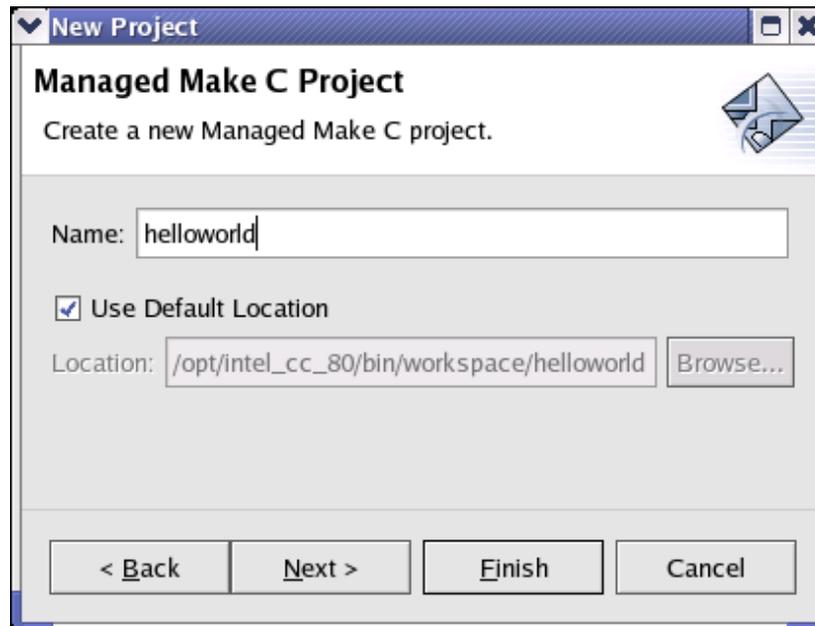
Creating a New Project

To create a simple `helloworld` project, follow these steps after starting Eclipse*:

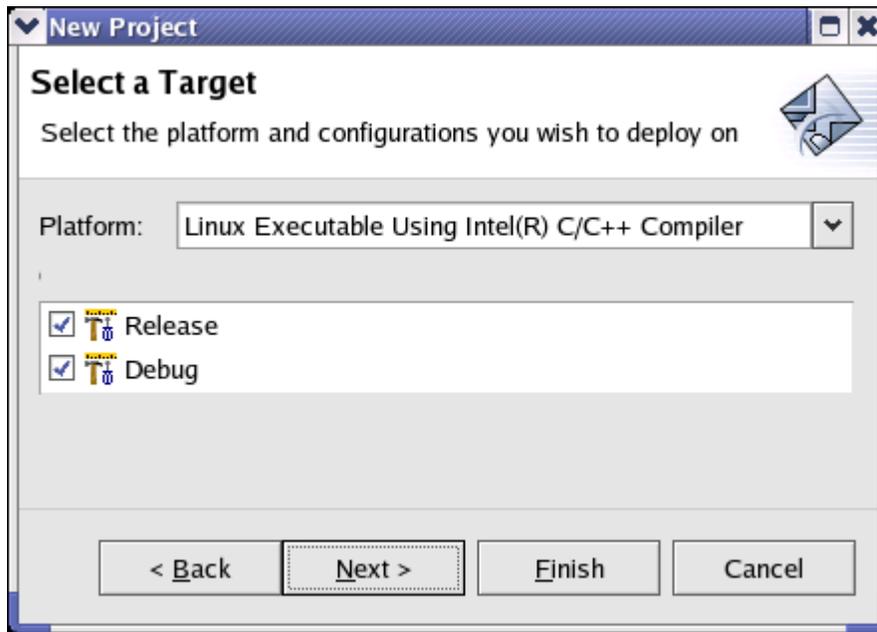
1. Select **Window > Open Perspective > C/C++ Development**.
2. From the Eclipse **File** menu, select **New > Project**. The **New Project** wizard opens with the **Select** dialog to specify the kind of project you want to create. In the left column, select **C** from the list. In the right column, select **Managed Make C Project**. Click **Next** to proceed. See also Standard and Managed Make Files.



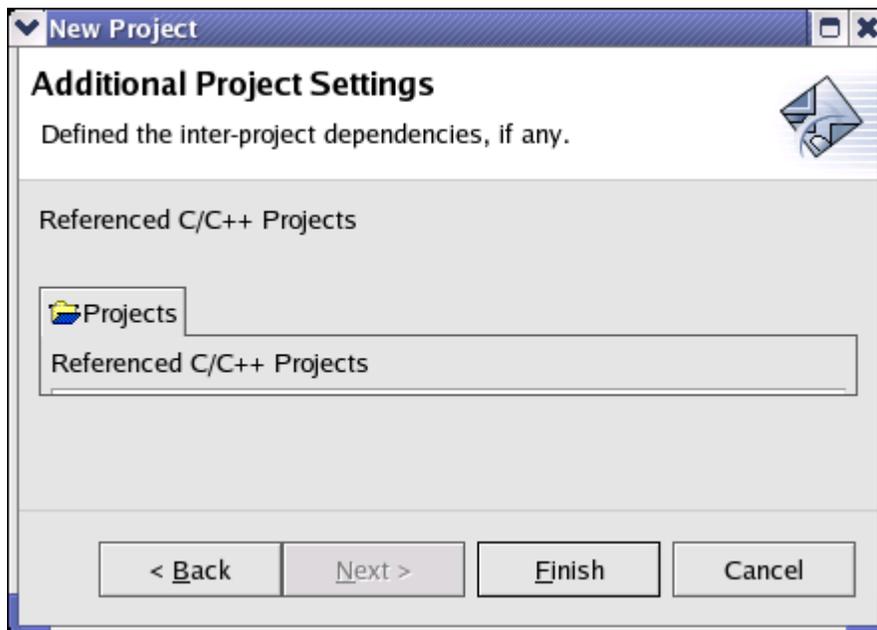
3. In the **Name** text box of the **Managed Make C Project** dialog, type `helloworld`. Check the **Use Default Location** box, if not already checked. Click **Next** to proceed.



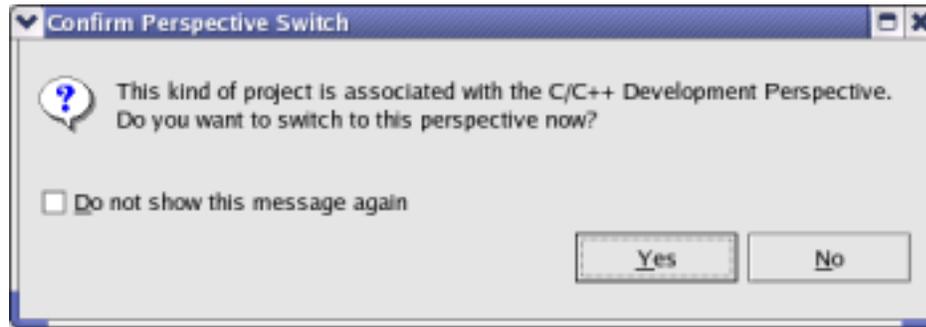
- From the **Select a Target** dialog, select **Linux Executable Using Intel(R) C/C++ Compiler** from the **Platform** drop-down list. Check the **Release** and **Debug** configuration boxes. Click **Next** to proceed.



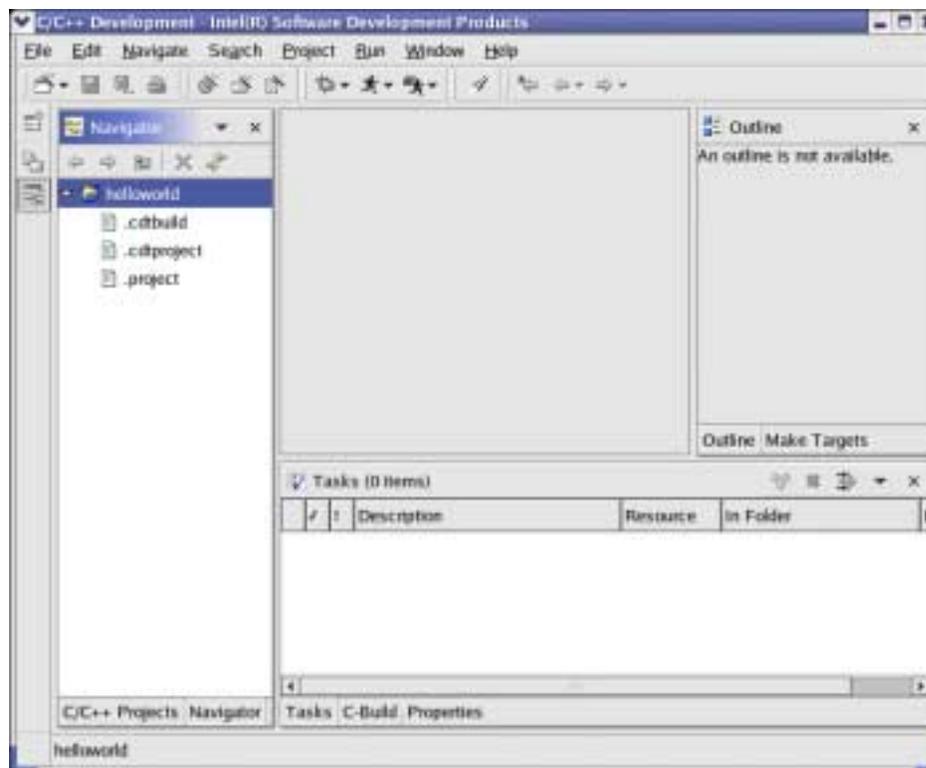
- The **Additional Project Settings** dialog lets you create dependencies between your new project and other existing projects. There should not be any other existing projects at this point. Click **Finish** to complete creation of your new `helloworld` project.



- If you are not currently in the **C/C++ Development Perspective**, you will see the **Confirm Perspective Switch** dialog. Click **Yes** to proceed.



- In the **Navigator** view, you should now see an entry for your `helloworld` project.



The next step is Adding a C Source File.

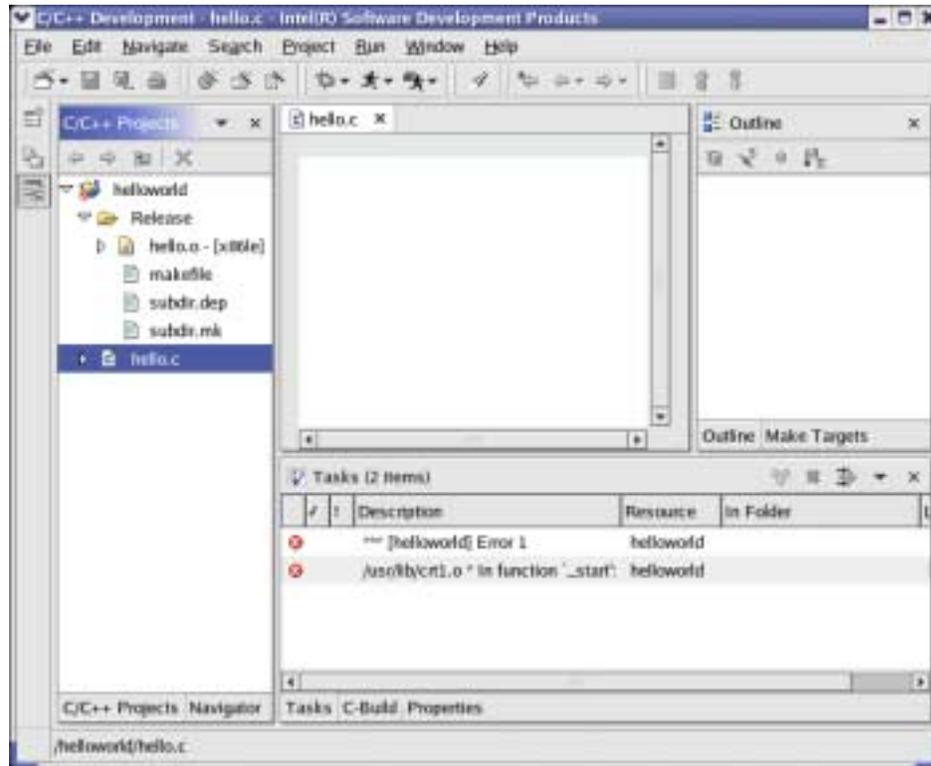
Adding a C Source File

After Creating a New Project, you can add source files, then build and run your completed project. Follow these steps to add a `hello.c` source file to the `helloworld` project.

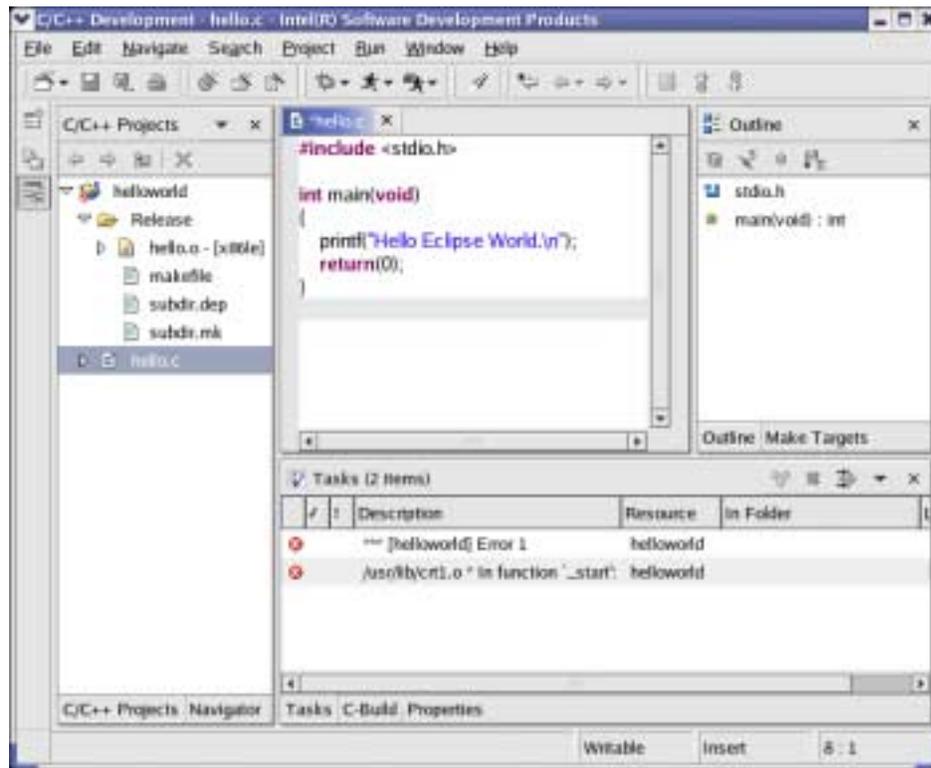
1. From the Eclipse* **File** menu, select **New > File**. Enter `hello.c` in the **File name** text box of the **New File** dialog. Click **Finish** to add the file to the `helloworld` project.



- Your Eclipse Preference settings in **Window > Preferences > Workbench** let you specify **Perform build automatically on resource modification**. If this preference is checked, Eclipse/CDT will attempt a build when `hello.c` is created. Since `hello.c` does not yet include code, errors are indicated in the **Tasks** view and **C-Build** view near the bottom of the screen. This is expected behavior, not a true error. Select **Window > Show View > C/C++ Projects** to view the project files.



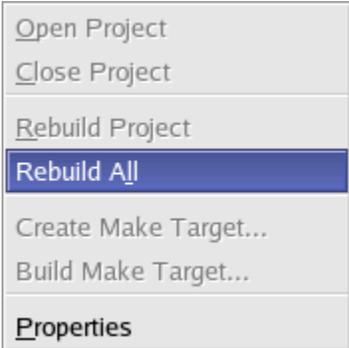
3. In the **Editor** view, add your code for `hello.c`. If you close `hello.c` in the **Editor** view, you can open it by double-clicking on `hello.c` in the **Navigator** view.



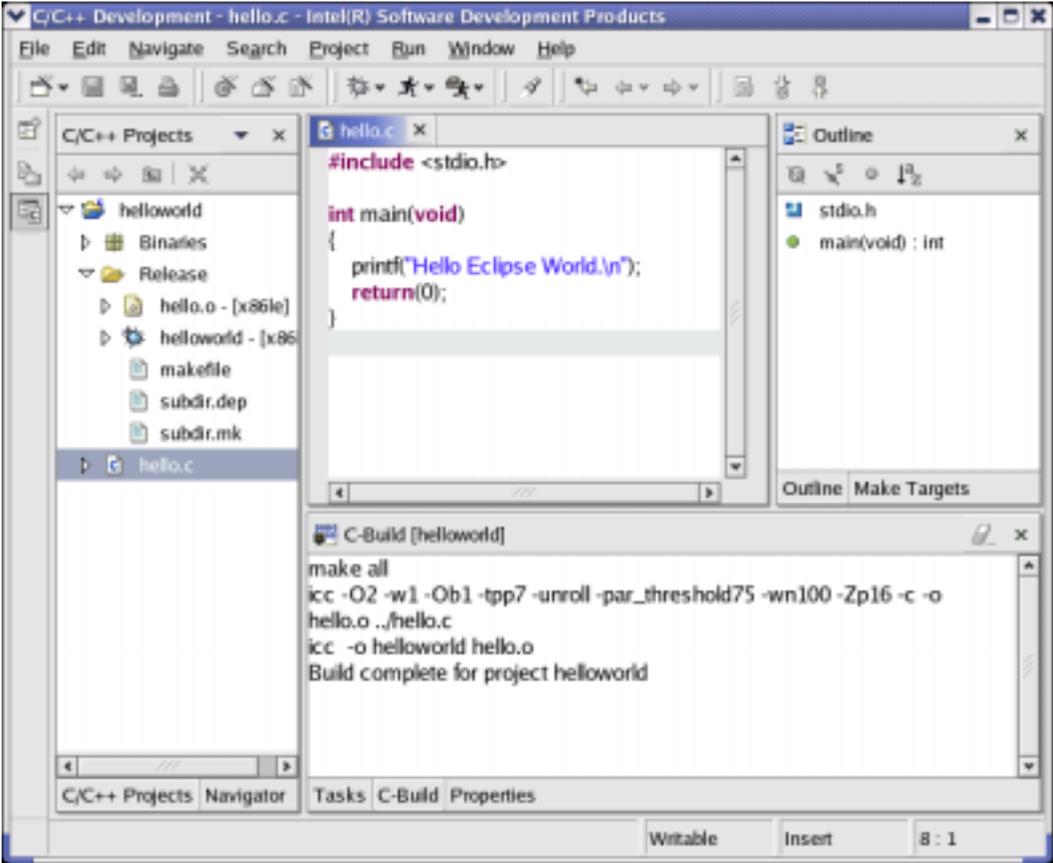
When your code is complete, save your file using **File > Save**, then proceed to Building a Project.

Building a Project

You can build your project by selecting **Rebuild All** from the Eclipse* **Project** menu.



See the **Build** results in the **C-Build** view.

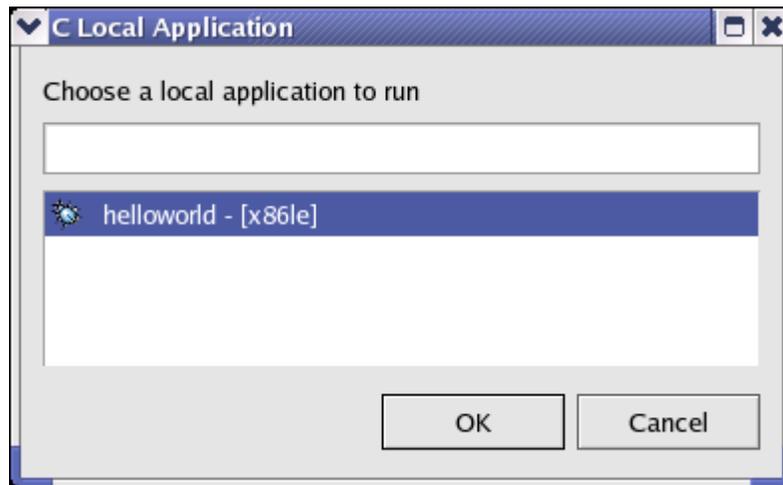


The final step is Running a Project.

Running a Project

After Building a Project, you can run your project by following these steps:

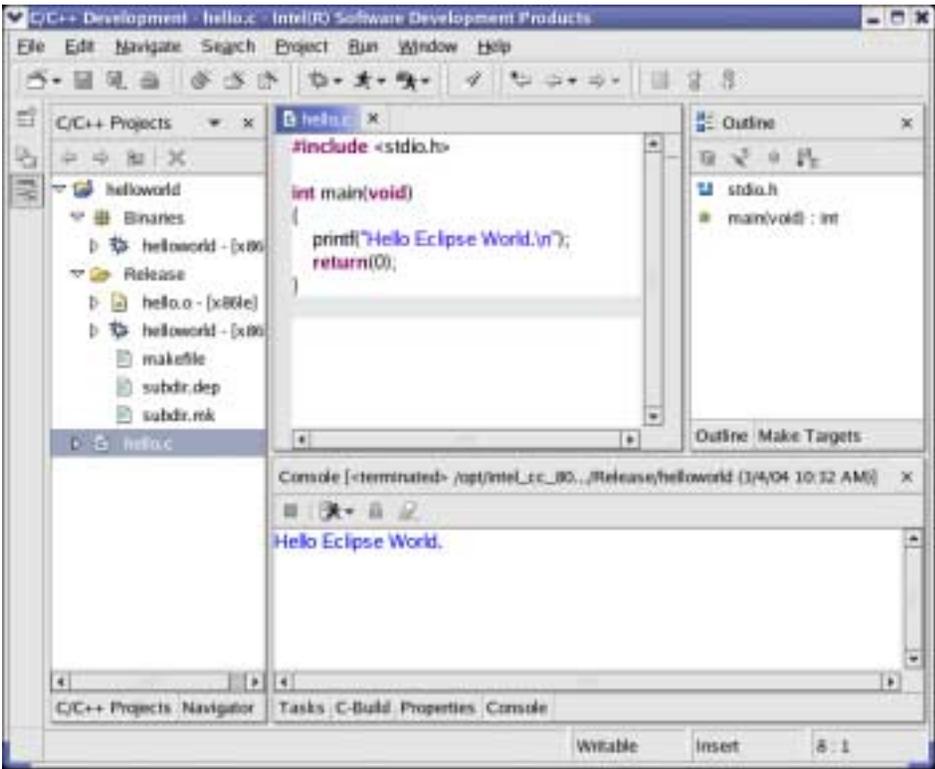
1. Select **Run > Run As > C Local Application**. When the **C Local Application** dialog appears, click **OK**.



2. On the **Launch Debug Configuration Selection** dialog, select **GDB Debugger**, then click **OK**.



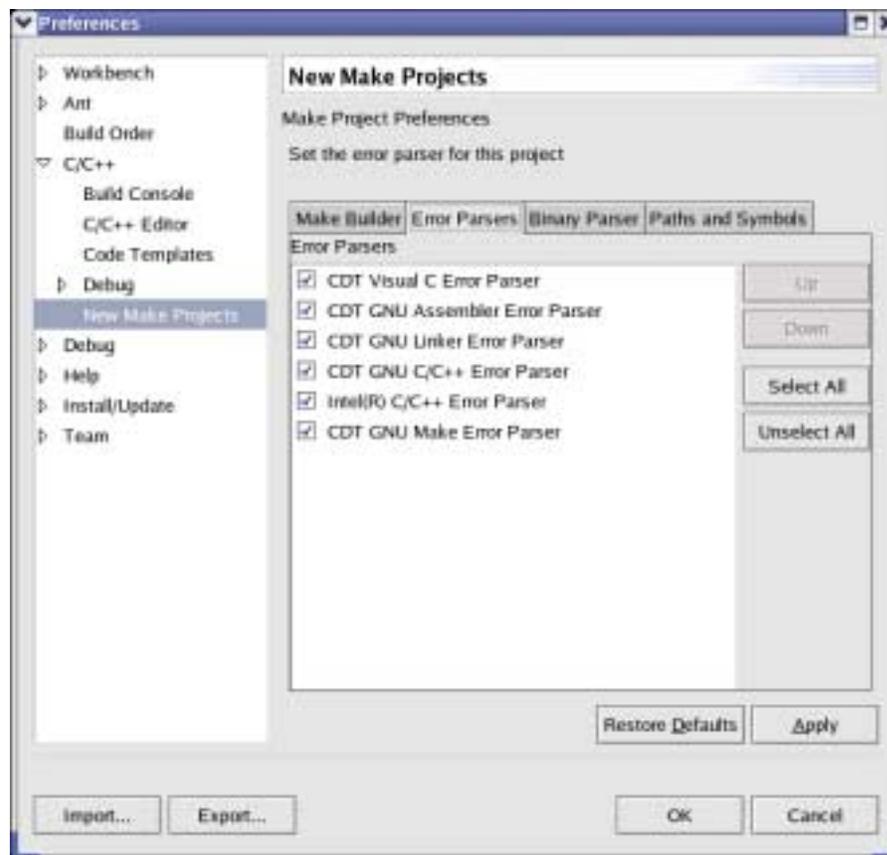
3. After the executable runs, the output of `hello.c` appears in the **Console** view.



Intel® C/C++ Error Parser

The Intel® C/C++ Error Parser lets you track compile-time errors in Eclipse*/CDT*. However, you must enable the Error Parser to see the results:

1. On the Eclipse toolbar, select **Window > Preferences**.
2. On the **Preferences** dialog, select **C/C++ > New Make Projects**.
3. Click the **Error Parsers** tab. Check the **Intel(R) C/C++ Error Parser** selection to enable this feature.



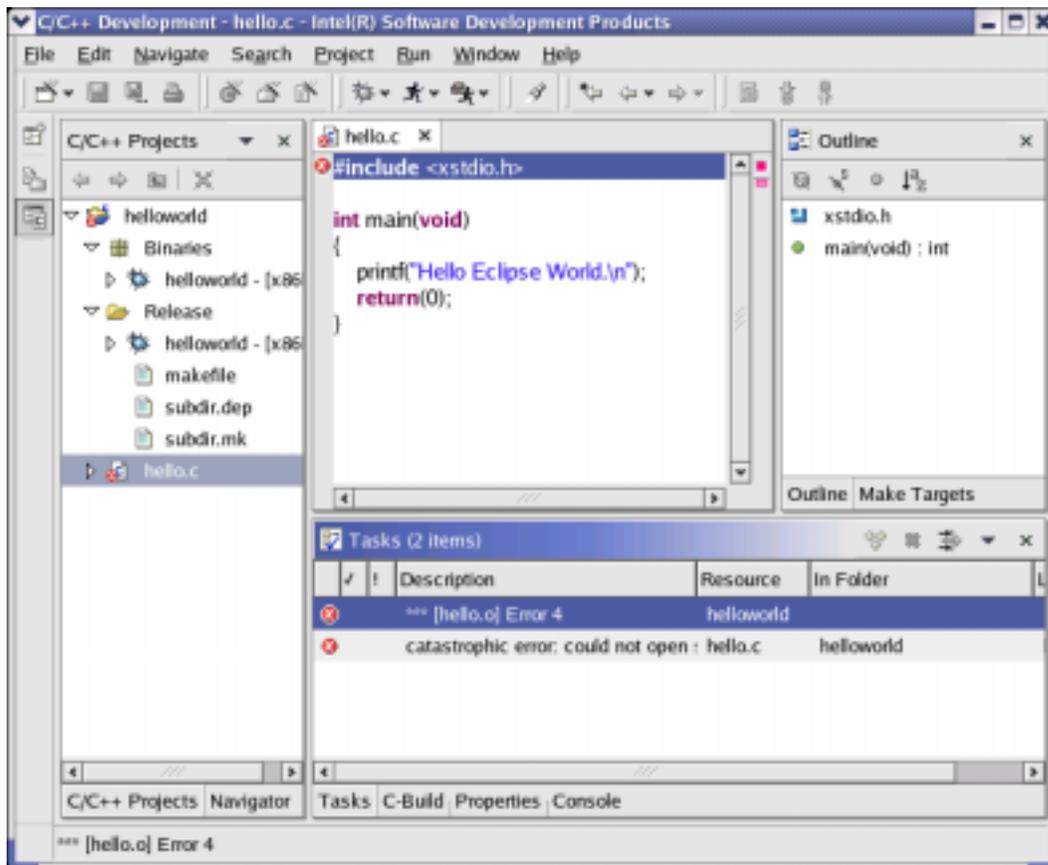
4. Click **OK** to update your choices and close the dialog.

Using the Intel C/C++ Error Parser

If you introduce an error into your `hello.c` program, such as:

```
#include <xstdio.h>
```

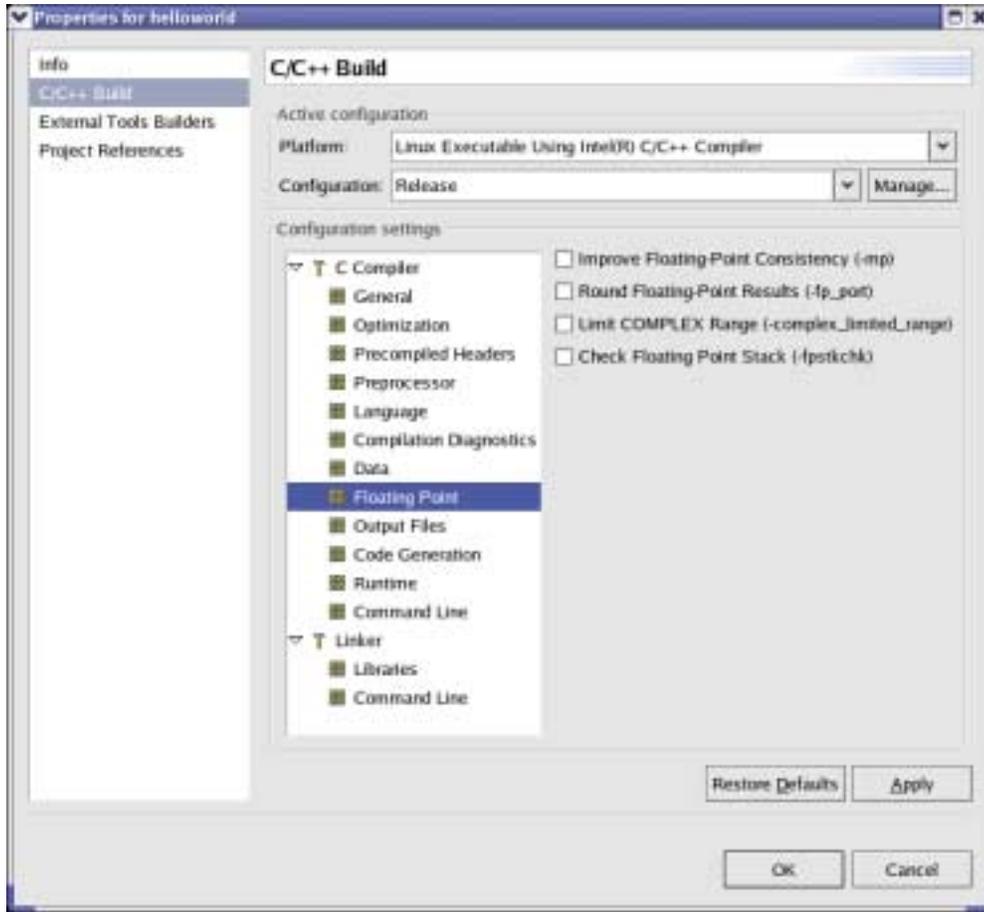
then compile `hello.c`, the error is reported in the **Tasks** view and a  marker appears in the source file at the line where the error was detected.



Setting Properties

The Intel® C++ Compiler integration with Eclipse*/CDT* lets you specify compiler, linker, and archiver options. Follow these steps to set options for your project:

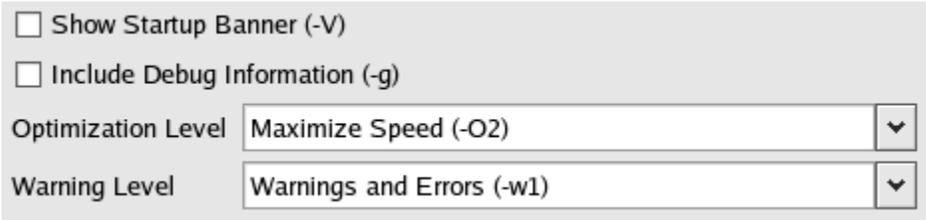
1. Select your project in the **C/C++ Projects** view.
2. From the Eclipse toolbar, select **Project > Properties > C/C++ Build**.
3. Under **Configuration settings**, click an option category for **C Compiler** or **Linker**. In the example that follows, the options in the **Floating Point** category are displayed.



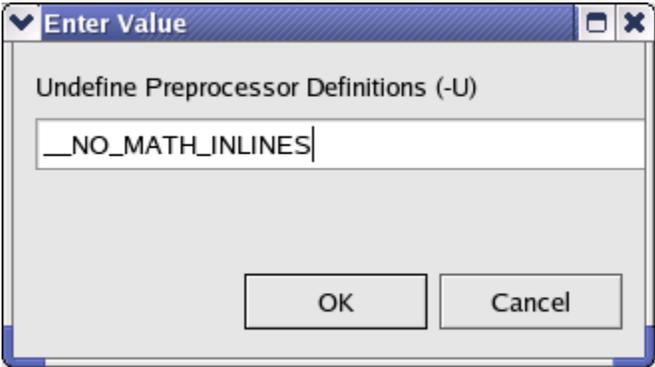
4. Check the option(s) you want to add to your project compilations, then open other categories if necessary.
5. Click **OK** to complete your selections.

To reset properties to their default setting, click **Restore Defaults**. The **Restore Defaults** button appears on each property page, but the **Restore Defaults** action applies to ALL property pages.

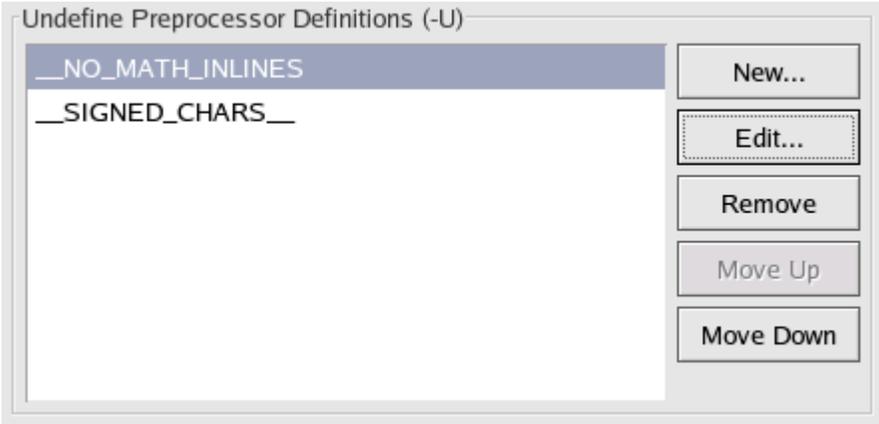
Some properties use check boxes, while others use drop-down lists to specify a compiler option.



Several options let you specify arguments. Click **New** to add an argument to the list. Enter a valid argument for the option, then click **OK**.



In this example, `__NO_MATH_INLINES` and `__SIGNED_CHARS__` are specified as arguments for the `-U` option.



If you want to specify an option that is not available from the **Properties** dialog, use the **Command Line** category. Enter the command line options in the **Additional Options** text box just as you would enter them on the command line.



For a complete list of options listed on the Properties page, see Properties for Supported Options.

Properties for Supported Options

The options listed in the following tables are supported under the corresponding Option Category.

Compiler Options

Option Category	Use	Option Name
General	Show startup banner	-V
	Include debug information	-g
	Optimization level	-O0 (default for Debug Config.)
		-O1
		-O2 (default for Release Config.)
		-O3
		-fast
	Warning level	-w0
		-w1 (default)
		-w2
Optimization	Provide frame pointers	-fp
	Disable prefetch insertion	-prefetch
	Enable interprocedural optimization for single file compilation	-ip
	Disable intrinsic inline expansion	-nolib_inline
	Inline function expansion	-Ob0
		-Ob1
		-Ob2
	Optimize for Intel® processor	-tpp5
		-tpp6
		-tpp7 (default)
Parallelization	-parallel	
Precompiled Headers	Automatic Processing for Precompiled Headers	-pch
Preprocessor	gcc compatibility options	-cxxlib-icc
		-cxxlib-gcc

Option Category	Use	Option Name	
		-fabi-version	
		-gcc-version	
	Additional include directories	-I	
	Ignore standard include path	-X	
	Preprocessor definitions	-D	
	Do not predefine <code>_GNUC_</code> , <code>_GNUC_MINOR_</code> , <code>_GNUC_PATCHLEVEL_</code> macros	-no-gcc	
	Undefine preprocessor definitions	-U	
Undefine all preprocessor definitions	-A-		
Language	Enable use of ANSI aliasing rules in optimizations	-ansi_alias	
	Disable C99 support	-c99-	
	Recognize the <code>restrict</code> keyword	-restrict	
	Process OpenMP* directives	-openmp -openmp_stubs	
Compilation Diagnostics	Treat warnings as errors	-Werror	
	Allow usage messages	-Wcheck	
	Enable equivalent of GNU* ANSI	-ansi	
	Strict ANSI conformance dialect	-strict_ansi	
	OpenMP report		-openmp_report0
			-openmp_report1
			-openmp_report2
	Auto-parallelizer report		-par_report0
			-par_report1
			-par_report2
			-par_report3
	Vectorizer report		-vec_report0
			-vec_report1
			-vec_report2
			-vec_report3
		-vec_report4	
	-vec_report5		

Option Category	Use	Option Name	
Data	Disable argument aliasing	-alias_args-	
	Assume no aliasing in program	-fno-alias	
	Allow gprel addressing of common data variables	-fno-common	
	Allocate as many bytes as needed for enumerated types	-fshort-enums	
	Change default bitfield type to unsigned	-funsigned-bitfields	
	Change default char type to unsigned	-funsigned-char	
	Store string literals in a writable section	-fwritable-strings	
	Disable placement of zero-initialized variables in .bss - use .data	-nobss_init	
	Default symbol visibility		-fvisibility=extern
			-fvisibility=default
			-fvisibility=protected
			-fvisibility=hidden
			-fvisibility=internal
	Structure member alignment		-Zp1
			-Zp2
-Zp4			
-Zp8			
-Zp16 (default)			
Floating Point	Improve floating-point consistency	-mp	
	Round floating-point results	-fp_port	
	Limit Complex range	-complex_limited_range	
	Check floating-point stack	-fpstkchk	
Output Files	Generate assembler source file	-S	
Code Generation	Generate position-independent code	-fpic	
	Use Intel® processor extensions	-axK	
		-axN	
		-axB	
		-axP	
	Require Intel® processor extensions	-xK	
-xN			

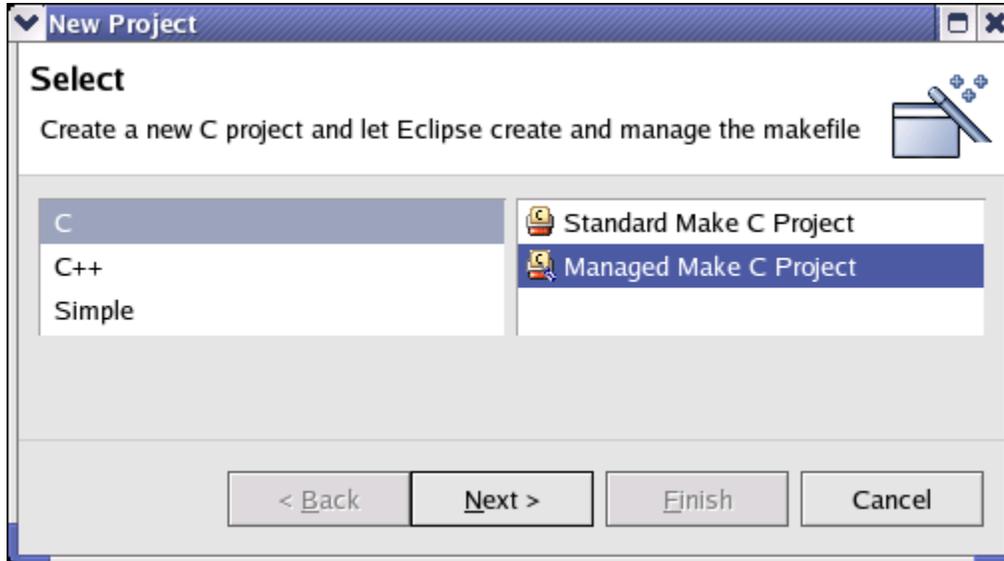
Option Category	Use	Option Name
		-xB
		-xP
Run-time	Generate traceback information	-traceback

Library Options

Libraries	Maximize speed across entire program	-fast
	Enable interprocedural optimization for single file compilation	-ip
	Link with static libraries	-static
	Link Intel® libcxa C++ library statically	-static-libcxa
	Link with dynamic libraries	-i_dynamic
	Use no C++ libraries	-no_cpprt
	Use no system libraries	-nodefaultlibs
	gcc compatibility options	-cxxlib-icc
		-cxxlib-gcc
		-fabi-version
		-gcc-version
	Process OpenMP directives	-openmp
		-openmp_stubs
	Additional libraries	-l
	Search directory for libraries	-L
Archiver options	-r	

Standard and Managed Makefiles

When you create a new Intel C project in Eclipse*/CDT*, you can select either **Standard Make C Project** or **Managed Make C Project**.



- Select **Standard Make C Project** if your project already includes a makefile.
- Use **Managed Make C Project** to build a makefile using Intel compiler-specific options assigned from property pages.

Exporting Makefiles

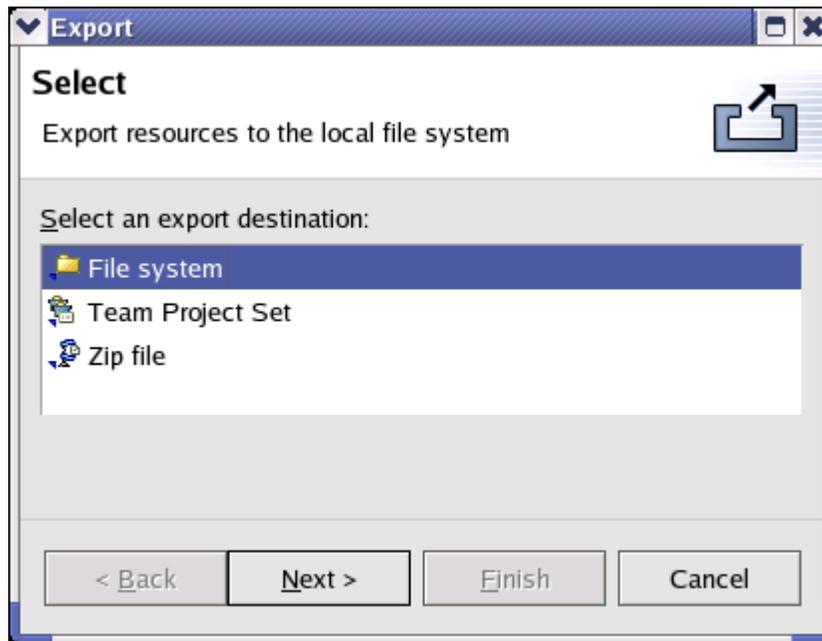
If you created a Managed Make C Project, you can use Eclipse* to build a makefile that includes Intel compiler options. See Setting Properties. When your project is complete, you can export your makefile and project source files to another directory, then build your project from the command line using make.

Exporting makefiles

To export your makefile:

1. Select your project in the Eclipse **C/C++ Projects** view.
2. From the Eclipse **File** menu, select **Export** to launch the Export Wizard.

- On the **Select** dialog of the Export Wizard, select **File system**, then click **Next**.

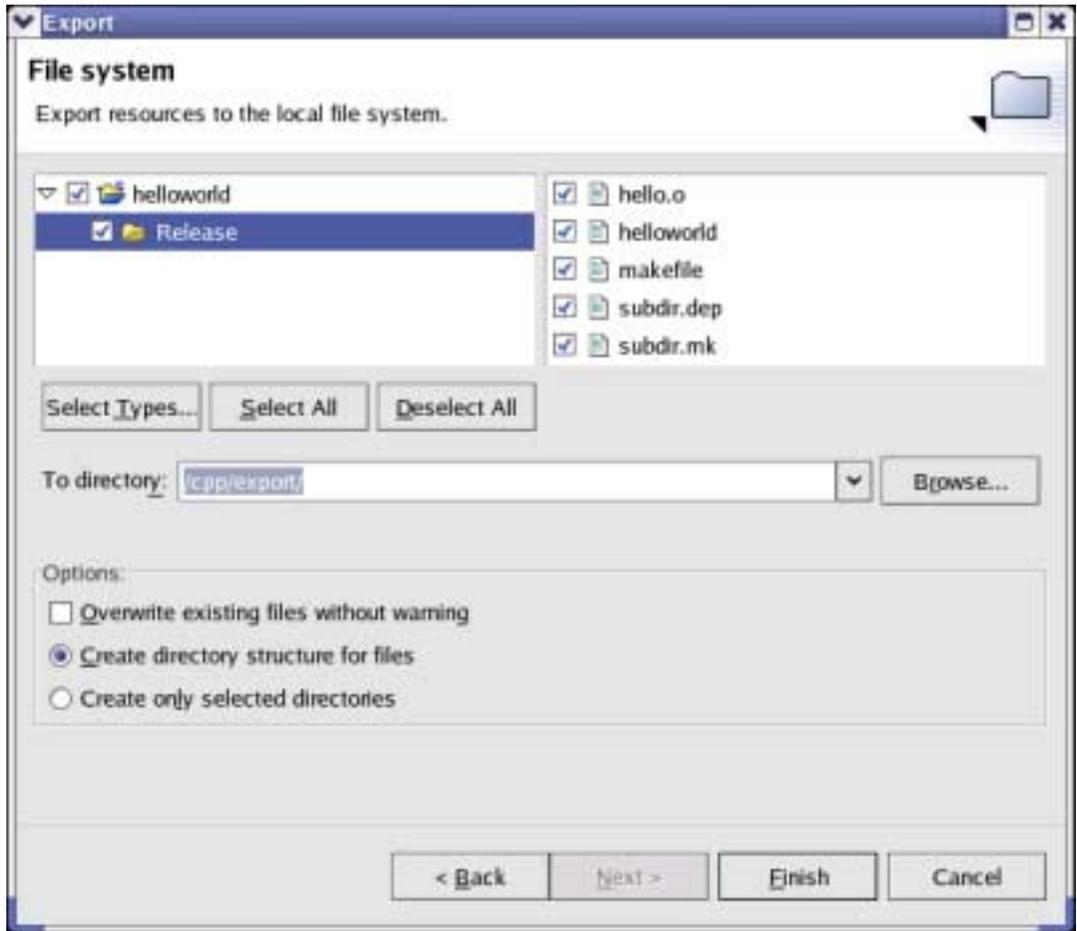


- On the **File system** dialog, check both the **helloworld** and **Release** directories in the left-hand pane. Be sure all the project sources in the right-hand pane are also checked.

 **Note**

You may deselect some files in the right-hand pane, such as the `hello.o` object file and `helloworld` executable. However, you must also select **Create directory structure for files** in the **Options** section to successfully create the export directory. This also applies to project files in the `helloworld` directory.

- Use the **Browse** button to target the export to an existing directory. Eclipse can also create a new directory for full paths entered in the **To directory** text box. If, for example, you specified `/cpp/export` as the export directory, Eclipse creates two new sub-directories:
 - `/cpp/export/helloworld`
 - `/cpp/export/helloworld/Release`



- Click **Finish** to complete the export.

Running make

In a terminal window, change to the `/cpp/export/helloworld/Release` directory, then run `make` by typing:

```
make clean all
```

You should see the following output:

```
rm -rf hello.o    helloworld
icc -O2 -w1 -Ob1 -tpp7 -unroll -par_threshold75 -wn100 -Zp16 -c
-o hello.o ../hello.c
icc -o helloworld hello.o
```

Compilation Options

This section describes the Intel® C++ Compiler options that determine the compilation process and output. By default, the compiler converts source code directly to an executable file. Appropriate options allow you to control the process by directing the compiler to produce:

- Preprocessed files (.i) with the `-P` option.
- Assembly files (.s) with the `-S` option.
- Object files (.o) with the `-c` option.
- Executable files (.out) by default.

You can also name the output file or designate a set of options that are passed to the linker. If you specify a phase-limiting option, the compiler produces a separate output file representing the output of the last phase that completes for each primary input file.

Preprocessor Options

This section describes the options you can use to direct the operations of the preprocessor. Preprocessing performs such tasks as macro substitution, conditional compilation, and file inclusion.

Option	Description
<code>-Aname [(values, ...)]</code>	Associates a symbol <i>name</i> with the specified sequence of <i>values</i> . Equivalent to an <code>#assert</code> preprocessing directive.
<code>-A-</code>	Causes all predefined macros and assertions to be inactive.
<code>-C</code>	Preserves comments in preprocessed source output.
<code>-Dname [(value)]</code>	Defines the macro <i>name</i> and associates it with the specified <i>value</i> . The default (<code>-Dname</code>) defines a macro with a <i>value</i> of 1.
<code>-E</code>	Directs the preprocessor to expand your source module and write the result to standard output.
<code>-EP</code>	Directs the preprocessor to expand your source module and write the result to standard output. Does not include <code>#line</code> directives in the output.
<code>-P</code>	Directs the preprocessor to expand your source module and store the result in a .i file in the current directory.
<code>-Uname</code>	Suppresses any automatic definition for the specified macro <i>name</i> .

Preprocessing Only

Use the `-E`, `-P` or `-EP` option to preprocess your source files without compiling them. When using these options, only the preprocessing phase of compilation is activated.

Using `-E`

When you specify the `-E` option, the compiler's preprocessor expands your source module and writes the result to `stdout`. The preprocessed source contains `#line` directives, which the compiler uses to determine the source file and line number. For example, to preprocess two source files and write them to `stdout`, enter the following command:

```
prompt>icpc -E prog1.cpp prog2.cpp
```

Using `-P`

When you specify the `-P` option, the preprocessor expands your source module and directs the output to a `.i` file instead of `stdout`. Unlike the `-E` option, the output from `-P` does not include `#line` number directives. By default, the preprocessor creates the name of the output file using the prefix of the source file name with a `.i` extension. You can change this by using the `-ofile` option. For example, the following command creates two files named `prog1.i` and `prog2.i`, which you can use as input to another compilation:

```
prompt>icpc -P prog1.cpp prog2.cpp
```



Caution

When you use the `-P` option, any existing files with the same name and extension are overwritten.

Using `-EP`

Using the `-EP` option directs the preprocessor to not include `#line` directives in the output. `-EP` is equivalent to `-E -P`.

```
prompt>icpc -EP prog1.cpp prog2.cpp
```

Preserving Comments in Preprocessed Source Output

Use the `-C` option to preserve comments in your preprocessed source output. Comments following preprocessing directives, however, are not preserved.

Preprocessing Directive Equivalents

You can use the `-A`, `-D`, and `-U` options as equivalents to preprocessing directives:

- `-A` is equivalent to a `#assert` preprocessing directive
- `-D` is equivalent to a `#define` preprocessing directive
- `-U` is equivalent to a `#undef` preprocessing directive

Using -A

Use the `-A` option to make an assertion. **Syntax:** `-Aname [(value)]`.

Argument	Description
<i>name</i>	Indicates an identifier for the assertion
<i>value</i>	Indicates a <i>value</i> for the assertion. If a <i>value</i> is specified, it should be quoted, along with the parentheses delimiting it.

For example, to make an assertion for the identifier `fruit` with the associated values `orange` and `banana` use the following command:

```
prompt>icpc -A"fruit(orange,banana)" prog1.cpp
```

Using -D

Use the `-D` option to define a macro. **Syntax:** `-Dname [=value]`.

Argument	Description
<i>name</i>	The name of the macro to define.
<i>value</i>	Indicates a value to be substituted for name. If you do not enter a value, name is set to 1. The value should be quoted if it contains non-alphanumerics.

For example, to define a macro called `SIZE` with the value `100` use the following command:

```
prompt>icpc -DSIZE=100 prog1.cpp
```

The `-D` option can also be used to define functions. For example:

```
prompt>icpc -D"f(x)=x" prog1.cpp
```

Using -U

Use the `-U` option to remove (undefine) a pre-defined macro. **Syntax:** `-Uname`.

Argument	Description
<i>name</i>	The name of the macro to undefine.



Note

If you use `-D` and `-U` in the same compilation, the compiler processes the `-D` option before `-U`, rather than processing them in the order they appear on the command line.

Predefined Macros

The Intel® C++ Compiler supports the predefined macros listed in the following table. The compiler also includes predefined macros specified by the ISO/ANSI standard. See Conformance to the C Standard.

Macro Name	Value	Architecture
<code>__BASE_FILE__</code>	Name of source file	Both
<code>__cplusplus</code>	1	Both
<code>__EDG__</code>	1	Both
<code>__EDG_VERSION__</code>	303	Both
<code>__ELF__</code>	1	Both
<code>__EXCEPTIONS</code>	Defined when <code>-fno-exceptions</code> is not used.	IA-32 only
<code>__GNUC__</code>	2 - if gcc version is less than 3.2 3 - if gcc version is 3.2, 3.3, or 3.4	Both
<code>__gnu_linux__</code>	1	Both
<code>__GNUC_MINOR__</code>	95 - if gcc version is less than 3.2 2 - if gcc version is 3.2 3 - if gcc version is 3.3 4 - if gcc version is 3.4	Both
<code>__GNUC_PATCHLEVEL__</code>	0	Both
<code>__GXX_ABI_VERSION</code>	102	Both
<code>__i386</code>	1	IA-32 only
<code>__i386__</code>	1	IA-32 only
<code>i386</code>	1	IA-32 only
<code>__ia64</code>	1	Itanium architecture only
<code>__ia64__</code>	1	Itanium architecture only
<code>ia64</code>	1	Itanium architecture only
<code>__INTEL_COMPILER</code>	810	Both
<code>__INTEL_COMPILER_BUILD_DATE</code>	YYYYMMDD	Both
<code>__INTEL_CXXLIB_ICC</code>	1 when <code>-cxxlib_icc</code> option is specified during compilation.	Both

Macro Name	Value	Architecture
__INTEL_RTTI__	1 when -fno-rtti is not specified.	Both
__INTEL_STRICT_ANSI__	1 when -strict_ansi is specified.	Both
_INTEGRAL_MAX_BITS	64	Itanium architecture only
__itanium__	1	Itanium architecture only
__linux	1	Both
__linux__	1	Both
linux	1	Both
__LONG_DOUBLE_SIZE__	80	IA-32 only
__LONG_MAX__	9223372036854775807L	Itanium architecture only
__lp64	1	Itanium architecture only
__LP64__	1	Itanium architecture only
_LP64	1	Itanium architecture only
__NO_INLINE__	1	Both
__NO_MATH_INLINES	1	Both
__NO_STRING_INLINES	1	Both
__OPTIMIZE__	1	Both
__PIC__	1 when -fPIC is used.	Both
__pic__	1 when -fPIC is used.	Both
_PGO_INSTRUMENT	1 when -prof_gen[x] is used.	Both
__PTRDIFF_TYPE__	int on IA-32 long on Itanium architecture	Both
__REGISTER_PREFIX__	(no value)	Both
__SIGNED_CHARS__	1	Both

Macro Name	Value	Architecture
<code>__SIZE_TYPE__</code>	unsigned on IA-32 unsigned long on Itanium architecture	Both
<code>__unix</code>	1	Both
<code>__unix__</code>	1	Both
<code>unix</code>	1	Both
<code>__USER_LABEL_PREFIX__</code>	(no value)	Both
<code>__VERSION__</code>	"Intel® C++ gcc 3.0 mode"	Both
<code>__WCHAR_T</code>	1	Both
<code>__WCHAR_TYPE__</code>	long int on IA-32 int on Itanium architecture	Both
<code>__WINT_TYPE__</code>	unsigned int	Both

Suppress Macro Definition

Use the `-Uname` option to suppress any macro definition currently in effect for the specified *name*. The `-U` option performs the same function as an `#undef` preprocessor directive. You can use the `-no-gcc` option to disable the `__GNU_C_MINOR__`, `__GNU_C_MINOR__`, and `__GNU_C_PATCHLEVEL__` macros.

Customizing the Compilation Environment

For IA-32 and the Intel® Itanium® architecture, you will need to set a compilation environment. To customize the environment used during compilation, you can specify:

- Environment Variables -- the paths where the compiler and other tools can search for specific files.
- Configuration Files -- the options to use with each compilation.
- Response Files -- the options and files to use for individual projects.
- Include Files -- the names and locations of source header files.

Environment Variables

You can customize your environment by specifying paths where the compiler can search for special files such as libraries and include files.

- `LD_LIBRARY_PATH` -- specifies the location for shared objects.
- `PATH` -- specifies the directories the system searches for binary executable files.
- `ICCCFG` -- specifies the configuration file for customizing compilations when invoking the compiler using `icc`.
- `ICPCCFG` -- specifies the configuration file for customizing compilations when invoking the compiler using `icpc`.
- Several environment variables are supported to specify the location for temporary files. The compiler searches for the following variables in the order specified: `TMP`, `TMPDIR`, and `TEMP`. If none of these variables are found, temporary files are stored in `/tmp`.
- `IA32ROOT` (IA32-based systems) -- points to the directory containing the `bin`, `lib`, `include` and substitute header directories.
- `IA64ROOT` (Itanium®-based systems) -- points to the directory containing the `bin`, `lib`, `include` and substitute header directories.

GNU* Environment Variables

The Intel C++ Compiler supports the following GNU environment variables:

- `CPATH` -- Path to include directory for C/C++ compilations
- `C_INCLUDE_PATH` -- Path include directory for C compilations
- `CPLUS_INCLUDE_PATH` -- Path include directory for C++ compilations.
- `LIBRARY_PATH` -- The value of `LIBRARY_PATH` is a colon-separated list of directories, much like `PATH`.
- `DEPENDENCIES_OUTPUT` -- If this variable is set, its value specifies how to output dependencies for Make based on the non-system header files processed by the compiler. System header files are ignored in the dependency output.
- `SUNPRO_DEPENDENCIES` -- This variable is the same as `DEPENDENCIES_OUTPUT`, except that system header files are not ignored.

Compilation Environment Options

The Intel® C++ Compiler installation includes shell scripts that you can use to set environment variables. See *Invoking the Compiler from the Command Line* for more information.

Configuration Files

You can decrease the time you spend entering command-line options and ensure consistency by using the configuration file to automate often-used command-line entries. You can insert any valid command-line option into the configuration file. The compiler processes options in the configuration file in the order they appear followed by the command-line options that you specify when you invoke the compiler.



Note

Options in the configuration file will be executed every time you run the compiler. If you have varying option requirements for different projects, see *Response Files*.

How to Use Configuration Files

The following example illustrates a basic configuration file. After you have written the `.cfg` file, simply ensure it is in the same directory as the compiler's executable file when you run the compiler. The text following the pound (`#`) character is recognized as a comment. The configuration file is `icc.cfg`.

```
## Sample configuration file.
## Define preprocessor macro MY_PROJECT.

-DMY_PROJECT

## Additional directories to be searched
## for INCLUDE files, before the default.

-I /project/include
```

Specifying the Location with ICCCFG

You can use the `ICCCFG` environment variable to specify the location of your configuration file:

```
ICCCFG=/cpp/config/my_options.cfg
```

Each time you invoke the compiler with `icc`, `my_options.cfg` is used as your configuration file. The `ICPCCFG` environment variable is supported for invoking the compiler with `icpc`.

See Environment Variables.

Response Files

Use response files to specify options used during particular compilations. Response files are invoked as an option on the command line. Options in a response file are inserted in the command line at the point where the response file is invoked.

Sample Response Files

```
# response file: response1.txt
# compile with these options

-axP
-pch

# end of response1 file
```

```
# response file: response2.txt
# compile with these options

-mp1
-strict_ansi

# end of response2 file
```

Use response files to decrease the time spent entering command-line options and to ensure consistency by automating command-line entries. Use individual response files to maintain options for specific projects to avoid editing the configuration file when changing projects.

Any number of options or file names can be placed on a line in the response file. Several response files can be referenced in the same command line.

The syntax for using response files is as follows:

```
prompt>icpc @response1.txt source1.cpp @response2.txt source2.cpp
```



Note

An "at" sign (@) must precede the name of the response file on the command line.

Include Files

Include directories are searched in the default system areas and whatever is specified by the `-Idirectory` option. For multiple search directories, multiple `-Idirectory` commands must be used. The compiler searches directories for include files in the following order:

- directory of the source file that contains the include
- directories specified by the `-I` option

How to Remove Include Directories

Use the `-X` option to prevent the compiler from searching the default system areas. You can use the `-X` option with the `-I` option to prevent the compiler from searching the default path for include files and direct it to use an alternate path.

For example, to direct the compiler to search the path `/alt/include` instead of the default path, do the following:

```
prompt>icpc -X -I/alt/include prog.cpp
```

See also Searching for Include Files.

Searching for Include Files

By default, the compiler searches for the standard include files in the directories specified in the `CPATH`, `C_INCLUDE_PATH`, and `CPLUS_INCLUDE_PATH` environment variables. You can indicate the location of include files in the configuration file.

How to Specify an Include Directory

Use the `-Idirectory` option to specify an additional directory in which to search for include files. For multiple search directories, multiple `-Idirectory` commands must be used. Included files are brought into the program with a `#include` preprocessor directive. The compiler searches directories for include files in the following order:

- directory of the source file that contains the include
- directories specified by the `-I` option
- directories specified in the `CPATH`, `C_INCLUDE_PATH`, and `CPLUS_INCLUDE_PATH` environment variables

How to Remove Include Directories

Use the `-X` option to prevent the compiler from searching the default path specified by the environment variables.

You can use the `-X` option with the `-I` option to prevent the compiler from searching the default path for include files and direct it to use an alternate path.

For example, to direct the compiler to search the path `/alt/include` instead of the default path, do the following:

```
prompt>icpc -X -I/alt/include source.cpp
```

Controlling Compilation

If no errors occur during processing, you can use the output files from a particular phase as input to a subsequent compiler invocation. The following table describes the options to control the output:

Option	Input	Output
<code>-P</code>	<ul style="list-style-type: none"> • Source files 	Preprocessed files (<code>.i</code> files).
<code>-E</code>	<ul style="list-style-type: none"> • Source files 	Preprocesses source file and directs output to <code>stdout</code> .
<code>-EP</code>	<ul style="list-style-type: none"> • Source files 	Preprocesses source file, directs output to <code>stdout</code> , and omits line numbers.
<code>-c</code>	<ul style="list-style-type: none"> • Source files • Preprocessed files 	Compile to object only (<code>.o</code>), do not link.
<code>-S</code>	<ul style="list-style-type: none"> • Source files • Preprocessed files 	Generate assemblable files with <code>.s</code> suffix and stops the compilation process.

Option	Input	Output
-syntax	<ul style="list-style-type: none"> • Source files • Preprocessed files 	Emits diagnostic list of syntax errors to <code>sdtout</code> . There is no output for source files free of syntax errors.
(Default)	<ul style="list-style-type: none"> • Source files • Preprocessed files • Assemblable files • Object files • Libraries 	Executable file (<code>.out</code> files).

Controlling Compilation Flow

Option	Description
-c	Stops the compilation process after an object file has been generated. The compiler generates an object file for each C or C++ source file or preprocessed source file. Also takes an assembler file and invokes the assembler to generate an object file.
-Kpic, -KPIC	Generate position-independent code.
-lname	Link with a library indicated in <i>name</i> .
-nobss_init	Places variables that are initialized with zeroes in the DATA section.
-P, -F	Stops the compilation process after C or C++ source files have been preprocessed and writes the results to files named according to the compiler's default file-naming conventions.
-S	Generates assemblable file only (with <code>.s</code> suffix), then stops the compilation.
-sox[-]	Enables [disables] the saving of compiler options and version information in the executable file. Default is <code>-sox-</code> .
-Zp{1 2 4 8 16}	Packs structures on 1, 2, 4, 8, or 16 byte boundaries.

Controlling Compilation Output

Option	Description
- <i>oname</i>	Produces an assembly file with the specified file <i>name</i> , or the default file name if <i>name</i> is not specified.
-S	Generates assemblable file only (with <code>.s</code> suffix), then stops the compilation.

Specifying Alternate Tools and Paths

You can direct the compiler to specify alternate tools for preprocessing, compilation, assembly, and linking. Further, you can invoke options specific to your alternate tools on the command line. The following sections explain how to use `-Qlocation` and `-Qoption` to do this.

How to Specify an Alternate Component

Use `-Qlocation` to specify an alternate path for a tool. This option accepts two arguments using the following syntax:

```
prompt>icpc -Qlocation,tool,path
```

tool	Description
cpp	Specifies the compiler front-end preprocessor.
c	Specifies the C++ compiler.
asm	Specifies the assembler.
ld	Specifies the linker.
gas	Specifies the GNU assembler.
gld	Specifies the GNU linker.

path is the complete path to the tool.

How to Pass Options to Other Programs

Use `-Qoption` to pass an option specified by *optlist* to a *tool*, where *optlist* is a comma-separated list of options. The syntax for this command is the following:

```
prompt>icpc -Qoption,tool,optlist
```

tool	Description
cpp	Specifies the compiler front-end preprocessor.
c	Specifies the C++ compiler.
asm	Specifies the assembler.
ld	Specifies the linker.

optlist indicates one or more valid argument strings for the designated program. If the argument is a command-line option, you must include the hyphen. If the argument contains a space or tab character, you must enclose the entire argument in quotation characters ("). You must separate multiple arguments with commas. The following example directs the linker to create a memory map when the compiler produces the executable file from the source.

```
prompt>icpc -Qoption,link,-map,proto.map proto.cpp
```

The `-Qoption,link` option in the preceding example is passing the `-map` option to the linker. This is an explicit way to pass arguments to other tools in the compilation process. Also, you can use the `-Xlinker val` to pass values (*val*) to the linker.

Monitoring Data Settings

The options described here provide monitoring of Intel compiler-generated code.

Specifying Structure Tag Alignments

You can specify an alignment constraint for structures and unions in two ways:

- Place a pack pragma in your source file, or
- Enter the alignment option on the command line

Both specifications change structure tag alignment constraints.

Flushing Denormal Values to Zero for Itanium-based Systems Only

Option `-ftz` flushes denormal results to zero when the application is in the gradual underflow mode. Use this option if the denormal values are not critical to application behavior. Flushing the denormal values to zero with `-ftz` may improve performance of your application. The default status of `-ftz` is OFF. By default, the compiler lets results gradually underflow.

The `-ftz` switch only needs to be used on the source containing function `main()`. The effect of the `-ftz` switch is to turn on FTZ mode for the process started by `main()`. The initial thread and any threads subsequently created by that process will operate in FTZ mode.



Note

The `-O3` option turns `-ftz` ON. Use `-ftz-` to disable flushing denormal results to zero.

Allocation of Zero-initialized Variables

By default, variables explicitly initialized with zeros are placed in the BSS section. But using the `-nobss_init` option, you can place any variables that are explicitly initialized with zeros in the DATA section if required.

Precompiled Header Files

The Intel® C++ Compiler supports precompiled header (PCH) files to significantly reduce compile times using the following options:

- `-pch`
- `-create_pch filename`
- `-use_pch filename`
- `-pch_dir dirname`



Caution

Depending on how you organize the header files listed in your sources, these options may increase compile times. See Organizing Source Files to learn how to optimize compile times using the PCH options.

-pch

The `-pch` option directs the compiler to use appropriate PCH files. If none are available, they are created as `sourcefile.pchi`. This option supports multiple source files, such as the ones shown in Example 1:

Example 1 command line:

```
prompt>icpc -pch source1.cpp source2.cpp
```

Example 1 output when .pch files do not exist:

```
"source1.cpp": creating precompiled header file "source1.pchi"
"source2.cpp": creating precompiled header file "source2.pchi"
```

Example 1 output when .pch files do exist:

```
"source1.cpp": using precompiled header file "source1.pchi"
"source2.cpp": using precompiled header file "source2.pchi"
```



Note

The `-pch` option will use PCH files created from other sources if the headers files are the same. For example, if you compile `source1.cpp` using `-pch`, then `source1.pchi` is created. If you then compile `source2.cpp` using `-pch`, the compiler will use `source1.pchi` if it detects the same headers.

-create_pch

Use the `-create_pch filename` option if you want the compiler to create a PCH file called `filename`. Note the following regarding this option:

- The `filename` parameter must be specified.
- The `filename` parameter can be a full path name.
- The full path to `filename` must exist.
- The `.pch` extension is not automatically appended to `filename`.
- This option cannot be used in the same compilation as `-use_pch filename`.
- The `-create_pch filename` option is supported for single source file compilations only.

Example 2 command line:

```
prompt>icpc -create_pch /pch/source32.pchi source.cpp
```

Example 2 output:

```
"source.cpp": creating precompiled header file "/pch/source32.pchi"
```

-use_pch filename

This option directs the compiler to use the PCH file specified by *filename*. It cannot be used in the same compilation as `-create_pch filename`. The `-use_pch filename` option supports full path names and supports multiple source files when all source files use the same `.pchi` file.

Example 3 command line:

```
prompt>icpc -use_pch /pch/source32.pchi source.cpp
```

Example 3 output:

```
"source.cpp": using precompiled header file /pch/source32.pchi
```

-pch_dir dirname

Use the `-pch_dir dirname` option to specify the path (*dirname*) to the PCH file. You can use this option with `-pch`, `-create_pch filename`, and `-use_pch filename`.

Example 4 command line:

```
prompt>icpc -pch -pch_dir /pch source32.cpp
```

Example 4 output:

```
"source32.cpp": creating precompiled header file /pch/source32.pchi
```

Organizing Source Files

If many of your source files include a common set of header files, place the common headers first, followed by the `#pragma hdrstop` directive. This pragma instructs the compiler to stop generating PCH files. For example, if `source1.cpp`, `source2.cpp`, and `source3.cpp` all include `common.h`, then place `#pragma hdrstop` after `common.h` to optimize compile times.

```
#include "common.h"
#pragma hdrstop
#include "noncommon.h"
```

When you compile using the `-pch` option:

```
prompt>icpc -pch source1.cpp source2.cpp source3.cpp
```

the compiler will generate one PCH file for all three source files:

```
"source1.cpp": creating precompiled header file "source1.pchi"
```

```
"source2.cpp": using precompiled header file "source1.pchi"
```

```
"source3.cpp": using precompiled header file "source1.pchi"
```

If you don't use `#pragma hdrstop`, a different PCH file is created for each source file if different headers follow `common.h`, and the subsequent compile times will be longer. `#pragma hdrstop` has no effect on compilations that do not use these PCH options.

Linking

This topic describes the options that let you control and customize the linking with tools and libraries and define the output of the `ld` linker. See the `ld` man page for more information on the linker.

Option	Description
<code>-Ldirectory</code>	Instruct the linker to search <i>directory</i> for libraries.
<code>-Qoption,tool,list</code>	Passes an argument list to another program in the compilation sequence, such as the assembler or linker.
<code>-shared</code>	Instructs the compiler to build a Dynamic Shared Object (DSO) instead of an executable.
<code>-shared-libcxa</code>	<code>-shared-libcxa</code> has the opposite effect of <code>-static-libcxa</code> . When it is used, the Intel-provided <code>libcxa</code> C++ library is linked in dynamically, allowing the user to override the static linking behavior when the <code>-static</code> option is used. Note: By default, all C++ standard and support libraries are linked dynamically.
<code>-i_dynamic</code>	Specifies that all Intel-provided libraries should be linked dynamically.
<code>-static</code>	Causes the executable to link all libraries statically, as opposed to dynamically. When <code>-static</code> is not used: <ul style="list-style-type: none"> • <code>/lib/ld-linux.so.2</code> is linked in • all other libs are linked dynamically When <code>-static</code> is used: <ul style="list-style-type: none"> • <code>/lib/ld-linux.so.2</code> is not linked in • all other libs are linked statically
<code>-static-libcxa</code>	By default, the Intel-provided <code>libcxa</code> C++ library is linked in dynamically. Use <code>-static-libcxa</code> on the command line to link <code>libcxa</code> statically, while still allowing the standard libraries to be linked in by the default behavior.
<code>-Bstatic</code>	This option is placed in the linker command line corresponding to its location on the user command line. This option is used to control the linking behavior of any library being passed in via the command line.

Option	Description
-Bdynamic	This option is placed in the linker command line corresponding to its location on the user command line. This option is used to control the linking behavior of any library being passed in via the command line.

Suppressing Linking

Use the `-c` option to suppress linking. For example, entering the following command produces the object files `file1.o` and `file2.o`:

```
prompt>icpc -c file1.cpp file2.cpp
```



Note

The preceding command does not link these files to produce an executable file.

Debugging

This section describes the basic command-line options that you can use as tools to debug your compilation and to display and check compilation errors. This section includes topics on:

- Preparing for Debugging
- Parsing for Syntax Only
- Optimizations and Debugging
- Options for Debug Information

Debuggers

The Intel® Debugger is included with the Intel® C++ Compiler, but installation is optional. The Intel Debugger (`idb`) includes an environment script, `idbvars.sh`, which you need to run before executing `idb`:

```
prompt>source /opt/intel_idb_80/bin/idbvars.sh
```

See the "*Intel Debugger (IDB) Manual*" (`idb_debugger_manual.htm`) for complete information on using the Intel Debugger.

You can also use the GNU Debugger (`gdb`) to debug programs compiled with the Intel C++ Compiler.

Preparing for Debugging

Use the `-g` option to direct the compiler to generate code to support symbolic debugging. For example:

```
prompt>icpc -g prog.cpp
```

The compiler does not support the generation of debugging information in assembly files. If you specify the `-g` option, the resulting object file will contain debugging information, but the assembly file will not.



Note

The `-g` option changes the default optimization from `-O2` to `-O0`.

Parsing for Syntax Only

Use the `-syntax` option to stop processing source files after they have been parsed for C++ language errors. This option provides a method to quickly check whether sources are syntactically and semantically correct. The compiler creates no output file. In the following example, the compiler checks `prog.cpp` and displays diagnostic information to the standard error output:

```
prompt>icpc -syntax prog.cpp
```

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:

Option	Description
<code>-O0</code>	Disables optimizations. Enables the <code>-fp</code> option.
<code>-g</code>	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>-O1</code> , <code>-O2</code> , or <code>-O3</code> is explicitly specified in the command line together with <code>-g</code> .
<code>-fp</code> IA-32 only	Disable using the EBP register as general purpose register.
Option	Effect on <code>-fp</code>
<code>-O1</code> , <code>-O2</code> , or <code>-O3</code>	Disables <code>-fp</code> .
<code>-O0</code>	Enables <code>-fp</code> .

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 `-O1`, `-O2`, or `-O3` is specified on the command line along with `-g`, which produces symbolic debug information in the object file.

Note that if you specify the `-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 `-O1`, `-O2`, or `-O3` is not specified. In this case `-g` turns on `-O0` which is what slows the program down. If both `-O2` and `-g` are specified, the code should run nearly the same speed as if `-g` were not specified.

Refer to the following table 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 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.

Options for Debug Information

The Intel® C++ Compiler provides basic debugging information and new features for enhanced debugging of optimized code. The basic debugging switches are listed in the following table.

Option	Description
<code>-debug all</code> <code>-debug full</code>	These options are equivalent to <code>-g</code> . They turn on production of basic debug information. They are off by default.
<code>-debug none</code>	This option turns off production of debug information. This option is on by default.

The Intel C++ Compiler improves debuggability of optimized code through enhanced support for:

- tracebacks
- variable locations
- breakpoints and stepping

The options described in the following table control emission of enhanced debug information. They must be used in conjunction with the `-g` option.

Option	Description
<code>-debug inline_info</code>	This option produces enhanced source position information for inlined code. This leads to greater accuracy when reporting the source location of any instruction. It also provides more information to debuggers for function call traceback. The Intel debugger, <code>idb</code> , has been enhanced to use the richer debug information to show simulated call frames for inlined functions. This option is off by default.
<code>-debug variable_locations</code>	This option produces additional debug information for scalar local variables using a feature of the DWARF object module format known as "location lists." The runtime locations of local scalar variables are specified more accurately using this feature, i.e. whether at a given position in the code, a variable value is found in memory or a machine register. The Intel debugger is able to process location lists and display values of local variables at runtime with improved accuracy. This option is off by default.
<code>-debug extended</code>	This option turns on the <code>-debug</code> options described previously: <ul style="list-style-type: none"> • <code>-debug inline_info</code> • <code>-debug variable_locations</code>

**Note**

When the compiler needs to choose between optimization and quality of debug information, optimization is given priority.

Creating and Using Libraries

The Intel® C++ Compiler uses the GNU* C Library, Dinkumware* C++ Library, and the Standard C++ Library. These libraries are documented at the following Internet locations:

- GNU C Library -- <http://www.gnu.org/software/libc/manual/>
- Dinkumware C++ Library -- http://www.dinkumware.com/htm_cpl/lib_cpp.html
- Standard C++ Library -- <http://gcc.gnu.org/onlinedocs/libstdc++/documentation.html>

Creating Libraries

Libraries are simply an indexed collection of object files that are included as needed in a linked program. Combining object files into a library makes it easy to distribute your code without disclosing the source. It also reduces the number of command-line entries needed to compile your project.

Static Libraries

Executables generated using static libraries are no different than executables generated from individual source or object files. Static libraries are not required at runtime, so you do not need to include them when you distribute your executable. At compile time, linking to a static library is generally faster than linking to individual source files.

To build a static library:

1. use the `-c` option to generate object files from the source files

```
prompt>icpc -c my_source1.cpp my_source2.cpp my_source3.cpp
```
2. use the GNU `ar` tool to create the library file from the object files

```
prompt>ar rc my_lib.a my_source1.o my_source2.o my_source3.o
```
3. compile and link your project with your new library

```
prompt>icpc main.cpp my_lib.a
```

If your library file and source files are in different directories, use the `-Ldir` option to indicate where your library is located:

```
prompt>icpc -L/cpp/libs main.cpp my_lib.a
```

If you are using Interprocedural Optimization, see [Creating a Library from IPO Objects using `xiar`](#).

Shared Libraries

Shared libraries, also referred to as dynamic libraries or Dynamic Shared Objects (DSO), are linked differently than static libraries. At compile time, the linker insures that all the necessary symbols are either linked into the executable, or can be linked at runtime from the shared library. Executables compiled from shared libraries are smaller, but the shared libraries must be included with the executable to function correctly. When multiple programs use the same shared library, only one copy of the library is required in memory.

To build a shared library:

1. use the `-fPIC` and `-c` options to generate object files from the source files

```
prompt>icpc -fPIC -c my_source1.cpp my_source2.cpp
my_source3.cpp
```
2. use the `-shared` option to create the library file from the object files

```
prompt>icpc -shared my_lib.so my_source1.o my_source2.o
my_source3.o
```
3. compile and link your project with your new library

```
prompt>icpc main.cpp my_lib.so
```

See also Intel® Shared Libraries and Compiling for Non-shared Libraries.

Default Libraries

The following libraries are supplied with the Intel® C++ Compiler:

Library	Description
libguide.a libguide.so	For OpenMP* implementation
libguide_stats.a libguide_stats.so	OpenMP static library for the parallelizer tool with performance statistics and profile information
libompstub.a	Library that resolves references to OpenMP subroutines when OpenMP is not in use
libsvml.a	Short vector math library
libirc.a	Intel support library for PGO and CPU dispatch
libimf.a libimf.so	Intel math library
libcprts.a libcprts.so libcprts.so.5	Dinkumware* C++ Library
libunwind.a libunwind.so libunwind.so.5	Unwinder library
libcxa.a libcxa.so libcxa.so.5	Intel run time support for C++ features
libcxaguard.a libcxaguard.so libcxaguard.so.5	Used for interoperability support with the <code>-cxxlib-gcc</code> option. See gcc Interoperability.

When you invoke the `-cxxlib-gcc` option, the following replacements occur:

- `libcprts` is replaced with `libstdc++` from the `gcc*` distribution (3.2 or newer)
- `libcxa` and `libunwind` are replaced by `libgcc` from the `gcc` distribution (3.2 or newer)



Caution

The Linux* system libraries and the compiler libraries are not built with the `-align` option. Therefore, if you compile with the `-align` option and make a call to a compiler distributed or system library, and have `long long`, `double`, or `long double` types in your interface, you will get the wrong answer due to the difference in alignment. Any code built with `-align` cannot make calls to libraries that use these types in their interfaces unless they are built with `-align` (in which case they will not work without `-align`).

Math Libraries

The Intel math library, `libimf.a`, contains optimized versions of math functions found in the standard C run-time library. The functions in `libimf.a` are optimized for program execution speed on Intel® Pentium® III and Pentium 4 processors. The Itanium® compiler also includes a `libimf.a` designed to optimize performance on Itanium-based systems. The Intel math library is linked by default.

See [Managing Libraries and Intel Math Library](#).

Intel® Shared Libraries

By default, the Intel® C++ Compiler links Intel-provided C++ libraries dynamically. The GNU* and Linux* system libraries are also linked dynamically.

Options for Shared Libraries

Option	Description
<code>-i_dynamic</code>	Use the <code>-i_dynamic</code> option to link Intel-provided C++ libraries dynamically (default). This has the advantage of reducing the size of the application binary, but it also requires the libraries to be on the systems where the application runs.
<code>-shared</code>	The <code>-shared</code> option instructs the compiler to build a Dynamic Shared Object (DSO) instead of an executable. For more details, refer to the <code>ld</code> man page documentation.
<code>-fpic</code>	Use the <code>-fpic</code> option when building shared libraries. It is required for the compilation of each object file included in the shared library.

See also [Linking](#).

Managing Libraries

The `LD_LIBRARY_PATH` environment variable contains a colon-separated list of directories in which the linker will search for library (.a) files. If you want the linker to search additional libraries, you can add their names to `LD_LIBRARY_PATH`, to the command line, or to a response file (see Note). In each case, the names of these libraries are passed to the linker before the names of the Intel libraries that the driver always specifies.



Response files are processed at the location they appear on the command line. If libraries are specified in the response file, references from object files seen after the response file will not be resolved in those libraries.

Modifying LD_LIBRARY_PATH

If you want to add a directory, `/libs` for example, to the `LD_LIBRARY_PATH`, you can do either of the following:

- command line: `prompt>export LD_LIBRARY_PATH=/libs:$LD_LIBRARY_PATH`
- startup file `export LD_LIBRARY_PATH=/libs:$LD_LIBRARY_PATH`

To compile `file.cpp` and link it with the library `mylib.a`, enter the following command:

```
prompt>icpc file.cpp mylib.a
```

The compiler passes file names to the linker in the following order:

1. the object file
2. any objects or libraries specified on the command line, in a response file, or in a configuration file
3. the Intel® Math Library, `libimf.a`

Compiling for Non-shared Libraries

This section includes information on:

- Global Symbols and Visibility Attributes
- Symbol Preemption
- Specifying Symbol Visibility Explicitly
 - Other Visibility-related Command-line Options

Global Symbols and Visibility Attributes

A global symbol is one that is visible outside the compilation unit (single source file and its include files) in which it is declared. In C/C++, this means anything declared at file level without the `static` keyword. For example:

```
int x = 5;           // global data definition
extern int y;       // global data reference
int five()          // global function definition
{ return 5; }
extern int four();  // global function reference
```

A complete program consists of a main program file and possibly one or more shareable object (.so) files that contain the definitions for data or functions referenced by the main program. Similarly, shareable objects might reference data or functions defined in other shareable objects. Shareable objects are so called because if more than one simultaneously executing process has the shareable object mapped into its virtual memory, there is only one copy of the read-only portion of the object resident in physical memory. The main program file and any shareable objects that it references are collectively called the components of the program.

Each global symbol definition or reference in a compilation unit has a visibility attribute that controls how (or if) it may be referenced from outside the component in which it is defined. There are five possible values for visibility:

- **EXTERNAL** - The compiler must treat the symbol as though it is defined in another component. For a definition, 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 reference by a function in another component).
- **INTERNAL** - The symbol cannot be referenced outside its defining component, either directly or indirectly.

Static local symbols (in C/C++, declared at file scope or elsewhere with the keyword `static`) usually have **HIDDEN** visibility--they cannot be referenced directly by other components (or, for that matter, other compilation units within the same component), but they might be referenced indirectly.



Note

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

Symbol Preemption

Sometimes you may need to use some of the functions or data items from a shareable object, but may wish to replace others with your own definitions. For example, you may want to use the standard C runtime library shareable object, `libc.so`, but to use your 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` call your definition of the routines and not the definitions present in `libc.so`. Your definition should override, or preempt, the definition within the shareable object.

This feature of shareable objects is called symbol preemption. When the runtime 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. Since the main program image is always loaded first, none of the symbols it defines will be preempted.

The possibility of symbol preemption inhibits many valuable compiler optimizations because symbols with default visibility are not bound to a memory address until runtime. 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; the GP-relative address is not known at compile time.

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

Specifying Symbol Visibility Explicitly

You can explicitly set the visibility of an individual symbol using the `visibility` attribute on a data or function declaration. For example:

```
int i __attribute__((visibility("default")));
void __attribute__((visibility("hidden"))) x () {...}
extern void y() __attribute__((visibility("protected"));
```

The `visibility` declaration attribute accepts one of the five keywords:

- external
- default
- protected
- hidden
- internal

The value of the `visibility` declaration attribute overrides the default set by the `-fvisibility`, `-fpic`, or `-fno-common` attributes.

If you have a number of symbols for which you wish to specify the same `visibility` attribute, you can set the visibility using one of the five command line options:

- `-fvisibility-external=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 white space (blanks, TAB characters, or newlines). For example, the command line option:

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

where file `prot.txt` contains:

```
a
  b c d
e
```

sets protected visibility for symbols `a`, `b`, `c`, `d`, and `e`. This has the same effect as

```
__attribute__((visibility="protected"))
```

on the declaration for each of the symbols. Note that these two ways to explicitly set visibility are mutually exclusive--you may use `__attribute__((visibility()))` on the declaration, or specify the symbol name in a file, but not both.

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

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

This option sets the visibility for symbols not specified in a visibility list file and that do not have `__attribute__((visibility()))` in their declaration. For example, the command line options:

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

where file `prot.txt` is as previously described, will cause 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.

Other Visibility-related Command-line Options

-fminshared

The `-fminshared` option specifies that the compilation unit will be part of a main program component and will not be linked as part of a shareable object. Since symbols defined in the main program cannot be preempted, this allows the compiler to treat symbols declared with default visibility as though they have protected visibility (i.e., `-fminshared` implies `-fvisibility=protected`). Also, 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

The `-fpic` option specifies full symbol preemption. Global symbol definitions as well as global symbol references get default (i.e., preemptable) visibility unless explicitly specified otherwise.

-fno-common

Normally a C/C++ file-scope declaration with no initializer and without the `extern` or `static` keyword

```
int i;
```

is represented as a common symbol. Such a symbol is treated as an external reference, except that if no other compilation unit has a global definition for the name, the linker allocates memory for it. The `-fno-common` option causes the compiler to treat what otherwise would be common symbols as global definitions and to allocate memory for the symbol at compile time. This may permit the compiler to use the more efficient GP-relative addressing mode when accessing the symbol.

gcc* Compatibility

C language object files created with the Intel® C++ Compiler are binary compatible with the GNU gcc* compiler and glibc*, the GNU C language library. You can use the Intel compiler or the gcc compiler to pass object files to the linker. However, to correctly pass the Intel libraries to the linker, use the Intel compiler. See Linking and Default Libraries for more information.

The Intel C++ Compiler provides many of the language extensions provided by the GNU C compiler, gcc, and the GNU C++ compiler, g++.

gcc Extensions to the C Language

GNU C includes several, non-standard features not found in ISO standard C. This version of the Intel C++ Compiler supports most these extensions (listed in the following table). See <http://www.gnu.org> for more information.

gcc Language Extension	Intel Support	GNU Description and Examples
Statements and Declarations in Expressions	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Statement-Exprs.html#Statement%20Exprs
Locally Declared Labels	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Local-Labels.html#Local%20Labels
Labels as Values	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Labels-as-Values.html#Labels%20as%20Values
Nested Functions	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Nested-Functions.html#Nested%20Functions
Constructing Function Calls	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Constructing-Calls.html#Constructing%20Calls
Naming an Expression's Type	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.2/gcc/Naming-Types.html#Naming%20Types
Referring to a Type with typeof	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/typeof.html#typeof
Generalized Lvalues	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Lvalues.html#Lvalues
Conditionals with Omitted Operands	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Conditionals.html#Conditionals
Double-Word Integers	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Long-Long.html#Long%20Long
Complex Numbers	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Complex.html#Complex

gcc Language Extension	Intel Support	GNU Description and Examples
Hex Floats	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Hex-Floats.html#Hex%20Floats
Arrays of Length Zero	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Zero-Length.html#Zero%20Length
Arrays of Variable Length	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Variable-Length.html#Variable%20Length
Macros with a Variable Number of Arguments.	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Variadic-Macros.html#Variadic%20Macros
Slightly Looser Rules for Escaped Newlines	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Escaped-Newlines.html#Escaped%20Newlines
String Literals with Embedded Newlines	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.3/gcc/Multi-line-Strings.html#Multi-line%20Strings
Non-Lvalue Arrays May Have Subscripts	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Subscripting.html#Subscripting
Arithmetic on void-Pointers	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Pointer-Arith.html#Pointer%20Arith
Arithmetic on Function-Pointers	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Pointer-Arith.html#Pointer%20Arith
Non-Constant Initializers	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Initializers.html#Initializers
Compound Literals	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Compound-Literals.html#Compound%20Literals
Designated Initializers	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Designated-Inits.html#Designated%20Inits
Cast to a Union Type	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Cast-to-Union.html#Cast%20to%20Union
Case Ranges	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Case-Ranges.html#Case%20Ranges
Mixed Declarations and Code	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Mixed-Declarations.html#Mixed%20Declarations
Declaring Attributes of Functions	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Function-Attributes.html#Function%20Attributes
Attribute Syntax	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Attribute-Syntax.html#Attribute%20Syntax

gcc Language Extension	Intel Support	GNU Description and Examples
Prototypes and Old-Style Function Definitions	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Function-Prototypes.html#Function%20Prototypes
C++ Style Comments	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/C---Comments.html#C++%20Comments
Dollar Signs in Identifier Names	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Dollar-Signs.html#Dollar%20Signs
ESC Character in Constants	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Character-Escapes.html#Character%20Escapes
Specifying Attributes of Variables	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Variable-Attributes.html#Variable%20Attributes
Specifying Attributes of Types	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Type-Attributes.html#Type%20Attributes
Inquiring on Alignment of Types or Variables	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Alignment.html#Alignment
Inline Function is As Fast As a Macro	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Inline.html#Inline
Assembler Instructions with C Expression Operands	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Extended-Asm.html#Extended%20Asm
Controlling Names Used in Assembler Code	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Asm-Labels.html#Asm%20Labels
Variables in Specified Registers	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Explicit-Reg-Vars.html#Explicit%20Reg%20Vars
Alternate Keywords	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Alternate-Keywords.html#Alternate%20Keywords
Incomplete enum Types	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Incomplete-Enums.html#Incomplete%20Enums
Function Names as Strings	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Function-Names.html#Function%20Names
Getting the Return or Frame Address of a Function	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Return-Address.html#Return%20Address
Using Vector Instructions Through Built-in Functions	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Vector-Extensions.html#Vector%20Extensions

gcc Language Extension	Intel Support	GNU Description and Examples
Other built-in functions provided by GCC	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Other-Builtins.html#Other%20Builtins
Built-in Functions Specific to Particular Target Machines	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Target-Builtins.html#Target%20Builtins
Pragmas Accepted by GCC	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Pragmas.html#Pragmas
Unnamed struct/union fields within structs/unions	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Unnamed-Fields.html#Unnamed%20Fields

g++* Extensions to the C++ Language

GNU C++ includes several, non-standard features not found in ISO standard C++. This version of the Intel C++ Compiler supports many of these extensions (listed in the following table). See <http://www.gnu.org> for more information.

g++ Language Extension	Intel Support	GNU Description and Examples
Minimum and Maximum operators in C++	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Min-and-Max.html#Min%20and%20Max
When is a Volatile Object Accessed?	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Volatiles.html#Volatiles
Restricting Pointer Aliasing	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Restricted-Pointers.html#Restricted%20Pointers
Vague Linkage	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Vague-Linkage.html#Vague%20Linkage
Declarations and Definitions in One Header	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/C---Interface.html#C++%20Interface
Where's the Template?	extern template supported	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Template-Instantiation.html#Template%20Instantiation
Extracting the function pointer from a bound pointer to member function	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Bound-member-functions.html#Bound%20member%20functions
C++-Specific Variable, Function, and Type Attributes	Yes	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/C---Attributes.html#C++%20Attributes
Java Exceptions	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Java-Exceptions.html#Java%20Exceptions

g++ Language Extension	Intel Support	GNU Description and Examples
Deprecated Features	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Deprecated-Features.html#Deprecated%20Features
Backwards Compatibility	No	http://gcc.gnu.org/onlinedocs/gcc-3.4.0/gcc/Backwards-Compatibility.html#Backwards%20Compatibility

**Note**

The Intel C++ Compiler supports gcc-style inline ASM if the assembler code uses AT&T* System V/386 syntax. See http://www.gnu.org/software/binutils/manual/gas-2.9.1/html_node/as_196.html for more information.

gcc* Interoperability

C++ compilers are interoperable if they can link object files and libraries generated by one compiler with object files and libraries generated by the second compiler, and the resulting executable runs successfully. The Intel® C++ Compiler has made significant improvements towards interoperability and is highly compatible with the GNU gcc* compiler. This section describes features of the Intel C++ Compiler that provide interoperability with gcc. These features include:

- Compiler Options for Interoperability
- Predefined Macros for Interoperability

See gcc Compatibility for a detailed list of compatibility features.

Compiler Options for Interoperability

The Intel® C++ Compiler options that affect gcc* interoperability include:

- `-gcc-name`
- `-gcc-version`
- `-cxxlib-gcc`
- `-cxxlib-icc`
- `-fabi-version`
- `-no-gcc` (see Predefined Macros for Interoperability)

-gcc-name option

The `-gcc-name=name` option, used with `-cxxlib-gcc`, lets you specify the location of gcc if the compiler cannot locate the gcc C++ libraries. Use this option when referencing a non-standard gcc installation.

-gcc-version option

The `-gcc-version=nnn` option provides compatible behavior with `gcc`, where `nnn` indicates the `gcc` version. The `-gcc-version` option is ON by default, and the value of `nnn` depends on the version of `gcc` installed on your system. This option selects the version of `gcc` with which you achieve ABI interoperability.

Installed Version of gcc	Default Value of -gcc-version
older than version 3.2	not set
3.2	320
3.3	330
3.4	340

-cxxlib-gcc option

The `-cxxlib-gcc[=GCC-root-dir]` option lets you to build your applications using the C++ libraries and header files included with the `gcc` compiler. They include:

- `libstdc++` standard C++ header files
- `libstdc++` standard C++ library
- `libgcc` C++ language support

Use the optional argument, `=GCC-root-dir`, to specify the top-level location for the `gcc` binaries and libraries.



Note

The Intel C++ Compiler is compatible with `gcc` 3.2, 3.3, and 3.4. The `-cxxlib-gcc` option is ON by default if you are using `gcc` 3.2, 3.3, or 3.4.

When you compile and link your application using the `-cxxlib-gcc` option, the resulting C++ object files and libraries can interoperate with C++ object files and libraries generated by `gcc` 3.2. This means that third-party C++ libraries built with `gcc` 3.2 will work with C++ code generated by the Intel Compiler.

The `-cxxlib-gcc` option can only be used on Linux distributions that include `gcc` 3.2. This is required for C++ ABI conformance.

By default, the Intel C++ Compiler uses headers and libraries included with the product. If you are linking with code compiled with `g++`, which was compiled against `gnu` C++ headers, then differences in the headers might cause incompatibilities that result in run-time errors.

If you build one shared library against the Intel C++ libraries, build a second shared library against the gnu C++ libraries, and use both libraries in a single application, you will have two C++ run-time libraries in use. Since the application might use symbols from both libraries, the following problems may occur:

- partially initialized libraries
- lost I/O operations from data put in unaccessed buffers
- other unpredictable results, such as jumbled output

The Intel C++ Compiler does not support more than one run-time library in one application.



Warning

If you successfully compile your application using more than one run-time library, the resulting program will likely be very unstable, especially when new code is linked against the shared libraries.

You should use the `-cxxlib-gcc` option if your application includes source files generated by `g++` and source files generated by the Intel C++ Compiler. This option directs the Intel compiler to use the `g++` header and library files to build one set of run-time libraries. As a result, your program should run correctly.

-cxxlib-icc option

The `-cxxlib-icc` option directs the Intel compiler to use the C++ run-time libraries and C++ header files included with the Intel compiler. They include:

- `libcprts` standard C++ headers
- `libcprts` standard C++ library
- `libcxa` and `libunwind` C++ language support



Note

The `-cxxlib-icc` option is ON by default if are using a `gcc` version less than 3.2.

-fabi-version

The `-fabi-version=n` option directs the compiler to select a specific ABI implementation. By default, the Intel compiler uses the ABI implementation that corresponds to the installed version of `gcc`. Both `gcc` 3.2 and 3.3 are not fully ABI-compliant.

Value of n	Description
<code>n=0</code>	Select most recent ABI implementation
<code>n=1</code>	Select <code>g++</code> 3.2 compatible ABI implementation
<code>n=2</code>	Select most conformant ABI implementation

See <http://www.codesourcery.com> for more information on ABI conformance.

See [Specifying Alternate Tools and Paths](#) for information on using `-Qlocation` to specify the location of the GNU assembler and linker.

Predefined Macros for Interoperability

The Intel® C++ Compiler and `gcc*` support the following predefined macros:

- `__GNUC__`
- `__GNUC_MINOR__`
- `__GNUC_PATCHLEVEL__`

You can specify the `-no-gcc` option to undefine these macros. If you need `gcc` interoperability (`-cxxlib-gcc`), do not use the `-no-gcc` compiler option.



Warning

Not defining these macros results in different paths through system header files. These alternate paths may be poorly tested or otherwise incompatible.

See also [Predefined Macros and GNU Environment Variables](#).

`gcc*` Built-in Functions

This version of the Intel® C++ compiler supports the following `gcc*` built-in functions:

```

__builtin_abs
__builtin_labs
__builtin_cos
__builtin_cosh
__builtin_fabs
__builtin_fabsf
__builtin_memcmp
__builtin_memcpy
__builtin_sin
__builtin_sinh
__builtin_sqrt
__builtin_sqrtf
__builtin_strcmp
__builtin_strlen
__builtin_strncmp
__builtin_abort
__builtin_prefetch
__builtin_constant_p
__builtin_printf
__builtin_fprintf
__builtin_fscanf
__builtin_scanf
__builtin_fputs
__builtin_memset
__builtin_strcat
__builtin_strcpy
__builtin_strncpy
__builtin_exit
__builtin_strchr
__builtin_strspn
__builtin_strcspn

```

```

__builtin_strstr
__builtin_strpbrk
__builtin_strrchr
__builtin_strncat
__builtin_alloca
__builtin_ffs
__builtin_index
__builtin_rindex
__builtin_bcmp
__builtin_bzero
__builtin_sinl
__builtin_cosl
__builtin_sqrtl
__builtin_fabsl
__builtin_frame_address (IA-32 only)
__builtin_return_address (IA-32 only)

```

For more information on gcc built-in functions, see <http://gcc.gnu.org/onlinedocs/gcc-3.4.1/gcc/Other-Builtins.html#Other%20Builtins>

gcc* Function Attributes

This version of the Intel® C++ Compiler supports the following gcc* function attributes:

- `noinline` - prevents a function from being inlined
- `always_inline` - inlines the function even if no optimization is specified
- `used - code` - must be emitted for the function even if the function is not referenced

Example

```
int round_sqrt(int) __attribute__ ((always_inline));
```

In this example, the function `round_sqrt()` is inlined even if no optimization is specified.

Thread-local Storage

The Intel® C++ Compiler supports the storage class keyword `__thread`, which can be used in variable definitions and declarations. Variables defined and declared this way are automatically allocated locally to each thread:

```

__thread int i;

__thread struct state s;

extern __thread char *p;

```



Note

The `__thread` keyword is only recognized when the GNU compatibility version is 3.3 or higher. You may need to specify the `-gcc-version=330` compiler option to enable thread-local storage.

See also <http://gcc.gnu.org/onlinedocs/gcc/Thread-Local.html>.

Language Conformance

Conformance Options

Option	Description
<code>-ansi</code>	Equivalent to GNU* ANSI
<code>-strict_ansi</code>	Strict ANSI conformance dialect

Conformance to the C Standard

You can set the Intel® C++ Compiler to accept either

- ANSI conformance equivalent to GNU* ANSI with the `-ansi` option, or
- Strict ANSI conformance dialect with the `-strict_ansi` option

The compiler is set by default to accept extensions and not be limited to the ANSI/ISO standard.

Understanding the ANSI/ISO Standard C Dialect

The Intel C++ Compiler provides conformance to the ANSI/ISO standard for C language compilation (ISO/IEC 9899:1990). This standard requires that conforming C compilers accept minimum translation limits. This compiler exceeds all of the ANSI/ISO requirements for minimum translation limits.

Macros Included with the Compiler

The ANSI/ISO standard for C language requires that certain predefined macros be supplied with conforming compilers. The following table lists the macros that the Intel C++ Compiler supplies in accordance with this standard:

The compiler includes predefined macros in addition to those required by the standard.

Macro	Value
<code>__DATE__</code>	The date of compilation as a string literal in the form <code>Mmm dd YYYY</code> .
<code>__FILE__</code>	A string literal representing the name of the file being compiled.
<code>__LINE__</code>	The current line number as a decimal constant.
<code>__STDC__</code>	The name <code>__STDC__</code> is defined when compiling a C translation unit.
<code>__STDC_HOSTED__</code>	The integer 1.
<code>__TIME__</code>	The time of compilation as a string literal in the form <code>hh:mm:ss</code> .

C99 Support

The following C99 features are supported in this version of the Intel C++ Compiler:

- restricted pointers (`restrict` keyword, available with `-restrict`). See Note.
- variable-length Arrays
- flexible array members
- complex number support (`_Complex` keyword)
- hexadecimal floating-point constants
- compound literals
- designated initializers
- mixed declarations and code
- macros with a variable number of arguments
- inline functions (`inline` keyword)
- boolean type (`_Bool` keyword)



Note

The `-restrict` option enables the recognition of the `restrict` keyword as defined by the ANSI standard. By qualifying a pointer with the `restrict` keyword, the user asserts that an object accessed via the pointer is only accessed via that pointer in the given scope. It is the user's responsibility to use the `restrict` keyword only when this assertion is true. In these cases, the use of `restrict` will have no effect on program correctness, but may allow better optimization.

These features are not supported:

- `#pragma STDC FP_CONTRACT`
- `#pragma STDC FENV_ACCESS`
- `#pragma STDC CX_LIMITED_RANGE`
- `long double` (128-bit representations)

Conformance to the C++ Standard

The Intel® C++ Compiler conforms to the ANSI/ISO standard (ISO/IEC 14882:1998) for the C++ language.

Exported Templates

The Intel® C++ Compiler supports exported templates using the following options:

Option	Description
<code>-export</code>	Enable recognition of exported templates. Supported in C++ mode only.
<code>-export_dir dir</code>	Specifies a directory name to be placed on the exported template search path. The directories are used to find the definitions of exported templates and are searched in the order in which they are specified on the command-line. The current directory is always the first entry on the search path.

Exported templates are templates declared with the `export` keyword. Exporting a class template is equivalent to exporting each of its static data members and each of its non-inline member functions. An exported template is unique because its definition does not need to be present in a translation unit that uses that template. For example, the following C++ program consists of two separate translation units:

```
// file1.cpp
#include <stdio.h>
static void trace() { printf("File 1\n"); }
export template<class T> T const& min(T const&, T const&);
int main() {
    trace();
    return min(2, 3);
}

// file2.cpp
#include <stdio.h>
static void trace() { printf("File 2\n"); }
export template<class T> T const& min(T const &a, T const &b) {
    trace();
    return a<b? a: b;
}
```

Note that these two files are separate translation units: one is not included in the other. That allows the two functions `trace()` to coexist (with internal linkage).

Usage

```
prompt>icpc -export -export_dir /usr2/export/ -c file1.cpp
```

```
prompt>icpc -export -export_dir /usr2/export/ -c file2.cpp
```

```
prompt>icpc -export -export_dir /usr2/export/ file1.o file2.o
```

Template Instantiation

The Intel® C++ Compiler supports `extern` template, which lets you specify that a template in a specific translation unit will not be instantiated because it will be instantiated in a different translation unit or different library. The compiler now includes additional support for:

- **inline template** -- instantiates the compiler support data for the class (i.e. the vtable) for a class without instantiating its members.
- **static template** -- instantiates the static data members of the template, but not the virtual tables or member functions.

You can now use the following options to gain more control over the point of template instantiation:

Option	Description
<code>-fno-implicit-templates</code>	Never emit code for non-inline templates which are instantiated implicitly (i.e. by use). only emit code for explicit instantiations.
<code>-fno-implicit-inline-templates</code>	Do not emit code for implicit instantiations of inline templates either. The default is to handle inlines differently so that compilations, with and without optimization, will need the same set of explicit instantiations.

Volume II: Optimizing Applications

Optimization Levels

This section discusses the command-line options `-O0`, `-O1`, `-O2`, 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`, and `-O3` optimize as follows:

- `-O1` -- code size and locality
- `-O2` -- code speed; this is the default option
- `-O3` -- enables `-O2` with more aggressive optimizations.

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

Setting Optimization Levels

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` options' behavior on each architecture.

Option	Effect
<code>-O0</code>	Disables optimizations.
<code>-O1</code>	Optimizes to favor code size and code locality. Disables loop unrolling. May improve performance for applications with very large code size, any branches, and execution time not dominated by code within loops. In most cases, <code>-O2</code> is recommended over <code>-O1</code> . IA-32 systems: Disables intrinsics inlining to reduce code size. Itanium-based systems: Disables software pipelining and global code scheduling.
<code>-O2</code> , <code>-O</code>	ON by default. Optimizes for code speed. This is the generally recommended optimization level. Itanium-based systems: Enables software pipelining.
<code>-O3</code>	Enables <code>-O2</code> optimizations and more aggressive optimizations such as loop and memory access transformations. 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. IA-32 systems: In conjunction with <code>-ax{K W N B P}</code> and <code>-x{K W N B P}</code> 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.

Option	Effect
<code>-fast</code>	<p>The <code>-fast</code> option enhances execution speed across the entire program by including the following options that can improve run-time performance:</p> <ul style="list-style-type: none"> • <code>-O3</code> (maximum speed and high-level optimizations.) • <code>-ipo</code> (enables interprocedural optimizations across files) • <code>-static</code> (prevents linking with shared libraries) • <code>-xP</code> (specific optimization for Intel Pentium 4 processor with Streaming SIMD Extensions 3). The <code>-fast</code> option does not include <code>-xP</code> when compiling on Itanium®-based systems. <p>To override one of the options set by <code>-fast</code>, specify that option after the <code>-fast</code> option on the command line. To target <code>-fast</code> optimizations for a specific processor, use one of the <code>-x</code> options. For example:</p> <pre>prompt>icpc -fast -xW source_file.cpp</pre> <p>The options set by <code>-fast</code> may change from release to release.</p>

Restricting Optimizations

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

Option	Description
<code>-O0</code>	Disables optimizations. Enables the <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.
<code>-g</code>	Specifying the <code>-g</code> option turns off the default <code>-O2</code> option and makes <code>-O0</code> the default unless <code>-O1</code> , <code>-O2</code> , or <code>-O3</code> is explicitly specified in the command line together with <code>-g</code> .
<code>-nolib_inline</code>	Disables inline expansion of intrinsic functions.

**Note**

You can turn off all optimizations for specific functions by using `#pragma optimize`. In the following example, all optimization is turned off for function `foo()`:

```
#pragma optimize("", off)
foo(){
  ...
}
```

Valid second arguments for `#pragma optimize` are "on" or "off." With the "on" argument, `foo()` is compiled with the same optimization as the rest of the program. The compiler ignores first argument values.

Floating-point Optimizations

Floating-point Arithmetic Precision

There are several compiler options that affect floating-point computations. In general, the options discussed here let you decide between performance and accuracy. To achieve greater performance, it may be necessary to sacrifice some degree of floating-point accuracy.

See also Floating-point Arithmetic Options for Itanium®-based Systems.

Options for IA-32 and Itanium®-based Systems

-mp Option

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. Floating point intermediate results are kept in full 10-byte internal precision. All spills and reloads of the X87 floating-point registers utilize this internal format to prevent accidental loss of precision.

For most programs, specifying this 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. Alternatives to `-mp` include `-xN` (for the Intel® Pentium 4 processor or newer) and `-mp1`.

- user variables declared as floating-point types are not assigned to registers.
- whenever an expression is spilled (moved from a register to memory), it is spilled as 80 bits (extended precision), not 64 bits (double precision).
- floating-point arithmetic comparisons conform to the IEEE 754 specification except for NaN behavior.
- 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 re-association.
- the compiler does not perform the constant-folding optimization 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.

- floating-point operations conform to ANSI C. When assignments to type `float` and `double` are made, the precision is rounded from 80 bits (extended) down to 32 bits (float) or 64 bits (double). When you do not specify `-mp`, the extra bits of precision are not always rounded before the variable is reused.
- sets the `-nolib_inline` option, which disables inline functions expansion.

-mp1 Option

Use the `-mp1` option to improve floating-point precision. `-mp1` disables fewer optimizations and has less impact on performance than `-mp`. `-mp1` prevents the compiler from performing optimizations which change NAN comparison semantics. It also causes all values used in comparisons to be truncated to declared precision prior to use in the comparison. It also makes sure to use library routines which give better precision results compared to the X87 transcendental instructions.

-complex_limited_range

This option enables the use of the basic algebraic expansions of some complex arithmetic operations. At the loss of some exponent range, the `-complex_limited_range` option can allow for some performance improvement in programs which utilize complex arithmetic. By default, the compiler disables this option by using `-complex_limited_range-`.

Options for IA-32 Only



Caution

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

-prec_div Option

With some optimizations, the Intel® C++ Compiler changes floating-point division computations into multiplication by the reciprocal of the denominator. For example, A/B is computed as $A \times (1/B)$ to improve the speed of the computation. However, for values of B greater than 2^{126} , the value of $1/B$ is "flushed" (changed) to 0. When it is important to maintain the value of $1/B$, use `-prec_div` to disable the floating-point division-to-multiplication optimization. The result of `-prec_div` is greater accuracy with some loss of performance.

-pcn Option

Use the `-pcn` option to enable floating-point significand precision control. Some floating-point algorithms are sensitive to the accuracy of the significand or fractional part of the floating-point value. For example, iterative operations like division and finding the square root can run faster if you lower the precision with the `-pcn` option. Set n to one of the following values to round the significand to the indicated number of bits:

- `-pc32`: 24-bit significand (single precision)
- `-pc64`: 53-bit significand (double precision)
- `-pc80`: 64-bit significand (long double precision)

The default value for n is 80, indicating long double precision. This option allows full optimization. Using this option does not have the negative performance impact of using the `-Op` option because only the fractional part of the floating-point value is affected. The range of the exponent is not affected. The `-pcn`

option causes the compiler to change the floating-point precision control when the `main()` function is compiled. The program that uses `-pcn` must use `main()` as its entry point, and the file containing `main()` must be compiled with `-pcn`.

-rcd Option

The Intel compiler uses the `-rcd` option to improve the performance of code that requires floating-point-to-integer conversions. The optimization is obtained by controlling the change of the rounding mode. The system default floating point rounding mode is round-to-nearest. This means that values are rounded during floating point calculations. However, the C 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-to-integer 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 C semantics.

-fp_port Option

The `-fp_port` option rounds floating-point results at assignments and casts. An impact on speed may result.

-fpstkchk Option

When a function call returns a floating-point value, the return value should be placed at the top of the FP stack. If the return value is unused, the compiler pops the value off the stack to keep the FP stack in the correct state. However, if the application leaves out the function's prototype or incorrectly prototypes the function, then the return value may remain on the stack. This may result in the FP stack filling up and eventually overflowing.

Generally, when the FP stack overflows, a NaN value is put into FP calculations, and the program's results differ. Unfortunately, the overflow point can be far away from the point of the actual bug. The `-fpchkstk` option places code that would access violate immediately after an incorrect call occurred, thus making it easier to locate these issues.

Floating-point Arithmetic Options for Itanium®-based Systems

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

- `-ftz[-]`
- `-IPF_fma[-]`
- `-IPF_fp_speculationmode`
- `-IPFflt_eval_method0`
- `-IPFfltacc[-]`(Default:`-IPFfltacc-`)
- `-IPFfp_relaxed[-]`

Flush Denormal Results to Zero

Use the `-ftz` option to flush denormal results to zero.

Contraction of FP Multiply and Add/Subtract Operations

`-IPF_fma[-]` enables [disables] the contraction of floating-point multiply and add/subtract operations into a single operation. Unless `-mp` is specified, the compiler contracts these operations whenever possible. The `-mp` option disables the contractions. Use `-IPF_fma` and `-IPF_fma-` to override the default compiler behavior. For example, a combination of `-mp` and `-IPF_fma` enables the compiler to contract operations (on Itanium®-based systems only):

```
prompt>icpc -mp -IPF_fma prog.cpp
```

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
- `safe`: enables the compiler to speculate on floating-point operations only when it is safe
- `strict`: disables the speculation of floating-point operations.
- `off`: disables the speculation on floating-point operations.



Note

`-IPF_fp_speculationsafe` is the default when `-O0` is specified.

FP Operations Evaluation

`-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.

Controlling Accuracy of the FP Results

`-IPFfltacc[-]` enables [disables] optimizations that affect floating-point accuracy. By default (`-IPFfltacc-`) the compiler may apply optimizations that reduce floating-point accuracy. You may use `-IPFfltacc` or `-mp` to improve floating-point accuracy, but at the cost of disabling some optimizations.

`-IPFfp_relaxed[-]` enables [disables] use of faster but slightly less accurate code sequences for math functions, such as divide and square root. As compared to strict IEEE* precision, using this option slightly reduces the accuracy of floating-point calculations performed by these functions, usually limited to the least significant digit.

Optimizing for Specific Processors

Processor Optimization for IA-32 only

The `-tpp{5|6|7}` options optimize your application's performance for a specific Intel processor. The resulting binary will also run on the other processors listed in the table. The Intel® C++ Compiler includes gcc*-compatible versions of the `-tpp` options. These options are listed in the **gcc* Version** column.

Option	gcc* Version	Optimizes for
<code>-tpp5</code>	<code>-mcpu=pentium</code>	Intel® Pentium® processors
<code>-tpp6</code>	<code>-mcpu=pentiumpro</code>	Intel Pentium Pro, Intel Pentium II, and Intel Pentium III processors
<code>-tpp7</code>	<code>-mcpu=pentium4</code>	Intel Pentium 4 processors, Intel Pentium M processors, and Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3)



Note

The `-tpp7` option is ON by default.

Example

The following invocations all result in a compiled binary optimized for Pentium 4. The same binary will also run on Pentium, Pentium Pro, Pentium II, and Pentium III processors.

```
prompt>icpc prog.cpp
prompt>icpc -tpp7 prog.cpp
prompt>icpc -mcpu=pentium4 prog.cpp
```

Processor Optimization (Itanium®-based Systems only)

The `-tpp{1|2}` options optimize your application's performance for a specific Intel® Itanium® processor. The resulting binary will also run on the processors listed in the table. The Intel® C++ Compiler includes gcc*-compatible versions of the `-tpp` options. These options are listed in the **gcc* Version** column.

Option	gcc* Version	Optimizes for
<code>-tpp1</code>	<code>-mcpu=itanium</code>	Itanium processors
<code>-tpp2</code>	<code>-mcpu=itanium2</code>	Itanium 2 processors



Note

The `-tpp2` option is ON by default.

Example

The following invocations all result in a compiled binary optimized for the Intel Itanium 2 processor. The same binary will also run on Intel Itanium processors.

```
prompt>icpc prog.cpp
```

```
prompt>icpc -tpp2 prog.cpp
```

```
prompt>icpc -mcpu=itanium2 prog.cpp
```

Processor-specific Optimization (IA-32 only)

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

Option	Specific Optimization 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. Programs compiled with this option 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. Programs compiled with this option 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 Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3). Programs compiled with this option 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 following invocation compiles `prog.cpp` 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.

```
prompt>icpc -xW prog.cpp
```

**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) 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 Optimizations (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 with this option will execute on any 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.
<code>-axP</code>	Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3). This option also enables new optimizations in addition to Intel processor specific-optimizations.

Example

The following compilation will generate 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 likely performance benefit
- a version optimized for Intel Pentium 4 processors, as long as there is a likely performance benefit

```
prompt>icpc -axKW prog.cpp
```

Manual CPU Dispatch (IA-32 only)

Use `__declspec(cpu_specific)` and `__declspec(cpu_dispatch)` in your code to generate instructions specific to the Intel processor on which the application is running, and also to execute correctly on other IA-32 processors.



Note

Manual CPU dispatch cannot be used to recognize Intel® Itanium® processors. The syntax of these extended attributes is as follows:

- `cpu_specific(cpu_id)`
- `cpu_dispatch(cpu_id-list)`

The values for `cpu_id` and `cpu_id-list` are shown in the following tables:

Processor	Values for <i>cpu_id</i>
x86 processors not provided by Intel Corporation	<code>generic</code>
Intel® Pentium® processors	<code>pentium</code>
Intel Pentium processors with MMX™ Technology	<code>pentium_mmx</code>
Intel Pentium Pro processors	<code>pentium_pro</code>
Intel Pentium II processors	<code>pentium_ii</code>
Intel Pentium III processors	<code>pentium_iii</code>
Intel Pentium III (exclude xmm registers)	<code>pentium_iii_no_xmm_regs</code>
Intel Pentium 4 processors	<code>pentium_4</code>
Intel Pentium M processors	<code>pentium_m</code>
Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3).	<code>pentium_4_sse3</code>
Values for <i>cpu_id-list</i>	
<code>cpu_id</code>	
<code>cpu_id-list, cpu_id</code>	

The attributes are not case sensitive. The body of a function declared with `__declspec(cpu_dispatch)` must be empty, and is referred to as a stub (an empty-bodied function).

Use the following guidelines to implement automatic processor dispatch support:

1. **Stub for `cpu_dispatch` must have a `cpuid` defined in `cpu_specific` elsewhere**
If the `cpu_dispatch` stub for a function `f` contains the `cpuid` `p`, then a `cpu_specific` definition of `f` with `cpuid` `p` must appear somewhere in the program; otherwise an unresolved external error is reported. A `cpu_specific` function definition need not appear in the same translation unit as the corresponding `cpu_dispatch` stub, unless the `cpu_specific` function is declared `static`. The `inline` attribute is disabled for all `cpu_specific` and `cpu_dispatch` functions.
2. **Must have a stub for `cpu_specific` function**
If a function `f` is defined as `__declspec(cpu_specific(p))`, then a `cpu_dispatch` stub must also appear for `f` within the program; and `p` must be in the `cpuid-list` of that stub; otherwise, that `cpu_specific` definition cannot be called nor generate an error condition.
3. **Overrides command line settings**
When a `cpu_dispatch` stub is compiled, its body is replaced with code that determines the processor on which the program is running, then dispatches the "best" `cpu_specific` implementation available as defined by the `cpuid-list`. The `cpu_specific` function optimizes to the specified Intel processor regardless of command-line option settings.

Processor Dispatch Example

Here is an example of how these features can be used:

```
#include <mmintrin.h>
/* Pentium processor function does not use intrinsics to add
two arrays. */

__declspec(cpu_specific(pentium))
void array_sum(int *r, int *a, int *b, size_t l)
{
    for (; length > 0; l--)
        *result++ = *a++ + *b++;
}

/* Implementation for a Pentium processor with MMX technology
uses
an MMX instruction intrinsic to add four elements
simultaneously. */

__declspec(cpu_specific(pentium_MMX))
void array_sum(int *r, int const *a, int *b, size_t l)
{
    __m64 *mmx_result = (__m64 *)result;
    __m64 const *mmx_a = (__m64 const *)a;
    __m64 const *mmx_b = (__m64 const *)b;

    for (; length > 3; length -= 4)
        *mmx_result++ = _mm_add_pi16(*mmx_a++, *mmx_b++);

    /* The following code, which takes care of excess elements,
is not
needed if the array sizes passed are known to be multiples
of four. */

    result = (unsigned short *)mmx_result;
    a = (unsigned short const *)mmx_a;
    b = (unsigned short const *)mmx_b;

    for (; length > 0; l--)
        *result++ = *a++ + *b++;
}

__declspec(cpu_dispatch(pentium, pentium_MMX))
void array_sum (int *r, int const *a, int *b, size_t l) )
{
    /* Empty function body informs the compiler to generate the
CPU-dispatch function listed in the cpu_dispatch clause. */
}
}
```

Processor-specific Runtime Checks, IA-32 Systems

The Intel® C++ Compiler optimizations take effect at run time. For IA-32 systems, the compiler enhances processor-specific optimizations by inserting a code segment in the program that performs the run-time checks described here.

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

To prevent execution errors, the compiler inserts code in the program to check for proper processor usage.

Programs compiled with options `-xN`, `-xB`, or `-xP` will check at run time whether they are being executed on the Intel® Pentium® 4 processor, Intel Pentium M processor, or the Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3), 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 the program `prog.cpp` for the Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3), issue the following command:

```
prompt>icpc -xP prog.cpp
```

The resulting executable aborts if it is executed on a processor that does not support the Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3), such as the Intel Pentium III or Intel Pentium 4 processor.

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 values of the flags flush-to-zero (FTZ) and denormals-as-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 Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3), 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.

Examples

- Executing a program on a Pentium III processor enables FTZ, but not DAZ.
- Executing a program on an Intel Pentium M processor or Intel Pentium 4 processor with Streaming SIMD Extensions 3 (SSE3) enables both FTZ and DAZ.

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

For non-Intel processors, you can set the flags manually with the following macros:

```
Enable FTZ: _MM_SET_FLUSH_ZERO_MODE(_MM_FLUSH_ZERO_ON)
```

```
Enable DAZ: _MM_SET_DENORMALS_ZERO_MODE(_MM_DENORMALS_ZERO_ON)
```

The prototypes for these macros are in `xmmintrin.h` (FTZ) and `pmmintrin.h` (DAZ).

Interprocedural Optimizations

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 and loop invariant code
Dead code elimination	Code size
Propagation of function characteristics	Call deletion and call movement
Multifile optimization	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 and register usage
Loop-invariant code motion	Further optimizations and 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.

`-auto_ilp32` for Itanium-based Systems

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.

IPO Compilation Model

For the topics in this section, the term IPO generally refers to multi-file IPO.

When you use the `-ipo` option, the compiler collects information from individual program modules of a program. Using this information, the compiler performs optimizations across modules. In order to do this, the `-ipo` option is applied to both the compilation phase and the link phase.

One of the main benefits of IPO is that it enables more inlining. For information on inlining and the minimum inlining criteria, see [Criteria for Inline Function Expansion and Controlling Inline Expansion of User Functions](#). Inlining and other optimizations are improved by profile information. For a description of how to use IPO with profile information for further optimization, see [Example of Profile-Guided Optimization](#).

Compilation Phase

When using IPO, as each source file is compiled, the compiler 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 IPO. Generating mock files instead of real object files reduces the time spent in the 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 or the `xild` tool. (See [Creating a Multifile IPO Executable with xild](#).)



Note

Failure to link "mock" objects with the `-ipo` option 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 invoke the linker, adding `-ipo` to the command line causes the compiler to be invoked a final time before the linker. The compiler performs IPO across all object files that have an IR. The compiler first analyzes all of the summary information, and then finishes compiling the pieces of the application for which it has IR. Having global information about the application while it is compiling individual pieces can improve the quality of optimization.



Note

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

When you use the `-ipo` option, the compiler attempts to detect 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.

Command Line for Creating an IPO Executable

The command line options to enable IPO for compilations targeted for both the IA-32 and Itanium® architectures are identical. To produce mock object files containing intermediate representation (IR), compile your source files with `-ipo` as follows:

```
prompt>icpc -ipo -c a.cpp b.cpp c.cpp
```

This produces `a.o`, `b.o`, and `c.o` object files. These files contain Intel compiler IR corresponding to the compiled source files `a.cpp`, `b.cpp`, and `c.cpp`. Using `-c` to stop compilation after generating `.o` files is required.

You can now optimize interprocedurally by adding `-ipo` to your link command line. The following example produces an executable named `app`:

```
prompt>icpc -oapp -ipo a.o b.o c.o
```

This command invokes the compiler on the objects containing IR and creates a new list of object(s) to be linked. The command then calls GCC `ld` to link the specified object files and produce `app`, as specified by the `-o` option. IPO is applied only to the object files that contain IR; otherwise the object file passes to link stage.



Note

For the above step, you can use the `xild` tool instead of `icpc`.

The two steps described above can be combined, as shown in the following:

```
prompt>icpc -ipo -oapp a.f b.f c.f
```

Generating Multiple IPO Object Files

For the most part, IPO generates a single object file for the link-time compilation. This can be clumsy for very large applications, perhaps even making it impossible to use `-ipo` on the application. The compiler provides two ways to avoid this problem. The first way is a size-based heuristic, which automatically causes the compiler to generate multiple object files for large link-time compilations. The second way is using one of two explicit command line controls that tell the compiler to do multi-object IPO:

- `-ipoN`, where *N* indicates the number of object files to generate.
- `-ipo_separate`, which tells the compiler to generate a separate IPO object file for each source file.

These options are alternatives to the `-ipo` option, that is, they indicate an IPO compilation. Explicitly requesting a multi-object IPO compilation turns the size-based heuristic off.

The number of files generated by the link-time compilation is invisible unless either the `-ipo_c` or `-ipo_s` option is used. In this case the compiler appends a number to the file name. For example, consider this example:

```
prompt>icpc -ipo_separate -ipo_c a.o b.o c.o
```

Here, `a.o`, `b.o`, and `c.o` all contain IR, so the compiler will generate `ipo_out.o`, `ipo_out1.o`, `ipo_out2.o`, and `ipo_out3.o`.

The first object file contains global symbols. The other object files correspond to the source files.

This naming convention is also applied to user-specified names. For example:

```
prompt>icpc -ipo_separate -ipo_c -o appl.o a.o b.o c.o
```

This will generate `appl.o`, `appl1.o`, `appl2.o`, and `appl3.o`.

Capturing Intermediate Outputs of IPO

The `-ipo_c` and `-ipo_s` options are useful either for analyzing the effects of IPO, or when using IPO on 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:

```
prompt>icpc -tpp6 -ipo_c -ofilename a.cpp b.cpp c.cpp
```

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:

```
prompt>icpc -tpp6 -ipo_s -ofilename a.cpp b.cpp c.cpp
```

The `-ipo_c` and `-ipo_s` options generate multiple outputs if multi-object IPO is being used. The name of the first file is taken from the value of the `-o` option. The name of subsequent files is derived from this file by appending a numeric value to the file name. For example, if the first object file is named `f00.o`, the second object file will be named `f001.o`.

The compiler generates a message indicating the name of each object or assembly file it is generating. These files can be added to the real link step to build the final application.

Creating a Multifile IPO Executable Using `xild`

Use the Intel linker, `xild`, instead of step 2 in Command Line for Creating an IPO Executable. The `xild` linker performs the following steps:

1. Invokes the compiler to perform IPO if objects containing IR are found.
2. Invokes GCC linker, `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 create `app` using IPO, use the option `-ofilename` as shown in the following example:

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

`xild` calls the 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 `app`.



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.

The `xild` command supports the `-ipo`, `-ipoN`, and `-ipo_separate` options.

Usage Rules

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

- Your source files were compiled with the `-ipo` 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 IPO. These options are described in the following table.

<code>-qipo_fa[file.s]</code>	Produces an assembly listing for the IPO compilation. You can 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 an object file for the IPO compilation. You can 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>	Adds code bytes to the assembly listing.
<code>-ipo_fsource-asm</code>	Adds high-level source code to the assembly listing.

<code>-ipo_fverbose-asm,</code> <code>-ipo_fnoverbose-asm</code>	Enables and disables, respectively, inserting comments containing version and options used in the assembly listing for <code>xild</code> .
---	--

If the `xild` invocation leads to an IPO multi-object compilation (either because the application is big, or because the user explicitly asked for multiple objects), the first `.s` file takes its name from the `-qipo_fa` option. The compiler derives the names of subsequent `.s` files by appending a number to the name, for example, `foo.s` and `foo1.s` for `-qipo_fafoo.s`. The same is true for the `-qipo_fo` option.

Code Layout and Multi-Object IPO

One of the optimizations performed during an IPO compilation is code layout. IPO analysis determines a layout order for all of the routines for which it has IR. If a single object is being generated, the compiler generates the layout simply by compiling the routines in the desired order.

For a multi-object IPO compilation, the compiler must tell the linker about the desired order. The compiler first puts each routine in a named text section (the first routine in `.text00001`, the second in `.text00002`, and so forth). It then generates a linker script that tells the linker to first link contributions from `.text00001`, then `.text00002`. This happens transparently when the same invocation is used for both the link-time compilation and the final link.

However, the linker script must be taken into account by the user if `-ipo_c` or `-ipo_S` is used. With these switches, the IPO compilation and actual link are done by different invocations. When this occurs, the compiler will issue an informational message indicating that it is generating an explicit linker script, `ipo_layout.script`.

When `ipo_layout.script` is generated, the typical response is to modify your link command to use this script:

```
--script=ipo_layout.script
```

If your application already requires a custom linker script, you can place the necessary contents of `ipo_layout.script` in your script. The layout-specific content of `ipo_layout.script` is at the beginning of the description of the `.text` section. For example, to describe the layout order for 12 routines:

```
.text:
{
*(.text00001) *(.text00002) *(.text00003) *(.text00004) *(.text00005)
*(.text00006) *(.text00007) *(.text00008) *(.text00009) *(.text00010)
*(.text00011) *(.text00012)
...
}
```

For applications that already require a linker script, you can add this section of the `.text` section description to the customized linker script. If you add these lines to your linker script, it is desirable to add additional entries to account for future development. This is harmless, since the `**(...)` syntax makes these contributions optional.

If you choose to not use the linker script your application will still build, but the layout order will be random. This may have an adverse affect on application performance, particularly for large applications.

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® C++ 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 a file with a `.il` suffix. 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 your 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.

```
prompt>xiar cru user.a a.o b.o
```

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:

```
prompt>icpc -ipo -c a.cpp b.cpp
```

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 tool `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.

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

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

Using -ip with -Qoption Specifiers

You can adjust the Intel® C++ 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 an `-ip` or `-ipo` specification, as follows:

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

where `tool` is C++ (`c`) 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.

-Qoption Specifiers

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

- 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 following table:

-Qoption Specifiers	Description
<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 <code>n</code> is 230.
<code>-ip_ninl_min_stats=n</code>	Sets the valid <code>min</code> number of intermediate language statements for a function that is expanded in line. The number <code>n</code> is a positive integer. The default value for <code>ip_ninl_min_stats</code> is: IA-32 compiler: <code>ip_ninl_min_stats = 7</code> Itanium® compiler: <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 <code>n</code> is a positive integer. The default value for <code>n</code> is 2000.

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

```
prompt>icpc -ip -Qoption,c,-ip_ninl_max_stats=5 source.cpp
```

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	Description
<code>-ip_no_inlining</code>	This option is only useful if <code>-ip</code> is also specified. In this 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.
<code>-ip_no_pinlining</code>	Disables partial inlining; can be used if <code>-ip</code> or <code>-ipo[<i>value</i>]</code> is also specified.

Criteria for Inline Function Expansion

Once the criteria are met, the compiler picks the routines whose inline expansion will provide the greatest benefit to program performance. The inlining heuristics used by the compiler differ, based on whether or not you use profile-guided optimizations (`-prof_use`). When you use profile-guided optimizations with `-ip` or `-ipo[value]`, 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 will not inline functions with more than 230 intermediate statements. You can change this value by specifying the option `-Qoption,c,-ip_ninl_max_stats=new_value`. Note: there is a higher limit for functions declared by the user as `inline` or `__inline`.
- The default inline heuristic will stop inlining when direct recursion is detected.
- The default heuristic will always inline 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`. This limit can be modified with the option `-Qoption,c,-ip_ninl_min_stats=new_value`.

If you do not use profile-guided optimizations with `-ip` or `-ipo[value]`, the compiler uses less aggressive inlining heuristics:

- Inline a function if the inline expansion will not increase the size of the final program.
- Inline a function if it is declared with the `inline` or `__inline` keywords.

Profile-guided Optimizations

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 use feedback from a previous compilation to be more selective 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 used 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 here help you determine if you can benefit by using IPO and PGO.

Profile-guided Optimizations Methodology

PGO works best for code with many frequently executed branches that are difficult to predict at compile time. An example is code that is heavy with 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 rarely mispredicted. Eliminating the interleaving of "hot" and "cold" code improves instruction cache behavior. For example, the use of PGO often enables the compiler to make better decisions about function inlining, thereby increasing the effectiveness of interprocedural optimizations.

PGO Phases

The PGO methodology requires three phases:

- Phase 1: Instrumentation compilation and linking with `-prof_gen[x]`
- Phase 2: Instrumented execution by running the executable
- Phase 3: Feedback compilation with `-prof_use`

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.

In cases where your code behavior differs greatly between executions, PGO may not provide noticeable benefits. You have to ensure that the benefit of the profile information is worth the effort required to maintain up-to-date profiles.

When using `-prof_gen[x]` with the `x` qualifier, extra source position is collected which enables code coverage tools, such as the Intel® C++ Compiler Code-coverage Tool. Without such tools, `-prof_genx` does not provide better optimization and may slow parallel compile times.

Basic PGO Options

Option	Description
<code>-prof_gen[x]</code>	Instructs the compiler to produce instrumented code in your object files in preparation for instrumented execution.
<code>-prof_use</code>	Instructs the compiler to produce a profile-optimized executable and merges available dynamic information (<code>.dyn</code>) files into a <code>pgopti.dpi</code> file.

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

The `-prof_gen[x]` 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_genx` compilations.

Generating a Profile-optimized Executable

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.

Disabling Function Splitting (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.

Example of Profile-guided Optimization

The three basic phases of PGO are:

- Instrumentation Compilation and Linking
- Instrumented Execution
- Feedback Compilation

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:

```
prompt>icpc -prof_gen -prof_dir /profdata -c a1.cpp a2.cpp a3.cpp
prompt>icpc a1.o a2.o a3.o
```

In place of the second command, you could use the linker directly to produce the instrumented program.

Instrumented Execution

Run your instrumented program with a representative set of data to create a dynamic information file.

```
prompt>./a.out
```

The resulting dynamic information file has a unique name and `.dyn` suffix every time you run `a.o`. 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.

Feedback Compilation

Compile and link the source files with `-prof_use` to use the dynamic information to optimize your program according to its profile:

```
prompt>icpc -prof_use -ipo a1.cpp a2.cpp a3.cpp
```

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 with `-ipo` for phase 3. This example used `-O2` in phase 1 and `-O2 -ipo` in phase 3.



Note

The compiler ignores the `-ipo` options with `-prof_gen[x]`. With the `x` qualifier, extra information is gathered.

PGO Environment Variables

The following table describes environment values to determine the directory to store dynamic information files or whether to overwrite `pgopti.dpi`. Refer to your operating system documentation for instructions on how to specify environment values.

Profile-guided Optimization Environment Variables

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_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 if <code>.dyn</code> files are newer than an existing <code>pgopti.dpi</code> file. 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.

Using profmerge to Relocate the Source Files

The compiler uses the full path to the source file to look up profile summary information. 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 `profmerge` with the `-src_old` and `-src_new` options. For example:

```
prompt>profmerge -prof_dir <p1> -src_old <p2> -src_new <p3>
```

where:

- `<p1>` is the full path to dynamic information file (`.dpi`).
- `<p2>` is the old full path to source files.
- `<p3>` is the new full path to source files.

This command will read the `pgopti.dpi` file. For each function represented in the `pgopti.dpi` file, whose source path begins with the `<p2>` prefix, `profmerge` replaces that prefix with `<p3>`. The `pgopti.dpi` file is updated with the new source path information.

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:

```
prompt>profmerge -prof_dir -src_old /src/prog_1 -src_new /src/prog_2
```

```
prompt>profmerge -prof_dir -src_old /proj_1 -src_new /proj_2
```

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.

PGO API Support Overview

Profile Information Generation Support lets you control of 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. The functions described in this section may be necessary in assuring that profile information is generated in the following situations:

- when the instrumented application exits using a non-standard exit routine
- when instrumented application is a non-terminating application where `exit()` is never called
- when you want control of when the profile information is generated

This section includes descriptions of the functions and environment variable that comprise Profile Information Generation Support. The functions are available by inserting `#include <pgouser.h>` at the top of any source file where the functions may be used.

The compiler sets a `define` for `_PGO_INSTRUMENT` when you compile with either `-prof_gen` or `-prof_genx`.

Dumping Profile Information

```
void _PGOPTI_Prof_Dump(void);
```

Description

This function dumps the profile information collected by the instrumented application. The profile information is recorded in a `.dyn` file.

Recommended Usage

Insert a single call to this function in the body of the function which terminates your application.

Normally, `_PGOPTI_Prof_Dump` should be called just once. It is also possible to use this function in conjunction with `_PGOPTI_Prof_Reset()` 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();

while(input_data)
{
    _PGOPTI_Prof_Reset();
    process_data(input_data);
    _PGOPTI_Prof_Dump();
    input_data = get_input_data();
}

```

Resetting the Dynamic Profile Counters

```
void _PGOPTI_Prof_Reset(void);
```

Description

This function resets the dynamic profile counters.

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

```
void _PGOPTI_Prof_Dump_And_Reset(void);
```

Description

This function may be called more than once. Each call will dump the profile information to a new `.dyn` file. The dynamic profile counters are then reset, and execution of the instrumented application continues.

Recommended Usage

Periodic calls to this function allow a non-terminating application to generate one or more profile information files. These files are merged during the feedback phase of profile-guided optimization. The direct use of this function allows your application to control precisely when the profile information is generated.

Interval Profile Dumping

```
void _PGOPTI_Set_Interval_Prof_Dump(int interval);
```

Description

This function activates Interval Profile Dumping and sets the approximate frequency at which dumps will occur. The `interval` parameter is measured in milliseconds and specifies the time interval at which profile dumping will occur. 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.



Note

- Setting `interval` to zero or a negative number will disable interval profile dumping.
- Setting `interval` to a very small value 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 collect substantial profile information.

Recommended Usage

Call this function at the start of a non-terminating 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.

Environment Variable

`PROF_DUMP_INTERVAL`

This environment variable may be used to initiate Interval Profile Dumping in an instrumented application. See the Recommended Usage of `_PGOPTI_Set_Interval_Prof_Dump` for more information.

Code-coverage Tool

The Intel® C++ Compiler 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 compiler Code-coverage Tool are:

- Visual presentation of the application's code coverage information with a 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. 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.	
<code>-maddr</code>	Sets the email address of the web-page owner.	
<code>-bcolor</code>	Sets the html color name or code of the uncovered blocks.	<code>#ffff99</code>
<code>-fcolor</code>	Sets the html color name or code of the uncovered functions.	<code>#ffc000</code>
<code>-pcolor</code>	Sets the html color name or code of the partially covered code.	<code>#f0f0f0</code>
<code>-ccolor</code>	Sets the html color name or code of the covered code.	<code>#ffffff</code>
<code>-ucolor</code>	Sets the html color name or code of the unknown code.	<code>#ffffff</code>

Visual Presentation of the Application's Code Coverage

Based on the profile information collected from running the instrumented binaries when testing an application, the 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 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 example that follows shows a 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.

generated by [Intel® Compilers code-coverage tool](#)

Coverage Summary of Sample_Project

Files				Functions				Blocks			
total	cvrd	uncvrd	cvrg%	total	cvrd	uncvrd	cvrg%	total	cvrd	uncvrd	cvrg%
3	2	1	66.67	19	5	14	26.32	134	25	109	18.66

Covered Files in Sample_Project						
Name	Functions			Blocks		
	total	cvrd	cvrg%	total	cvrd	cvrg%
SAMPLE2.C	7	1	14.29	50	2	4.00
SAMPLE.C	5	4	80.00	34	23	67.65

Uncovered Files in Sample_Project			
Name	Functions		Blocks
	total	cvrd	total
SAMPLE3.C	7	0	50

generated by [Intel® Compilers code-coverage tool](#) Web-Page Owner: [Intel](#)

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.

This example shows the coverage source view of `SAMPLE.C`.

The screenshot shows a web browser displaying the code coverage information for `SAMPLE.C`. The browser window title is "Intel® Compilers code-coverage information for D:\COVERAGE\IA32\COMPILER\SAMPLE\SAMPLES\SAMPLE - Microsoft Internet Explorer". The address bar shows the file path: `D:\Coverage\ia32\compiler\sample\sample5\CodeCoverage\0_COVERAGE_IA32_COMPILER_SAMPLE_SAMPLES_SAMPLE_C.HTML`.

The left sidebar contains two panels:

- uncovered functions**: Lists `g2` with 8 blocks.
- covered functions**: Lists `f1` (66.67 (4/6)), `f2` (83.33 (5/6)), `g1` (100.00 (8/8)), and `main` (100.00 (15/15)).

The main content area shows the source code for `SAMPLE.C` with line numbers 9 to 40. The code is color-coded to indicate coverage:

- Lines 11, 12, 18, and 19 are highlighted in yellow, indicating they are uncovered.
- Lines 34, 35, 36, 37, and 38 are highlighted in pink, indicating they are covered.

```

9) void f1 (int n)
10) {
11)     if |(n == 1) || (n == 0)| (
12)         printf ("1 or 0\n");
13)     )
14) }
15)
16) void f2 (int n)
17) {
18)     if |(n == 1) || (n == 0)| (
19)         printf ("1 or 0\n");
20)     )
21) }
22)
23) void g1 (int n)
24) {
25)     int j, k;
26)
27)     for (j = 0; j < n; j ++| (
28)         a ++;
29)     )
30) }
31)
32) void g2 (int n)
33) {
34)     int j, k;
35)
36)     for (j = 0; j < n; j ++| (
37)         a ++;
38)     )
39) }
40)

```

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 table 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 color by using the options mentioned for each coverage category in the preceding table.

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 compiler 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 divide 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:

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



Note

Each line of the 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` contains `mod1.cpp`, 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.cpp`, then only those modules with the name `mod1.cpp` 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, providing useful information for both coverage and performance tuning.

The coverage tool can be configured to generate information about dynamic execution counts. This configuration requires 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 (macros, for example), 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.

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 `printf` 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

To run the Code-Coverage Tool on an application, developers must provide the following three items:

- The application sources
- The `.SPI` file generated by Intel Compilers when compiling the application for the instrumented binaries through the `-prof_genx` option
- The `.DPI` file generated by the Intel Compiler's `profmerge` tool that result from merging the dynamic profile information files `*.DYN` or the `.DPI` file generated implicitly by Intel Compilers, when compiling the application with the `-prof_use` option

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 Code-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® compiler Test-prioritization Tool enables profile-guided optimizations to select and prioritize application 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 lets you 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 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
<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> . Default is <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.

Option	Description
<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 the Intel compilers when compiling the application for the instrumented binaries with the `-prof_genx` option.
- The `.dpi` files generated by the Intel compiler `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 you make 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.



Note

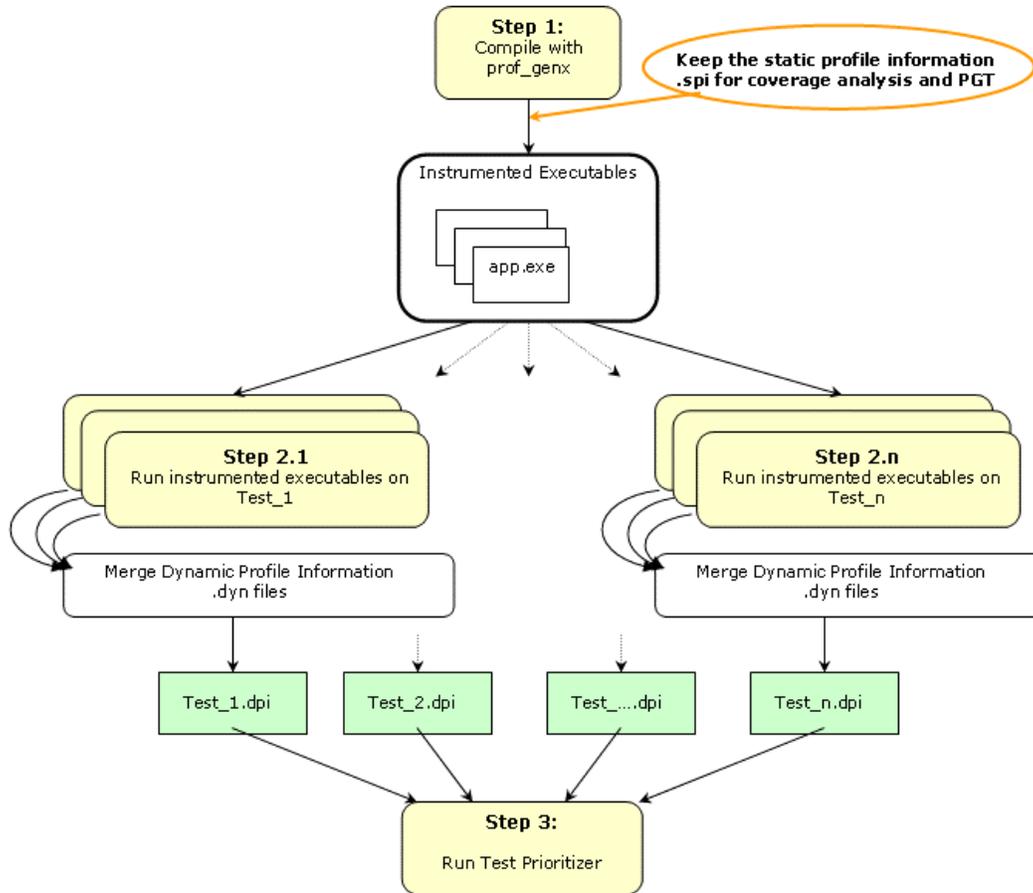
For successful tool execution, you should:

- 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.c`) for IA-32 systems.

1. Set

```
PROF_DIR=/myApp/prof_dir
```

2. Issue command

```
prompt>icpc -prof_genx myApp.c
```

This command compiles the program and generates an instrumented binary 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` 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 the preceding example: 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 previous 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 determine 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.

High-level Language Optimizations (HLO)

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 required to turn 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 the `-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 the 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 an `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 distribution
- Loop fusion
- Scalar replacement
- Absence of loop-carried memory dependency with `IVDEP` directive
- Runtime Data Dependencies checking (Itanium®-based systems only)

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
- Dead code elimination

In addition to the loop transformations listed for both IA-32 and Itanium® architectures above, the Itanium architecture enables implementation of collapsing techniques.

Scalar Replacement

The goal of scalar replacement, which is enabled by `-scalar_rep`, 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` and `-scalar_rep` are 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.

The `-scalar_rep` compiler option enables (default) scalar replacement performed during loop transformations. The `-scalar_rep-` option disables this scalar replacement.

Loop Unrolling with `-unroll`

The `-unroll[n]` option is used in the following way:

- `-unroll n` specifies the maximum number of times you want to unroll a loop. The following example unrolls a loop at most four times:

```
prompt>icpc -unroll4 a.cpp
```

To disable loop unrolling, specify `n` as 0. The following example disables loop unrolling:

```
prompt>icpc -unroll0 a.cpp
```

- `-unroll` (`n` omitted) lets the compiler decide whether to perform unrolling or not. This is the default; the compiler uses default heuristics or defines `n`.
- `-unroll0` (`n = 0`) disables the loop unroller.

The Itanium® compiler currently recognizes only `n = 0`; any other value is ignored.

Benefits and Limitations of Loop Unrolling

The benefits of loop unrolling are as follows:

- 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 and 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 processors, until they have a maximum of 16 iterations
 - Pentium III or Pentium II processors, until they have a maximum of 4 iterations

A potential limitation is that excessive unrolling, or unrolling of very large loops, can lead to increased code size.

For more information on how to optimize with `-unroll[n]`, refer to the *Intel® Pentium® 4 and Intel® Xeon™ Processor Optimization Reference Manual*.

Absence of Loop-carried Memory Dependency

For Itanium®-based applications, the `-ivdep_parallel` option indicates there is absolutely no loop-carried memory dependency in the loop where the `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.

```
#pragma ivdep
for (i=1; i<n; i++)
{
    e[ix[2][i]] = e[ix[2][i]]+1.0;
    e[ix[3][i]] = e[ix[3][i]]+2.0;
}
```

The following example shows that using this option and the `IVDEP` directive ensures there is no loop-carried dependency for the store into `a()`.

```
#pragma ivdep
for (j=0; j<n; j++)
{
    a[b[j]] = a[b[j]] + 1;
}
```

PREFETCH Directive

The `PREFETCH` directive is supported on Itanium®-based systems only.

Syntax:

```
#pragma prefetch var:hint:distance
```

where `hint` value can be 0 (T0), 1 (NT1), 2 (NT2), or 3 (NTA)

Example:

```

for (i=i0; i!=i1; i+=is) {
    float sum = b[i];
    int ip = srow[i];
    int c = col[ip];

    #pragma NOPREFETCH col
    #pragma PREFETCH value:1:80
    #pragma PREFETCH x:1:40

    for(; ip<srow[i+1]; c=col[++ip])
        sum -= value[ip] * x[c];
    y[i] = sum;
}

```

Prefetching

The goal of prefetch insertion optimization is to reduce cache misses by providing hints to the processor about when data should be loaded into the cache. The prefetch optimization is enabled or disabled by the `-prefetch[-]` compiler option.

`-prefetch` enables (default) prefetch insertion optimization. Note that `-O3` must be specified for this option to work.

To disable prefetch insertion optimization, use `-prefetch-`.

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 the Intel® Pentium® 4 and Intel Xeon Processor Optimization Reference Manual.

In addition to the `-prefetch` option, the `_mm_prefetch` intrinsic and `PREFETCH` compiler directive are also available. The intrinsic prefetches data from the specified address on one memory cache line. The compiler directive enables a data prefetch from memory.

Key Tuning Techniques

Use the following techniques to tune your applications for Itanium®-based systems:

- Compile your program with the `-O3` and `-Qipo` options. Use profile guided optimization (PGO) whenever possible.
- Identify hot spots in your code.
- Turn on Optimization Report.
- Check why loops are not software pipelined.
 - Use `#pragma ivdep` to indicate there is no dependence. You might need to compile with the `-ivdep_parallel` option to absolutely specify no loop carried dependence.
 - Use `#pragma swp` to enable software pipelining (useful for lop-sided controls and unknown loop count).
 - Use `#pragma loop count(n)` when needed.
 - Use of `-ansi-alias` is helpful. For example, for `**p = *q`, the ANSI rule indicates the pointer and float data do not overlap.
 - Add the `restrict` keyword to insure there is no aliasing.
 - Use `-alias_args-` to indicate arguments are not aliased.
 - Use `-fno_alias` only if pointers get traced back to the same base pointer.
 - Use `#pragma distribute point` to split large loops (normally this is done automatically).
- For C code, do not use `unsigned int` for loop indexes. HLO may skip optimization due to possible subscripts overflow. If upper bounds are pointer references, assign it to a local variable whenever possible.
 - Is prefetch distance correct? Use `#pragma prefetch` to override the distance when it is needed.

Parallel Programming

For parallel programming, the Intel® C++ Compiler supports both the OpenMP* 2.0 API and an automatic parallelization capability. The following table lists the options that perform OpenMP and auto-parallelization support.

Option	Description
<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>-openmp_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.
<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=100.
<code>-par_report{0 1 2 3}</code>	Controls the auto-parallelizer's diagnostic levels. Default: <code>-par_report1</code> .



Note

When both `-openmp` and `-parallel` are specified on the command line, the `-parallel` option is honored only in routines that do not contain OpenMP directives. For routines that contain OpenMP directives, only the `-openmp` option is honored.

Vectorization

The vectorizer is a component of the Intel® C++ 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 program to process 2, 4, 8, or 16 elements in one operation, depending on the data type.

This section provides guidelines, option descriptions, and examples for the Intel C++ Compiler vectorization on IA-32 systems only. The following list summarizes this section's contents.

- a quick reference of vectorization functionality and features
- descriptions of compiler switches to control vectorization
- descriptions of the C++ language features to control vectorization
- discussion and general guidelines on vectorization levels:
 - automatic vectorization
 - vectorization with user intervention
- examples demonstrating typical vectorization issues and resolutions

Vectorizer Options

Option	Description
<code>-ax{K W N B P}</code>	Enables the vectorizer and generates specialized and generic IA-32 code. The generic code is usually slower than the specialized code.
<code>-x{K W N B P}</code>	Turns on the vectorizer and generates processor-specific specialized code.
<code>-vec_reportn</code>	Controls the vectorizer's level of diagnostic messages: <ul style="list-style-type: none"> <code>n = 0</code> no diagnostic information is displayed. <code>n = 1</code> display diagnostics indicating loops successfully vectorized (default). <code>n = 2</code> same as <code>n = 1</code>, plus diagnostics indicating loops not successfully vectorized. <code>n = 3</code> same as <code>n = 2</code>, plus additional information about any proven or assumed dependences.

Usage

If you use `-c`, `-ipo` with `-vec_report{n}` option or `-c`, `-x{K|W|N|B|P}` or `-ax{K|W|N|B|P}` with `-vec_report{n}`, the compiler issues a warning and no report is generated.

To produce a report when using the aforementioned 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 a vectorization report:

- prompt>`icpc -x{K|W|N|B|P} -vec_report3 file.cpp`
- prompt>`icpc -x{K|W|N|B|P} -ipo -ipo_obj -vec_report3 file.cpp`
- prompt>`icpc -c -x{K|W|N|B|P} -ipo -ipo_obj -vec_report3 file.cpp`

The following commands do not generate a vectorization report:

- prompt>`icpc -c -x{K|W|M|B|P} -vec_report3 file.cpp`
- prompt>`icpc -x{K|W|N|B|P} -ipo -vec_report3 file.cpp`
- prompt>`icpc -c -x{K|W|N|B|P} -ipo -vec_report3 file.cpp`

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.

Note that in some cases successful loop parallelization (either automatically or by means of OpenMP* directives) may affect the messages reported by the compiler for loop vectorization; for example, under the `-vec_report2` option indicating loops not successfully vectorized.

Vectorization Key Programming Guidelines

The goal of vectorizing compilers is to exploit single-instruction multiple data (SIMD) processing automatically. 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 for loop bodies:

- use straight-line code (a single basic block)
- use 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
- use only assignment statements

Avoid the following in loop bodies:

- function calls
- unvectorizable operations
- mixing vectorizable types in the same loop
- data-dependent loop exit conditions

Preparing your code for vectorization

To make your code vectorizable, you will often need to make some changes to your loops. However, you should make only the changes needed to enable vectorization and no others. In particular, you should avoid these common changes:

- do not unroll your loops, the compiler does this automatically
- do not decompose one loop with several statements in the body into several single-statement loops

Restrictions

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 two memory references at 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 vectorizations. The following topics 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. The "Data-dependent Loop" example shows some code that exhibits data dependence. The value of each element of an array is dependent on itself and its two neighbors.

Data-dependent Loop

```
float data[N];
int i;

for (i=1; i<N-1; i++)
{
    data[i]=data[i-1]*0.25+data[i]*0.5+data[i+1]*0.25;
}
```

The loop in this 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 in the following example:

Data Dependence Vectorization Patterns

```
for(i=0; i<100; i++)
a[i]=b[i];
has access pattern
read b[0]
write a[0]
read b[1]
write a[1]
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 the loop shown, the value of `data[1]` read during the second iteration was written into the first iteration. For vectorization, the iterations must be done in parallel, without changing the semantics of the original loop.

Data Dependence Theory

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,
- for array references, the relationship between the subscripts.

For array references, the Intel® C++ Compiler's data dependence analyzer is organized as a series of tests that progressively increase in power as well as 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. Multi-dimensional 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 used are the fast GCD test, proving independence if the greatest common divisor of the coefficients of loop indices cannot evenly divide the constant term, and the extended bounds test, which tests potential overlap for the extreme values of subscript expressions.

If all simple tests fail to prove independence, the compiler will eventually resort to a powerful hierarchical dependence solver that uses Fourier-Motzkin elimination to solve the data dependence problem in all dimensions.

Loop Constructs

Loops can be formed with the usual `for` and `while` constructs. However, the loops must have a single entry and a single exit to be vectorized.

Correct Usage

```
while(i<n)
{
    // If branch is inside body of loop
    a[i]=b[i]*c[i];
    if(a[i]<0.0)
    {
        a[i]=0.0;
    }
    i++;
}
```

Incorrect Usage

```
while(i<n)
{
    if (condition) break;
    // 2nd exit.
    ++i;
}
```

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. The following examples illustrate countable and non-countable loop constructs.

Correct Usage for Countable Loop

```
// Exit condition specified by "N-1b+1"
count=N;

...

while(count!=1b)
{
    // 1b is not affected within loop
    a[i]=b[i]*x;
    b[i]=[i]+sqrt(d[i]);
    --count;
}
```

Correct Usage for Countable Loop

```
// Exit condition is "(n-m+2)/2"
i=0;
for(l=m; l<n; l+=2)
{
    a[i]=b[i]*x;
    b[i]=c[i]+sqrt(d[i]);
    ++i;
}
```

Incorrect Usage for Non-Countable Loop

```
i=0;

// Iterations dependent on a[i]
while(a[i]>0.0)
{
    a[i]=b[i]*c[i];
    ++i;
}
```

Types of Loops Vectorized

For integer loops, MMX(TM) technology and Streaming SIMD Extensions 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 if the final stored value is a 16-bit integer. Also, note that because the MMX(TM) instructions and Streaming SIMD Extensions instruction sets are not fully orthogonal (byte shifts, 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, the Streaming SIMD Extensions provide SIMD instructions for the arithmetic operators `+`, `-`, `*`, and `/`. Also, the Streaming SIMD Extensions provide SIMD instructions for the binary `MIN`, `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® C++ Compiler.

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.

Before Vectorization

```
i=0;
while(i<n)
{
    // Original loop code
    a[i]=b[i]+c[i];
    ++i;
}
```

After Vectorization

```

// The vectorizer generates the following two loops
i=0;

while(i<(n-n%4))
{
    // Vector strip-mined loop
    // Subscript [i:i+3] denotes SIMD execution
    a[i:i+3]=b[i:i+3]+c[i:i+3];
    i=i+4;
}

while(i<n)
{
    // Scalar clean-up loop
    a[i]=b[i]+c[i];
    ++i;
}

```

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 contain float operations (typically on arrays). Supported arithmetic operations include addition, subtraction, multiplication, division, negation, square root, max, and min. Operation on double precision types is not permitted unless optimizing for a Pentium® 4 processor system.

Integer Array Operations

The statements within the loop body may contain `char`, `unsigned char`, `short`, `unsigned short`, `int`, and `unsigned int`. Calls to functions such as `sqrt` and `fabs` are also supported. Arithmetic operations are limited to addition, subtraction, bitwise AND, OR, and XOR operators, division (16-bit only), multiplication (16-bit only), min, and max. 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 allowed. In particular, note that the special `__m64` and `__m128` datatypes are not vectorizable. The loop body cannot contain any function calls. Use of the Streaming SIMD Extensions intrinsics (`__mm_add_ps`) are not allowed.

Language Support and Directives

This topic addresses language features that better help to vectorize code. The `declspec(align(n))` declaration enables you to overcome hardware alignment constraints. The `restrict` qualifier and the pragmas address the stylistic issues due to lexical scope, data dependence, and ambiguity resolution.

Language Support

Feature	Description
<code>__declspec(align(n))</code>	Directs the compiler to align the variable to an n -byte boundary. Address of the variable is $address \bmod n=0$.
<code>__declspec(align(n,off))</code>	Directs the compiler to align the variable to an n -byte boundary with offset off within each n -byte boundary. Address of the variable is $address \bmod n=off$.
<code>restrict</code>	Permits the disambiguator flexibility in alias assumptions, which enables more vectorization.
<code>__assume_aligned(a,n)</code>	Instructs the compiler to assume that array a is aligned on an n -byte boundary; used in cases where the compiler has failed to obtain alignment information.
<code>#pragma ivdep</code>	Instructs the compiler to ignore assumed vector dependencies.
<code>#pragma vector{aligned unaligned always}</code>	Specifies how to vectorize the loop and indicates that efficiency heuristics should be ignored.
<code>#pragma novector</code>	Specifies that the loop should never be vectorized

Multi-version Code

Multi-version code is generated by the compiler in cases where data dependence analysis fails to prove independence for a loop due to the occurrence of pointers with unknown values. This functionality is referred to as dynamic dependence testing.

Pragma Scope

These pragmas control the vectorization of only the subsequent loop in the program, but the compiler does not apply them to any nested loops. Each nested loop needs its own `pragma` preceding it in order for the `pragma` to be applied. You must place a `pragma` only before the loop control statement.

#pragma vector always

Syntax: `#pragma vector always`

Definition: This `pragma` instructs the compiler to override any efficiency heuristic during the decision to vectorize or not. `#pragma vector always` will vectorize non-unit strides or very unaligned memory accesses.

Example:

```
for(i = 0; i <= N; i++)
{
    a[32*i] = b[99*i];
}
```

#pragma ivdep

Syntax: `#pragma ivdep`

Definition: This `pragma` 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 `pragma` overrides that decision. Only use this when you know that the assumed loop dependences are safe to ignore.

The loop in this example will not vectorize with the `ivdep` `pragma`, since the value of `k` is not known (vectorization would be illegal if `k < 0`).

Example:

```
#pragma ivdep
for (i = 0; i < m; i++)
{
    a[i] = a[i + k] * c;
}
```

#pragma vector

Syntax: `#pragma vector{aligned | unaligned}`

Definition: The vector loop `pragma` means the loop should be vectorized, if it is legal to do so, ignoring normal heuristic decisions about profitability. When the `aligned` (or `unaligned`) qualifier is used with this `pragma`, the loop should be vectorized using `aligned` (or `unaligned`) operations. Specify one and only one of `aligned` or `unaligned`.

**Caution**

If you specify `aligned` as an argument, you must be absolutely sure that the loop will be vectorizable using this instruction. Otherwise, the compiler will generate incorrect code.

The loop in the following example uses the `aligned` qualifier to request that the loop be vectorized with aligned instructions, as the arrays are declared in such a way that the compiler could not normally prove this would be safe to do so.

Example:

```
void foo (float *a)
{
    #pragma vector aligned
    for (i = 0; i < m; i++)
    {
        a[i] = a[i] * c;
    }
}
```

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 (but several other strategies are supported as well). If, in the loop, 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.

Alignment Strategies Example

```
float *a;
// alignment unknown
for (i = 0; i < 100; i++)
{
    a[i] = a[i] + 1.0f;
}

// dynamic loop peeling
p = a & 0x0f;
if (p != 0)
{
    p = (16 - p) / 4;
    for (i = 0; i < p; i++)
    {
        a[i] = a[i] + 1.0f;
    }
}

// loop with a aligned (will be vectorized accordingly)
for (i = p; i < 100; i++)
{
    a[i] = a[i] + 1.0f;
}
```

#pragma novector**Syntax:** #pragma novector

Definition: The `novector` loop pragma specifies that the loop should never be vectorized, even if it is legal to do so. In this example, suppose you know the trip count (`ub - lb`) is too low to make vectorization worthwhile. You can use `#pragma novector` to tell the compiler not to vectorize, even if the loop is considered vectorizable.

Example:

```
void foo (int lb, int ub)
{
    #pragma novector
    for (j = lb; j < ub; j++)
    {
        a[j] = a[j] + b[j];
    }
}
```

#pragma vector nontemporal**Syntax:** #pragma vector nontemporal

Definition: `#pragma vector nontemporal` results in streaming stores on Pentium® 4 based systems. An example loop (float type) together with the generated assembly are shown in the example. For large N, significant performance improvements result on a Pentium 4 systems over a non-streaming implementation.

Example:

```
#pragma vector nontemporal
for (i = 0; i < N; i++)
    a[i] = 1;
.B1.2:
movntps XMMWORD PTR _a[eax], xmm0
movntps XMMWORD PTR _a[eax+16], xmm0
add eax, 32
cmp eax, 4096
jl .B1.2
```

Dynamic Dependence Testing Example

```

float *p, *q;
for (i = L; i <= U; i++)
{
    p[i] = q[i];
}
...
pL = p + 4*L;
pH = p + 4*U;
qL = q + 4*L;
qH = q + 4*U;
if (pH < qL || pL > qH)
{
    // loop without data dependence
    for (i = L; i <= U; i++)
    {
        p[i] = q[i];
    } else {
        for (i = L; i <= U; i++)
        {
            p[i] = q[i];
        }
    }
}

```

Vectorization Examples

This section contains a few simple examples of some common issues in vector programming.

Argument Aliasing: A Vector Copy

The following loop example, a vector copy operation, vectorizes because the compiler can prove `dest[i]` and `src[i]` are distinct.

Vectorizable Copy Due To Unproven Distinction

```

void vec_copy(float *dest, float *src, int len)
{
    int i;
    for(i=0; i<len; i++;)
    {
        dest[i]=src[i];
    }
}

```

The `restrict` keyword in the following example indicates that the pointers refer to distinct objects. Therefore, the compiler allows vectorization without generation of multi-version code.

Using restrict to Prove Vectorizable Distinction

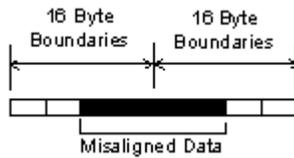
```
void vec_copy(float *restrict dest, float *restrict src, int
len)
{
    int i;
    for(i=0; i<len; i++)
    {
        dest[i]=src[i];
    }
}
```

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 sixteen.

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



For example, if you know that elements `a[0]` and `b[0]` are aligned on a 16-byte boundary, then the following loop can be vectorized with the alignment option on (`#pragma vector aligned`):

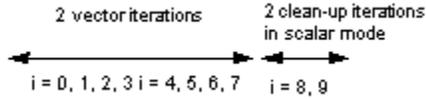
Alignment of Pointers is Known

```
float *a, *b;
int i;

for(int i=0; i<10; i++)
{
    a[i]=b[i];
}
```

After vectorization, the loop is executed as shown here:

Vector and Scalar Clean-up Iterations



Both the vector iterations $a[0:3] = b[0:3]$; and $a[4:7] = b[4:7]$; can be implemented with aligned moves if both the elements $a[0]$ and $b[0]$ (or, likewise, $a[4]$ and $b[4]$) are 16-byte aligned.



Caution

If you specify the vectorizer with incorrect alignment options, the compiler will generate unexpected behavior. Specifically, using aligned moves on unaligned data, will result in an illegal instruction exception.

Data Alignment Examples

This example contains a loop that vectorizes but only with unaligned memory instructions. The compiler can align the local arrays, but because lb is not known at compile-time. The correct alignment cannot be determined.

Loop Unaligned Due to Unknown Variable Value at Compile Time

```
void f(int lb)
{
    float z2[N], a2[N], y2[N], x2;
    for(i=lb; i<N; i++)
    {
        a2[i]=a2[i]*x2+y2[i];
    }
}
```

If you know that lb is a multiple of 4, you can align the loop with `#pragma vector aligned` as shown in the example that follows:

Alignment Due to Assertion of Variable as Multiple of 4

```
void f(int lb)
{
    float z2[N], a2[N], y2[N], x2;
    assert(lb%4==0);

    #pragma vector aligned

    for(i=lb; i<N; i++)
    {
        a2[i]=a2[i]*x2+y2[i];
    }
}
```

Loop Interchange and Subscripts: Matrix Multiply

Matrix multiplication is commonly written as shown in the following example:

Typical Matrix Multiplication

```
for(i=0; i<N; i++)
{
    for(j=0; j<n; j++)
    {
        for(k=0; k<n; k++)
        {
            c[i][j]=c[i][j]+a[i][k]*b[k][j];
        }
    }
}
```

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 shown in the "Matrix Multiplication With Stride-1" example.



Caution

Interchanging is not always possible because of dependencies, which can lead to different results.

Matrix Multiplication With Stride-1

```
for(i = 0; i<N; i++)
{
    for(k=0; k<n; k++)
    {
        for(j=0; j<n; j++)
        {
            c[i][j]=c[i][j]+a[i][k]*b[k][j];
        }
    }
}
```

Auto-parallelization

The auto-parallelization feature of the Intel® C++ 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 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

```
for (i=1; i<100; i++)
{
  a[i] = a[i] + b[i] * c[i];
}
```

Transformed Parallel Code

```
/* Thread 1 */
for (i=1; i<50; i++)
{
  a[i] = a[i] + b[i] * c[i];
}

/* Thread 2 */
for (i=50; i<100; i++)
{
  a[i] = a[i] + b[i] * c[i];
}
```

Programming with Auto-parallelization

The auto-parallelization feature implements some concepts of OpenMP*, such as worksharing construct (with the `parallel for` directive). 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, function calls, ambiguous indirect references, or global references.

Auto-parallelization Data Flow

For auto-parallelization processing, the compiler performs the following steps:

1. Data flow analysis
2. Loop classification
3. Dependence analysis
4. High-level parallelization
5. Data partitioning
6. 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 runtime routines for thread creation and synchronization

Auto-parallelization: Enabling, Options, 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 follows:

```
prompt>icpc -c -parallel prog.cpp
```

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.

Option	Description
<code>-parallel</code>	Enables the auto-parallelizer
<code>-par_threshold{1-100}</code>	Controls the work threshold needed for auto-parallelization. Default: <code>n=100</code> .
<code>-par_report{1 2 3}</code>	Controls the diagnostic messages from the auto-parallelizer.

Auto-parallelization Environment Variables

Variable	Description	Default
<code>OMP_NUM_THREADS</code>	Controls the number of threads used.	Number of processors currently installed in the system while generating the executable
<code>OMP_SCHEDULE</code>	Specifies the type of runtime scheduling.	<code>static</code>

Auto-parallelization Threshold Control and Diagnostics

Threshold Control

The `-par_threshold[n]` option sets a threshold for the 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. This option is used for loops whose computation work volume cannot be determined at compile time. The threshold is usually relevant when the loop trip count is unknown at compile time.

The `-par_threshold[n]` option has the following functionality:

- `-par_threshold100` is executed by default, so loops get auto-parallelized only if profitable parallel execution is almost certain.
- If you specify `-par_threshold` with designating a value for `n`, the compiler uses the default value `n=100`.
- The intermediate 1 to 99 values represent the percentage probability for profitable speed-up. For example, `n=50` directs the compiler to 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 dependencies inhibiting auto-parallelization (reasons for not parallelizing).

Example of Parallelization Diagnostics Report

This example shows output generated by `-par_report3`:

```
prompt>icpc -c -parallel -par_report3 prog.cpp
```

Sample Output

```
program prog
procedure: prog
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
```

where the program `prog.cpp` is as follows:

Sample prog.c

```
/* Assumed side effects */
for (i=1; i<10000; i++)
{
  a[i] = foo(i);
}

/* Actual dependence */
for (i=1; i<10000; i++)
{
  a[i] = a[i-1] + i;
}
```

Troubleshooting Tips

- Use `-par_threshold0` to see if the compiler assumed there was not enough computational work
- Use `-par_report3` to view diagnostics
- Use `-ipo[value]` to eliminate assumed side-effects done to function calls

Parallelization with OpenMP*

The Intel® C++ Compiler supports the OpenMP* C++ version 2.0 API specification. 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.

The Intel C++ 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 C++ Compiler provides Intel-specific extensions to the OpenMP C++ 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 of the options of the OpenMP feature in the Intel C++ Compiler.

For complete information on the OpenMP standard, visit the OpenMP Web site at <http://www.openmp.org>. For OpenMP* C++ version 2.0 API specifications, see <http://www.openmp.org/specs/>.

Parallel Processing with OpenMP

To compile with OpenMP, you need to prepare your program by annotating the code with OpenMP directives. The Intel C++ Compiler first processes the application and produces a multithreaded version of the code which is then compiled. The output is an executable program with the parallelism implemented by threads that execute parallel regions or constructs.

Targeting a Processor Run-time Check

While parallelizing a loop, the Intel compiler's loop parallelizer, OpenMP, tries to determine the optimal set of configurations for a given processor. At run time, a check is performed to determine for which IA-32 processor OpenMP should optimize a given loop. See detailed information in the Processor-specific Runtime Checks, IA-32 Systems.

Performance Analysis

For performance analysis of your program, you can use the Intel® VTune(TM) Performance Analyzer to show performance information. You can obtain detailed information about which portions of the code require the largest amount of time to execute and where parallel performance problems are located.

Parallel Processing Thread Model

This topic explains the processing of the parallelized program and adds more definitions of the terms used in parallel programming.

The Execution Flow

As previously mentioned, a program containing OpenMP* C++ 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 the OpenMP C++ API, the `#pragma omp parallel` directive defines 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 `#pragma omp parallel` directive reaches completion, the threads in the team synchronize, 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 and use directives to control execution in any of the called routines. For example:

```
int main(void)
{
    ...
    #pragma omp parallel
    {
        phase1();
    }
}

void phase1(void)
{
    ...
    #pragma omp for private(i) shared(n)
    for(i=0; i < n; i++)
    {
        some_work(i);
    }
}
```

This is an orphaned directive because the parallel region is not lexically present.

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 scope variables by using the `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.

Pseudo Code of the Parallel Processing Model

A sample pseudo 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.

```

main() {
    ...
    #pragma omp parallel // Begin a Parallel Construct, form
    {
        ... // (each team member executes
        ... // the same code)
        //
    #pragma omp sections // Begin a Worksharing Construct
    {
        #pragma omp section // One unit of work
        {...} //
        #pragma omp section // Another unit of work
        {...} //
    } // Wait until both units of work
    ... // More Replicated Code

```

```

//
#pragma omp for // Begin a Worksharing Construct;
nowait
  for(...) { // each iteration is unit of work
//
  ... // Work is distributed among the team
members
//
} // End of Worksharing Construct;
// nowait was specified, so
// threads proceed
//
#pragma omp critical // Begin a Critical Section
{
//
... // Replicated Code, but only one
// thread can execute it at a
} // given time
... // More Replicated Code
//
#pragma omp barrier // Wait for all team members to arrive
... // More Replicated Code
//
} // End of Parallel Construct;
// disband team and continue
// serial execution
//
... // Possibly more Parallel constructs
//
} // End serial execution
```

Compiling with OpenMP, Directive Format, and Diagnostics

To run the Intel® C++ Compiler in OpenMP* mode, invoke the compiler with the `-openmp` option:

```
prompt>icpc -openmp file.cpp
```

Before you run the multithreaded code, you can set the number of desired threads in the OpenMP environment variable, `OMP_NUM_THREADS`. See OpenMP Environment Variables for further information.

-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.

OpenMP Directive Format and Syntax

An OpenMP directive has the form:

```
#pragma omp directive-name [clause, ...] newline
```

where:

- `#pragma omp` -- Required for all OpenMP directives.
- `directive-name` -- A valid OpenMP directive. Must appear after the `pragma` and before any clauses.
- `clause` -- Optional. Clauses can be in any order, and repeated as necessary unless otherwise restricted.
- `newline` -- Required. Proceeds the structured block which is enclosed by this directive.

OpenMP Diagnostics

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. are successfully handled.

The default is `-openmp_report1`.

OpenMP* Directives and Clauses

OpenMP Directives

Directive Name	Description
<code>parallel</code>	Defines a parallel region.
<code>for</code>	Identifies an iterative work-sharing construct that specifies a region in which the iterations of the associated loop should be executed in parallel.
<code>sections</code>	Identifies a non-iterative work-sharing construct that specifies a set of constructs that are to be divided among threads in a team.
<code>single</code>	Identifies a construct that specifies that the associated structured block is executed by only one thread in the team.
<code>parallel for</code>	A shortcut for a <code>parallel</code> region that contains a <code>single for</code> directive. The <code>parallel</code> or <code>for</code> OpenMP directive must be immediately followed by a <code>for</code> statement. If you place other statement or an OpenMP directive between the <code>parallel</code> or <code>for</code> directive and the <code>for</code> statement, the Intel C++ Compiler issues a syntax error.
<code>parallel sections</code>	Provides a shortcut form for specifying a parallel region containing a single <code>sections</code> directive.
<code>master</code>	Identifies a construct that specifies a structured block that is executed by the master thread of the team.
<code>critical[lock]</code>	Identifies a construct that restricts execution of the associated structured block to a single thread at a time.
<code>barrier</code>	Synchronizes all the threads in a team.
<code>atomic</code>	Ensures that a specific memory location is updated atomically.
<code>flush</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.
<code>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</code>	Makes the named file-scope or namespace-scope variables specified private to a thread but file-scope visible within the thread.

OpenMP Clauses

Clause	Description
<code>private</code>	Declares variables to be <code>private</code> to each thread in a team.
<code>firstprivate</code>	Provides a superset of the functionality provided by the <code>private</code> clause.
<code>lastprivate</code>	Provides a superset of the functionality provided by the <code>private</code> clause.
<code>shared</code>	Shares variables among all the threads in a team.
<code>default</code>	Enables you to affect the data-scope attributes of variables.
<code>reduction</code>	Performs a reduction on scalar variables.
<code>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>if</code>	If the <code>if(scalar_logical_expression)</code> clause is present, the enclosed code block is executed in parallel only if the <code>scalar_logical_expression</code> evaluates to <code>TRUE</code> . Otherwise the code block is serialized.
<code>schedule</code>	Specifies how iterations of the <code>for</code> loop are divided among the threads of the team.
<code>copyin</code>	Provides a mechanism to assign the same name to <code>threadprivate</code> variables for each thread in the team executing the parallel region.

OpenMP* Support Libraries

The Intel® C++ 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.



Note

The `libguide.lib` library is linked dynamically, regardless of command-line options, to avoid performance issues that are hard to debug.

Execution modes

The Intel 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.

Serial

The serial mode forces parallel applications to run on a single processor.

Turnaround

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.

Throughput

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. Throughput mode is the default.

OpenMP* Environment Variables

This topic describes the OpenMP* environment variables (with the `OMP_` prefix) and Intel-specific environment variables (with the `KMP_` prefix).

Standard Environment Variables

Variable	Description	Default
<code>OMP_SCHEDULE</code>	Sets the runtime 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
KMP_LIBRARY	Selects the OpenMP run-time library throughput. The options for the variable value are: <code>serial</code> , <code>turnaround</code> , or <code>throughput</code> indicating the execution mode. The default value of <code>throughput</code> is used if this variable is not specified.	<code>throughput</code> (execution mode)
KMP_STACKSIZE	Sets the number of bytes to allocate for each parallel thread to use as its private stack. Use the optional suffix <code>b</code> , <code>k</code> , <code>m</code> , <code>g</code> , or <code>t</code> , to specify bytes, kilobytes, megabytes, gigabytes, or terabytes.	IA-32: 2m Itanium® compiler: 4m

OpenMP* Run-time Library Routines

OpenMP* provides several run-time library functions to assist you in managing your program in parallel mode. Many of these functions have corresponding environment variables that can be set as defaults. The run-time library functions enable you to dynamically change these factors to assist in controlling your program. In all cases, a call to a run-time library function overrides any corresponding environment variable.

The following table specifies the interfaces to these routines. The names for the routines are in user name space. The `omp.h` and `omp_lib.h` header files are provided in the `INCLUDE` directory of your compiler installation.

There are definitions for two different locks, `omp_lock_kind` and `omp_nest_lock_kind`, which are used by the functions in the table that follows:

Execution Environment Routines

Function	Description
<code>omp_set_num_threads(<i>nthreads</i>)</code>	Sets the number of threads to use for subsequent parallel regions.
<code>omp_get_num_threads()</code>	Returns the number of threads that are being used in the current parallel region.
<code>omp_get_max_threads()</code>	Returns the maximum number of threads that are available for parallel execution.
<code>omp_get_thread_num()</code>	Returns the unique thread number of the thread currently executing this section of code.
<code>omp_get_num_procs()</code>	Returns the number of processors available to the program.

Function	Description
<code>omp_in_parallel()</code>	Returns TRUE if called within the dynamic extent of a parallel region executing in parallel; otherwise returns FALSE.
<code>omp_set_dynamic(<i>dynamic_threads</i>)</code>	Enables or disables dynamic adjustment of the number of threads used to execute a parallel region. If <i>dynamic_threads</i> is TRUE, dynamic threads are enabled. If <i>dynamic_threads</i> is FALSE, dynamic threads are disabled. Dynamic threads are disabled by default.
<code>omp_get_dynamic()</code>	Returns TRUE if dynamic thread adjustment is enabled, otherwise returns FALSE.
<code>omp_set_nested(<i>nested</i>)</code>	Enables or disables nested parallelism. If <i>nested</i> is TRUE, nested parallelism is enabled. If <i>nested</i> is FALSE, nested parallelism is disabled. Nested parallelism is disabled by default.
<code>omp_get_nested()</code>	Returns TRUE if nested parallelism is enabled, otherwise returns FALSE.

Lock Routines

Function	Description
<code>omp_init_lock(<i>lock</i>)</code>	Initializes the lock associated with <i>lock</i> for use in subsequent calls.
<code>omp_destroy_lock(<i>lock</i>)</code>	Causes the lock associated with <i>lock</i> to become undefined.
<code>omp_set_lock(<i>lock</i>)</code>	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.
<code>omp_unset_lock(<i>lock</i>)</code>	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> .
<code>omp_test_lock(<i>lock</i>)</code>	Attempts to set the lock associated with <i>lock</i> . If successful, returns TRUE, otherwise returns FALSE.
<code>omp_init_nest_lock(<i>lock</i>)</code>	Initializes the nested lock associated with <i>lock</i> for use in the subsequent calls.
<code>omp_destroy_nest_lock(<i>lock</i>)</code>	Causes the nested lock associated with <i>lock</i> to become undefined.

Function	Description
<code>omp_set_nest_lock(lock)</code>	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.
<code>omp_unset_nest_lock(lock)</code>	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> .
<code>omp_test_nest_lock(lock)</code>	Attempts to set the nested lock associated with <i>lock</i> . If successful, returns the nesting count, otherwise returns zero.

Timing Routines

Function	Description
<code>omp_get_wtime()</code>	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.
<code>omp_get_wtick()</code>	Returns a double-precision value equal to the number of seconds between successive clock ticks.

Intel Extensions

The Intel® C++ Compiler implements the following groups of functions as extensions to the OpenMP* run-time library:

- getting and setting stack size for parallel threads
- memory allocation

The Intel extensions 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 functions with caution because using them requires the use of the `-openmp_stubs` command-line option to execute the program sequentially. These functions are also generally not recognized by other vendor's OpenMP-compliant compilers, which may cause the link stage to fail for these other compilers.



Note

The following functions require the pre-processor directive `#include <omp.h>`.

Stack Size

In most cases, directives can be used in place of extensions. For example, the stack size of the parallel threads may be set using the `KMP_STACKSIZE` environment variable rather than the `kmp_set_stacksize_s()` function.



Note

A run-time call to an Intel extension takes precedence over the corresponding environment variable setting. See the definitions of stack size functions in the Stack Size table.

Memory Allocation

The Intel® C++ Compiler implements a group of memory allocation functions as extensions to the OpenMP run-time library to enable threads to allocate memory from a heap local to each thread. These functions are `kmp_malloc()`, `kmp_calloc()`, and `kmp_realloc()`. The memory allocated by these functions must also be freed by the `kmp_free()` function. 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 functions in the Memory Allocation table.

Stack Size

Function	Description
<code>kmp_get_stacksize_s()</code>	Returns the number of bytes that will be allocated for each parallel thread to use as its private stack. This value can be changed with <code>kmp_set_stacksize_s()</code> prior to the first parallel region or with the <code>KMP_STACKSIZE</code> environment variable.
<code>kmp_get_stacksize()</code>	This function is provided for backwards compatibility only. Use <code>kmp_get_stacksize_s()</code> for compatibility across different families of Intel processors.
<code>kmp_set_stacksize_s(size)</code>	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.
<code>kmp_set_stacksize(size)</code>	This function is provided for backward compatibility only; use <code>kmp_set_stacksize_s()</code> for compatibility across different families of Intel processors.

Memory Allocation

Function	Description
<code>kmp_malloc(size)</code>	Allocate memory block of <i>size</i> bytes from thread-local heap.
<code>kmp_calloc(nelem, elsize)</code>	Allocate array of <i>nelem</i> elements of size <i>elsize</i> from thread-local heap.
<code>kmp_realloc(ptr, size)</code>	Reallocate memory block at address <i>ptr</i> and <i>size</i> bytes from thread-local heap.
<code>kmp_free(ptr)</code>	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> .

Intel Workqueuing Model

The workqueuing model lets you parallelize control structures that are beyond the scope of those supported by the OpenMP* model, while attempting to fit into the framework defined by OpenMP. In particular, the workqueuing model is a flexible mechanism for specifying units of work that are not pre-computed at the start of the worksharing construct. For `single`, `for`, and `sections` constructs all work units that can be executed are known at the time the construct begins execution. The workqueuing pragmas `taskq` and `task relax` this restriction by specifying an environment (the `taskq`) and the units of work (the `tasks`) separately.

Workqueuing Constructs

taskq Pragma

The `taskq` pragma specifies the environment within which the enclosed units of work (`tasks`) are to be executed. From among all the threads that encounter a `taskq` pragma, one is chosen to execute it initially. Conceptually, the `taskq` pragma causes an empty queue to be created by the chosen thread, and then the code inside the `taskq` block is executed single-threaded. All the other threads wait for work to be enqueued on the conceptual queue. The `task` pragma specifies a unit of work, potentially executed by a different thread. When a `task` pragma is encountered lexically within a `taskq` block, the code inside the `task` block is conceptually enqueued on the queue associated with the `taskq`. The conceptual queue is disbanded when all work enqueued on it finishes, and when the end of the `taskq` block is reached.

Control Structures

Many control structures exhibit the pattern of separated work iteration and work creation, and are naturally parallelized with the workqueuing model. Some common cases are:

- `while` loops
- C++ iterators
- recursive functions.

while Loops

If the computation in each iteration of a `while` loop is independent, the entire loop becomes the environment for the `taskq` pragma, and the statements in the body of the `while` loop become the units of work to be specified with the `task` pragma. The conditional in the `while` loop and any modifications to the control variables are placed outside of the `task` blocks and executed sequentially to enforce the data dependencies on the control variables.

C++ Iterators

C++ Standard Template Library (STL) iterators are very much like the `while` loops just described, whereby the operations on the data stored in the STL are very distinct from the act of iterating over all the data. If the operations are data-independent, they can be done in parallel as long as the iteration over the work is sequential. This type of `while` loop parallelism is a generalization of the standard OpenMP* worksharing for loops. In the worksharing for loops, the loop increment operation is the iterator and the body of the loop is the unit of work. However, because the `for` loop iteration variable frequently has a closed form solution, it can be computed in parallel and the sequential step avoided.

Recursive Functions

Recursive functions also can be used to specify parallel iteration spaces. The mechanism is similar to specifying parallelism using the `sections` pragma, but is much more flexible because it allows arbitrary code to sit between the `taskq` and the `task` pragmas, and because it allows recursive nesting of the function to build a conceptual tree of `taskq` queues. The recursive nesting of the `taskq` pragmas is a conceptual extension of OpenMP worksharing constructs to behave more like nested OpenMP parallel regions. Just like nested parallel regions, each nested workqueuing construct is a new instance and is encountered by exactly one thread. However, the major difference is that nested workqueuing constructs do not cause new threads or teams to be formed, but rather re-use the threads from the team. This permits very easy multi-algorithmic parallelism in dynamic environments, such that the number of threads need not be committed at each level of parallelism, but instead only at the top level. From that point on, if a large amount of work suddenly appears at an inner level, the idle threads from the outer level can assist in getting that work finished. For example, it is very common in server environments to dedicate a thread to handle each incoming request, with a large number of threads awaiting incoming requests. For a particular request, its size may not be obvious at the time the thread begins handling it. If the thread uses nested workqueuing constructs, and the scope of the request becomes large after the inner construct is started, the threads from the outer construct can easily migrate to the inner construct to help finish the request.

Since the workqueuing model is designed to preserve sequential semantics, synchronization is inherent in the semantics of the `taskq` block. There is an implicit team barrier at the completion of the `taskq` block for the threads that encountered the `taskq` construct to ensure that all of the tasks specified inside of the `taskq` block have finished execution. This `taskq` barrier enforces the sequential semantics of the original program. Just like the OpenMP worksharing constructs, it is assumed you are responsible for ensuring that either no dependences exist or that dependencies are appropriately synchronized between the task blocks, or between code in a task block and code in the `taskq` block outside of the task blocks.

The syntax, semantics, and allowed clauses are designed to resemble OpenMP* worksharing constructs. Most of the clauses allowed on OpenMP worksharing constructs have a reasonable meaning when applied to the workqueuing pragmas.

taskq Construct

```
#pragma intel omp taskq [clause[[,]clause]...]
    structured-block
```

where `clause` can be any of the following:

- `private` (`variable-list`)
- `firstprivate` (`variable-list`)
- `lastprivate` (`variable-list`)
- `reduction` (`operator` : `variable-list`)
- `ordered`
- `nowait`

private

The `private` clause creates a private, default-constructed version for each object in `variable-list` for the `taskq`. It also implies `captureprivate` on each enclosed task. The original object referenced by each variable has an indeterminate value upon entry to the construct, must not be modified within the dynamic extent of the construct, and has an indeterminate value upon exit from the construct.

firstprivate

The `firstprivate` clause creates a private, copy-constructed version for each object in `variable-list` for the `taskq`. It also implies `captureprivate` on each enclosed task. The original object referenced by each variable must not be modified within the dynamic extent of the construct and has an indeterminate value upon exit from the construct.

lastprivate

The `lastprivate` clause creates a private, default-constructed version for each object in `variable-list` for the `taskq`. It also implies `captureprivate` on each enclosed task. The original object referenced by each variable has an indeterminate value upon entry to the construct, must not be modified within the dynamic extent of the construct, and is copy-assigned the value of the object from the last enclosed task after that task completes execution.

reduction

The `reduction` clause performs a reduction operation with the given operator in enclosed task constructs for each object in `variable-list`. `operator` and `variable-list` are defined the same as in the OpenMP Specifications.

ordered

The `ordered` clause performs ordered constructs in enclosed `task` constructs in original sequential execution order. The `taskq` directive, to which the `ordered` is bound, must have an `ordered` clause present.

nowait

The `nowait` clause removes the implied barrier at the end of the `taskq`. Threads may exit the `taskq` construct before completing all the `task` constructs queued within it.

task Construct

```
#pragma intel omp task [clause[[,]clause]...]
    structured-block
```

where `clause` can be any of the following:

- `private(variable-list)`
- `captureprivate(variable-list)`

private

The `private` clause creates a private, default-constructed version for each object in `variable-list` for the `task`. The original object referenced by the variable has an indeterminate value upon entry to the construct, must not be modified within the dynamic extent of the construct, and has an indeterminate value upon exit from the construct.

captureprivate

The `captureprivate` clause creates a private, copy-constructed version for each object in `variable-list` for the `task` at the time the `task` is enqueued. The original object referenced by each variable retains its value but must not be modified within the dynamic extent of the `task` construct.

Combined parallel and taskq Construct

```
#pragma intel omp parallel taskq [clause[[,]clause]...]
    structured-block
```

where `clause` can be any of the following:

- `if(scalar-expression)`
- `num_threads(integer-expression)`
- `copyin(variable-list)`
- `default(shared | none)`
- `shared(variable-list)`
- `private(variable-list)`
- `firstprivate(variable-list)`
- `lastprivate(variable-list)`
- `reduction(operator : variable-list)`
- `ordered`

Clause descriptions are the same for `parallel` and `taskq` construct.

Example Function

The `test1` function is a natural candidate to be parallelized using the workqueuing model. You can express the parallelism by annotating the loop with a parallel `taskq` pragma and the work in the loop body with a `task` pragma. The parallel `taskq` pragma specifies an environment for the `while` loop in which to enqueue the units of work specified by the enclosed `task` pragma. Thus, the loop's control structure and the enqueueing are executed single-threaded, while the other threads in the team participate in dequeuing the work from the `taskq` queue and executing it. The `captureprivate` clause ensures that a private copy of the link pointer `p` is captured at the time each task is being enqueued, hence preserving the sequential semantics.

```
void test1(LIST p)
{
    #pragma intel omp parallel taskq shared(p)
    {
        while (p != NULL)
        {
            #pragma intel omp task captureprivate(p)
            {
                do_work1(p);
            }
            p = p->next;
        }
    }
}
```

Examples of OpenMP* Usage

The following examples show how to use the OpenMP* feature.

A Simple Difference Operator

This example shows a simple parallel loop where the amount of work in each iteration is different. Dynamic scheduling is used to get good load balancing. The `for` has a `nowait` because there is an implicit `barrier` at the end of the parallel region.

```
void for_1 (float a[], float b[], int n)
{
    int i, j;
    #pragma omp parallel shared(a,b,n) private(i,j)
    {
        #pragma omp for schedule(dynamic,1) nowait
        for(i = 1; i < n; i++)
        {
            for(j = 0; j <= i; j++)
                b[j + n*i] = (a[j + n*i] + a[j + n*(i-1)])/2.0;
        }
    }
}
```

Two Difference Operators

The following example uses two parallel loops fused to reduce fork/join overhead. The first `for` has a `nowait` because all the data used in the second loop is different than all the data used in the first loop.

```
void for_2 (float a[], float b[], float c[], \
float d[], int n, int m)
{
    int i, j;
    #pragma omp parallel shared(a,b,c,d,n,m) private(i,j)
    {
        #pragma omp for schedule(dynamic,1) nowait
        for(i = 1; i < n; i++)
        {
            for(j = 0; j <= i; j++)
                b[j + n*i] = ( a[j + n*i] + a[j + n*(i-1)] )/2.0;
        }

        #pragma omp for schedule(dynamic,1) nowait
        for(i = 1; i < m; i++)
        {
            for(j = 0; j <= i; j++)
                d[j + m*i] = ( c[j + m*i] + c[j + m*(i-1)] )/2.0;
        }
    }
}
```

Optimization Support Features

This section describes language extensions to the Intel® C++ Compiler that let you optimize your source code directly. Examples are included of optimizations supported by Intel extended directives and library routines that enhance and/or help analyze performance.

Compiler Directives

This section discusses the language extended directives used in:

- Software Pipelining
- Loop Count and Loop Distribution
- Loop Unrolling
- Prefetching
- Vectorization

Pipelining for Itanium®-based Applications

The `swp` and `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:

```
#pragma swp
```

```
#pragma noswp
```

Example of swp Directive

```
#pragma swp
for (i=0; i<m ; i++)
{
    if (a[i]==0)
    {
        b[i]=a[i]+1;
    }
    else
    {
        b[i]=a[i]*2;
    }
}
```

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. You can request and view the optimization report to see whether software pipelining was applied (see [Optimizer Report Generation](#)).

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:

```
#pragma loop count (n)
```

where `n` is an integer constant. The value of `loop count` affects heuristics used in software pipelining, vectorization and loop-transformations.

Example of loop count (n) Directive

```
#pragma loop count (10000)

for(i=0; i<m; i++)
{
    //swp likely to occur in this loop
    a[i]=b[i]+1.2;
}
```

distribute point Directive

The `distribute point` directive indicates to the compiler a preference of performing loop distribution. The syntax for this directive is:

```
#pragma distribute point
```

Loop distribution may cause large loops be distributed into smaller ones. This may enable software pipelining for more loops. 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. Only one `distribute` directive is supported when placed inside the loop.

Example of distribute point Directive

```
#pragma distribute point

for(i=1; i<m; i++)
{
    b[i]=a[i]+1;
    ...

    //Compiler will automatically
    //decide where to distribute.
    //Data dependency is observed.

    c[i]=a[i]+b[i];
    ...
    d[i]=c[i]+1;
}

for(i=1; i<m; i++)
{
    b[i]=a[i]+1;
    ...

    #pragma distribute point

    //Distribution will start here,
    //ignoring all loop-carried dependency.

    sub(a,n);
    c[i]=a[i]+b[i];
    ...
    d[i]=c[i]+1;
}
```

Loop Unrolling Support

unroll Directive

The `unroll` directive (`unroll(n) | nounroll`) tells the compiler how many times to unroll a counted loop. The syntax for this directive is:

```
#pragma unroll

#pragma unroll(n)

#pragma nounroll
```

where `n` is an integer constant from 0 through 255. The `unroll` directive must precede the `for` statement for each `for` 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. The directive can be applied only for the innermost nested loop. If applied to the outer loops, it is ignored. The compiler generates correct code by comparing `n` and the loop count.

Example of unroll Directive

```
#pragma unroll(4)

for(i=1; i<m; i++)
{
    b[i]=a[i]+1;
    d[i]=c[i]+1;
}
```

Prefetching Support

prefetch Directive

The `prefetch` and `noprefetch` directives assert that the data prefetches are generated or not generated for some memory references. This affects the heuristics used in the compiler. The syntax for this directive is:

```
#pragma noprefetch

#pragma prefetch

#pragma prefetch a,b
```

If the expression `a[j]` is used within a loop, by placing `prefetch a` in front of the loop, the compiler will insert prefetches for `a[j+d]` within the loop, where `d` is determined by the compiler. This directive is supported when option `-O3` is on.

Example of prefetch Directive

```
#pragma noprefetch b
#pragma prefetch a

for(i=0; i<m; i++)
{
    a[i]=b[i]+1;
}
```

Vectorization Support (IA-32)

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.

vector always Directive

The `vector always` directive instructs the compiler to override any efficiency heuristic during the decision to vectorize or not, and will vectorize non-unit strides or very unaligned memory accesses.

Example of vector always Directive

```
#pragma vector always

for(i=0; i<=N; i++)
{
    a[32*i]=b[99*i];
}
```

ivdep Directive

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. The loop in the following example will not vectorize with the `ivdep`, since the value of `k` is not known (vectorization would be illegal if $k < 0$).

Example of ivdep Directive

```
#pragma ivdep

for(i=0; i<m; i++)
{
    a[i]=a[i+k]*c;
}
```

vector aligned Directive

The `vector aligned` directive means the loop should be vectorized, if it is legal to do so, ignoring normal heuristic decisions about profitability. When the `aligned` or `unaligned` qualifier is used, the loop should be vectorized using `aligned` or `unaligned` operations. Specify either `aligned` or `unaligned`, but not both.



If you specify `aligned` as an argument, you must be absolutely sure that the loop will be vectorizable using this instruction. Otherwise, the compiler will generate incorrect code. The loop in the following example uses the `aligned` qualifier to request that the loop be vectorized with `aligned` instructions, as the arrays are declared in such a way that the compiler could not normally prove this would be safe to do so.

Example of vector aligned Directive

```
#void foo(float *a)
{
    #pragma vector aligned
    for(i=0; i<m; i++)
    {
        a[i]=a[i]*c;
    }
}
```

The compiler includes several alignment strategies in case the alignment of data structures is not known at compile time. A simple example follows, but several other strategies are supported as well. If, in the following loop, 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.

Example of Alignment Strategies

```

float *a;

//Alignment unknown
for(i=0; i<100; i++)
{
    a[i]=a[i]+1.0f;
}

//Dynamic loop peeling
p=a & 0x0f;
if(p!=0)
{
    p=(16-p)/4;
    for(i=0; i<p; i++)
    {
        a[i]=a[i]+1.0f;
    }
}

//Loop with a aligned.
//Will be vectorized accordingly.
for(i=p; i<100; i++)
{
    a[i]=a[i]+1.0f;
}

```

novector Directive

The `novector` directive specifies that the loop should never be vectorized, even if it is legal to do so. In this example, suppose you know the trip count (`ub - lb`) is too low to make vectorization worthwhile. You can use `novector` to tell the compiler not to vectorize, even if the loop is considered vectorizable.

Example of novector Directive

```

void foo(int lb, int ub)
{
    #pragma novector
    for(j=lb; j<ub; j++)
    {
        a[j]=a[j]+b[j];
    }
}

```

Optimizer Report Generation

The Intel® C++ Compiler provides options to generate and manage optimization reports:

- `-opt_report` generates an optimization report and directs it to `stderr`. By default, the compiler does not generate optimization reports.
- `-opt_report_filefilename` generates an optimization report and directs it to a file specified in *filename*.
- `-opt_report_level{min/med/max}` specifies the detail level of the optimization report. The *min* argument provides the minimal summary and *max* produces the full report. The default is `-opt_report_levelmin`.
- `-opt_report_routinefileroutine_substring` generates reports from all routines with names containing the *substring* as part of their name. If not specified, reports from all routines are generated. By default, the compiler generates reports for all routines.

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 following optimizer reports are supported.

Optimizer Logical Name	Optimizer Full Name
<code>ipo</code>	Interprocedural Optimizer
<code>hlo</code>	High Level Optimizer
<code>ilo</code>	Intermediate Language Scalar Optimizer
<code>ecg</code>	Code Generator
<code>omp</code>	Open MP
<code>all</code>	All phases

When one of the logical names for optimizers is specified, all reports from that optimizer are generated.

For example, `-opt_report_phaseipo -opt_report_phaseecg` generates 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 one of the optimizer logical names. For example:

Optimizer_optimization	Full Name
<code>ipo_inline</code>	Interprocedural Optimizer, inline expansion of functions
<code>ipo_constant_propagation</code>	Interprocedural Optimizer, constant propagation
<code>ipo_function_reorder</code>	Interprocedural Optimizer, function reorder
<code>ilo_constant_propagation</code>	Intermediate Language Scalar Optimizer, constant propagation
<code>ilo_copy_propagation</code>	Intermediate Language Scalar Optimizer, copy propagation
<code>ecg_software_pipelining</code>	Code Generator, software pipelining

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 available for report generation.

Timing Your Application

How fast your application executes is one indication of performance. When timing the speed of applications, consider the following circumstances:

- 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 system (processor 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. Certain overhead functions, like loading external programs, might influence short timings considerably.
- If your program displays a lot of text, consider redirecting the output from the program. Redirecting output from the program will change the times reported because of reduced screen I/O.

The following program illustrates a model for program timing:

```

/* Sample Timing */
#include <stdio.h>
#include <stdlib.h>
#include <time.h>
int main(void)
{
    clock_t start, finish;
    long loop;
    double duration, loop_calc;
    start = clock();
    for(loop=0; loop <= 2000; loop++)
    {
        loop_calc = 123.456 * 789;

        //printf() included to facilitate example
        printf("\nThe value of loop is: %d", loop);
    }
    finish = clock();
    duration = (double)(finish - start)/CLOCKS_PER_SEC;
    printf("\n%2.3f seconds\n", duration);
}

```

Reference

Compiler Limits

The following table shows the size or number of each item that the compiler can process. All capacities shown in the table are tested values; the actual number can be greater than the number shown.

Item	Tested Values
Control structure nesting (block nesting)	512
Conditional compilation nesting	512
Declarator modifiers	512
Parenthesis nesting levels	512
Significant characters, internal identifier	2048
External identifier name length	64K
Number of external identifiers/file	128K
Number of identifiers in a single block	2048
Number of macros simultaneously defined	128K
Number of parameters to a function call	512
Number of parameters per macro	512
Number of characters in a string	128K
Bytes in an object	512K
Include file nesting depth	512
Case labels in a switch	32K
Members in one structure or union	32K
Enumeration constants in one enumeration	8192
Levels of structure nesting	320
Size of arrays	2 GB

Key Files

Key Files Summary for IA-32 Compiler

The following tables list and briefly describe files that are installed for use by the IA-32 version of the compiler.

/bin Files

File	Description
codecov	Code-coverage tool
iccvars.sh iccvars.csh	Batch file to set environment variables
icc icpc	Scripts that check for license file and call compiler driver
iccbin icpcbin	Compiler drivers
mcpcom	Intel® C++ Compiler
iccbin icpcbin	Compiler drivers
profmerge	Utility used for Profile Guided Optimizations
proforder	Utility used for Profile Guided Optimizations
tselect	Test-prioritization tool
xiar	Tool used for Interprocedural Optimizations
xild	Tool used for Interprocedural Optimizations

/include Files

File	Description
dvec.h	SSE 2 intrinsics for Class Libraries
emm_func.h	Header file for SSE2 intrinsics (used by emmintrin.h)
emmintrin.h	Principal header file for SSE2 intrinsics
float.h	IEEE 754 version of standard float.h
fvec.h	SSE intrinsics for Class Libraries
iso646.h	Standard header file
ivec.h	MMX(TM) instructions intrinsics for Class Libraries
limits.h	Standard header file

File	Description
mathf.h	Principal header file for legacy Intel Math Library
mathimf.h	Principal header file for current Intel Math Library
mmintrin.h	Intrinsics for MMX instructions
omp.h	Principal header file OpenMP*
omp_lib.h	Header file for OpenMP
pgouser.h	For use in the instrumentation compilation phase of profile-guided optimizations
pmmintrin.h	Principal header file SSE3 intrinsics
proto.h	
sse2mmx.h	Principal header file for Streaming SIMD Extensions 2 intrinsics
stdarg.h	Replacement header for standard stdarg.h
stdbool.h	Defines <code>_Bool</code> keyword
stddef.h	Standard header file
syslimits.h	
varargs.h	Replacement header for standard varargs.h
xarg.h	Header file used by stdargs.h and varargs.h
xmm_func.h.h	Header file for Streaming SIMD Extensions
xmm_utils.h	Utilities for Streaming SIMD Extensions
xmmintrin.h	Principal header file for Streaming SIMD Extensions intrinsics

/lib Files

Library	Description
libguide.a libguide.so	For OpenMP* implementation
libguide_stats.a libguide_stats.so	OpenMP static library for the parallelizer tool with performance statistics and profile information
libompstub.a	Library that resolves references to OpenMP subroutines when OpenMP is not in use
libsvml.a	Short vector math library
libirc.a	Intel support library for PGO and CPU dispatch
libimf.a	Intel math library
libimf.so	Intel math library

Library	Description
libcprts.a libcprts.so libcprts.so.3	Dinkumware* C++ Library
libunwind.a libunwind.so libunwind.so.3	Unwinder library
libcxa.a libcxa.so libcxa.so.3	Intel run time support for C++ features
libcxaguard.a libcxaguard.so libcxaguard.so.3	Used for interoperability support with the <code>-cxxlib-gcc</code> option. See gcc Interoperability.

Key Files Summary for Itanium® Compiler

The following tables list and briefly describe files that are installed for use by the Itanium® compiler.

/bin Files

File	Description
codecov	Code-coverage tool
iccvars.sh	Batch file to set environment variables
icc.cfg	Configuration file for use from command line
icc icpc	Scripts that check for license file and call compiler driver
iccbin icpcbin	Compiler drivers
mcpcom	Intel® C++ Compiler
iccbin icpcbin	Compiler drivers
profmerge	Utility used for Profile Guided Optimizations
proforder	Utility used for Profile Guided Optimizations
tselect	Test-prioritization tool
xiar	Tool used for Interprocedural Optimizations
xild	Tool used for Interprocedural Optimizations

/include Files

File	Description
<code>emmintrin.h</code>	Principal header file for SSE2 intrinsics
<code>float.h</code>	IEEE 754 version of standard <code>float.h</code>
<code>fvec.h</code>	SSE intrinsics for Class Libraries
<code>ia64intrin.h</code>	
<code>ia64regs.h</code>	Standard header file
<code>iso646.h</code>	Standard header file
<code>ivec.h</code>	MMX(TM) instructions intrinsics for Class Libraries
<code>limits.h</code>	Standard header file
<code>mathimf.h</code>	Principal header file for current Intel Math Library
<code>mmintrin.h</code>	Intrinsics for MMX instructions
<code>omp.h</code>	Principal header file OpenMP*
<code>pgouser.h</code>	For use in the instrumentation compilation phase of profile-guided optimizations
<code>proto.h</code>	
<code>sse2mmx.h</code>	Principal header file for Streaming SIMD Extensions 2 intrinsics
<code>stdarg.h</code>	Replacement header for standard <code>stdarg.h</code>
<code>stdbool.h</code>	Defines <code>_Bool</code> keyword
<code>stddef.h</code>	Standard header file
<code>syslimits.h</code>	
<code>varargs.h</code>	Replacement header for standard <code>varargs.h</code>
<code>xarg.h</code>	Header file used by <code>stdarg.h</code> and <code>varargs.h</code>
<code>xmmintrin.h</code>	Principal header file for Streaming SIMD Extensions intrinsics

/lib Files

File	Description
libcprts.a	C++ standard language library
libcxa.so	C++ language library indicating I/O data location
libirc.a	Intel-specific library (optimizations)
libm.a	Math library
libguide.a	OpenMP library
libguide.so	Shared OpenMP library
libmofl.a	Multiple Object Format Library, used by the Intel assembler
libmofl.so	Shared Multiple Object Format Library, used by the Intel assembler
libunwinder.a	Unwinder library
libintrins.a	Intrinsic functions library

Diagnostics and Messages

This section describes the various messages that the compiler produces. These messages include the sign-on message and diagnostic messages for remarks, warnings, or errors. The compiler always displays any diagnostic message, along with the erroneous source line, on the standard output.

This section also describes how to control the severity of diagnostic messages.

Diagnostic Messages

Option	Description
-w0	Display errors (same as -w)
-w1	Display warnings and errors (DEFAULT)
-w2	Display remarks, warnings, and errors

Language Diagnostics

These messages describe diagnostics that are reported during the processing of the source file. These diagnostics have the following format:

```
filename (linenum): type [#nn]: message
```

filename	Indicates the name of the source file currently being processed.
linenum	Indicates the source line where the compiler detects the condition.
type	Indicates the severity of the diagnostic message: warning, remark, error, or catastrophic error.
[#nn]	The number assigned to the error (or warning) message. Hard errors or catastrophes are not assigned a number.
message	Describes the diagnostic.

The following is an example of a warning message:

```
tantst.cpp(3): warning #328: Local variable "increment" never used.
```

The compiler can also display internal error messages on the standard error. If your compilation produces any internal errors, contact your Intel representative. Internal error messages are in the following form:

```
FATAL COMPILER ERROR: message
```

Suppressing Warning Messages with lint Comments

The UNIX `lint` program attempts to detect features of a C or C++ program that are likely to be bugs, non-portable, or wasteful. The compiler recognizes three `lint`-specific comments:

1. `/*ARGSUSED*/`
2. `/*NOTREACHED*/`
3. `/*VARARGS*/`

Like the `lint` program, the compiler suppresses warnings about certain conditions when you place these comments at specific points in the source.

Suppressing Warning Messages or Enabling Remarks

Use the `-w` or `-wn` option to suppress warning messages or to enable remarks during the preprocessing and compilation phases. You can enter the option with one of the following arguments:

Option	Description
<code>-w0</code>	Display only errors (same as <code>-w</code>)
<code>-w1</code>	Display warnings and errors (DEFAULT)
<code>-w2</code>	Display remarks, warnings, and errors

For some compilations, you might not want warnings for known and benign characteristics, such as the K&R C constructs in your code. For example, the following command compiles `newprog.cpp` and displays compiler errors, but not warnings:

```
prompt>icpc -W0 newprog.cpp
```

Use the `-ww`, `-we`, or `-wd` option to indicate specific diagnostics.

Option	Description
<code>-wwL1[L2, ..., Ln]</code>	Changes the severity of diagnostics L1 through Ln to warning.
<code>-weL1[L2, ..., Ln]</code>	Changes the severity of diagnostics L1 through Ln to error.
<code>-wdL1[L2, ..., Ln]</code>	Disables diagnostics L1 through Ln.

Example

```
/* test.c */
int main()
{
    int x=0;
}
```

If you compile `test.c` using the `-Wall` option (enable all warnings), the compiler will emit warning #177:

```
prompt>icc -Wall test.c
```

```
remark #177: variable 'x' was declared but never referenced
```

To disable warning #177, use the `-wd` option:

```
prompt>icc -Wall -wd177 test.c
```

Likewise, using the `-we` option will result in a compile-time error:

```
prompt>icc -Wall -we177 test.c
```

```
error #177: variable 'x' was declared but never referenced
```

```
compilation aborted for test.c
```

Limiting the Number of Errors Reported

Use the `-wnn` option to limit the number of error messages displayed before the compiler aborts. By default, if more than 100 errors are displayed, compilation aborts.

Option	Description
<code>-wnn/i</code>	Limit the number of error diagnostics that will be displayed prior to aborting compilation to <i>n</i> . Remarks and warnings do not count towards this limit.

For example, the following command line specifies that if more than 50 error messages are displayed during the compilation of `a.cpp`, compilation aborts.

```
prompt>icpc -wn50 -c a.cpp
```

Remark Messages

These messages report common, but sometimes unconventional, use of C or C++. The compiler does not print or display remarks unless you specify level 4 for the `-w` option, as described in Suppressing Warning Messages or Enabling Remarks. Remarks do not stop translation or linking. Remarks do not interfere with any output files. The following are some representative remark messages:

- `function declared implicitly`
- `type qualifiers are meaningless in this declaration`
- `controlling expression is constant`

Intel Math Library

The Intel® C++ Compiler includes a mathematical software library containing highly optimized and very accurate mathematical functions. These functions are commonly used in scientific or graphic applications, as well as other programs that rely heavily on floating-point computations. Support for C99 `_Complex` data types is included by using the `-c99` compiler option. The `mathimf.h` header file includes prototypes for the library functions. See [Using the Intel Math Library](#). For a complete list of the functions available, refer to the [Function List](#) in this section.

Math Libraries for IA-32 and Itanium®-based Systems

The math library linked to an application depends on the compilation or linkage options specified.

Library	Description
<code>libimf.a</code>	Default static math library.
<code>libimf.so</code>	Default shared math library.

Using the Intel Math Library

To use the Intel math library, include the header file, `mathimf.h`, in your program. Here are two example programs that illustrate the use of the math library.

Example Using Real Functions

```

// real_math.c

#include <stdio.h>
#include <mathimf.h>

int main() {

float  fp32bits;
double fp64bits;
long double fp80bits;
long double pi_by_four = 3.141592653589793238/4.0;

// pi/4 radians is about 45 degrees.

fp32bits = (float) pi_by_four; // float approximation to pi/4
fp64bits = (double) pi_by_four; // double approximation to
pi/4
fp80bits = pi_by_four; // long double (extended)
approximation to pi/4

// The sin(pi/4) is known to be 1/sqrt(2) or approximately
.7071067

printf("When x = %8.8f, sinf(x) = %8.8f \n", fp32bits,
sinf(fp32bits));
printf("When x = %16.16f, sin(x) = %16.16f \n", fp64bits,
sin(fp64bits));
printf("When x = %20.20Lf, sinl(x) = %20.20f \n", fp80bits,
sinl(fp80bits));

return 0;
}

```

Compiling `real_math.c`:

```
prompt>icc real_math.c
```

The output of `a.out` will look like this:

```

When x = 0.78539816, sinf(x) = 0.70710678
When x = 0.7853981633974483, sin(x) = 0.7071067811865475
When x = 0.78539816339744827900, sinl(x) =
0.70710678118654750275

```

Example Using Complex Functions

```
// complex_math.c

#include <stdio.h>
#include <mathimf.h>

int main()
{

float  _Complex c32in,c32out;
double _Complex c64in,c64out;
double pi_by_four= 3.141592653589793238/4.0;

c64in = 1.0 + __I__ * pi_by_four;

// Create the double precision complex number 1 + (pi/4) * i
// where i is the imaginary unit.

c32in = (float _Complex) c64in;

// Create the float complex value from the double complex
value.

c64out = cexp(c64in);
c32out = cexpf(c32in);

// Call the complex exponential,
// cexp(z) = cexp(x+iy) = e^ (x + i y) = e^x * (cos(y) + i
sin(y))

printf("When z = %7.7f + %7.7f i, cexpf(z) = %7.7f + %7.7f i
\n"
,crealf(c32in),cimagf(c32in),crealf(c32out),cimagf(c32out));
printf("When z = %12.12f + %12.12f i, cexp(z) = %12.12f +
%12.12f i \n"
,creal(c64in),cimag(c64in),creal(c64out),cimagf(c64out));

return 0;
}
```

prompt>**icc complex_math.c**

The output of a.out will look like this:

```
When z = 1.0000000 + 0.7853982 i, cexpf(z) = 1.9221154 +
1.9221156 i
When z = 1.0000000000000 + 0.785398163397 i, cexp(z) =
1.922115514080 + 1.922115514080 i
```



Note

`_Complex` data types are supported in C but not in C++ programs.

Exception Conditions

If you call a math function using argument(s) that may produce undefined results, an error number is assigned to the system variable `errno`. Math function errors are usually domain errors or range errors.

Domain errors result from arguments that are outside the domain of the function. For example, `acos` is defined only for arguments between -1 and +1 inclusive. Attempting to evaluate `acos(-2)` or `acos(3)` results in a domain error, where the return value is `QNaN`.

Range errors occur when a mathematically valid argument results in a function value that exceeds the range of representable values for the floating-point data type. Attempting to evaluate `exp(1000)` results in a range error, where the return value is `INF`.

When domain or range error occurs, the following values are assigned to `errno`:

- domain error (EDOM): `errno = 33`
- range error (ERANGE): `errno = 34`

The following example shows how to read the `errno` value for an EDOM and ERANGE error.

```
// errno.c
#include <errno.h>
#include <mathimf.h>
#include <stdio.h>

int main(void)
{
    double neg_one=-1.0;
    double zero=0.0;

    // The natural log of a negative number is considered a
    domain error - EDOM
    printf("log(%e) = %e and errno(EDOM) = %d
    \n",neg_one,log(neg_one),errno);

    // The natural log of zero is considered a range error -
    ERANGE
    printf("log(%e) = %e and errno(ERANGE) = %d
    \n",zero,log(zero),errno);
}
```

The output of `errno.c` will look like this:

```
log(-1.000000e+00) = nan and errno(EDOM) = 33
log(0.000000e+00) = -inf and errno(ERANGE) = 34
```

For the math functions in this section, a corresponding value for `errno` is listed when applicable.

Other Considerations

Some math functions are inlined automatically by the compiler. The functions actually inlined may vary and may depend on any vectorization or processor-specific compilation options used. For more information, see Criteria for Inline Expansion of Functions.

A change of the default precision control or rounding mode may affect the results returned by some of the mathematical functions. See Floating-point Arithmetic Precision.

It's necessary to include the `-c99` compiler option when compiling programs that require support for `_Complex` data types.

Trigonometric Functions

The Intel Math library supports the following trigonometric functions:

ACOS

Description: The `acos` function returns the principal value of the inverse cosine of x in the range $[0, \pi]$ radians for x in the interval $[-1, 1]$.

errno: EDOM, for $|x| > 1$

Calling interface:

```
double acos(double x);
long double acosl(long double x);
float acosf(float x);
```

ACOSD

Description: The `acosd` function returns the principal value of the inverse cosine of x in the range $[0, 180]$ degrees for x in the interval $[-1, 1]$.

errno: EDOM, for $|x| > 1$

Calling interface:

```
double acosd(double x);
long double acosdl(long double x);
float acosdf(float x);
```

ASIN

Description: The `asin` function returns the principal value of the inverse sine of x in the range $[-\pi/2, +\pi/2]$ radians for x in the interval $[-1, 1]$.

errno: EDOM, for $|x| > 1$

Calling interface:

```
double asin(double x);
long double asinl(long double x);
float asinf(float x);
```

ASIND

Description: The `asind` function returns the principal value of the inverse sine of `x` in the range `[-90,90]` degrees for `x` in the interval `[-1,1]`.

errno: EDOM, for $|x| > 1$

Calling interface:

```
double asind(double x);
long double asindl(long double x);
float asindf(float x);
```

ATAN

Description: The `atan` function returns the principal value of the inverse tangent of `x` in the range `[-pi/2, +pi/2]` radians.

Calling interface:

```
double atan(double x);
long double atanl(long double x);
float atanf(float x);
```

ATAN2

Description: The `atan2` function returns the principal value of the inverse tangent of `y/x` in the range `[-pi, +pi]` radians.

errno: EDOM, for `x = 0` and `y = 0`

Calling interface:

```
double atan2(double y, double x);
long double atan2l(long double y, long double x);
float atan2f(float y, float x);
```

ATAND

Description: The `atand` function returns the principal value of the inverse tangent of `x` in the range `[-90,90]` degrees.

Calling interface:

```
double atand(double x);
long double atandl(long double x);
float atandf(float x);
```

ATAN2D

Description: The `atan2d` function returns the principal value of the inverse tangent of y/x in the range $[-180, +180]$ degrees.

errno: EDOM, for $x = 0$ and $y = 0$.

Calling interface:

```
double atan2d(double x, double y);
long double atan2dl(long double x, long double y);
float atan2df(float x, float y);
```

COS

Description: The `cos` function returns the cosine of x measured in radians. This function may be inlined with the Itanium® compiler.

Calling interface:

```
double cos(double x);
long double cosl(long double x);
float cosf(float x);
```

COSD

Description: The `cosd` function returns the cosine of x measured in degrees.

Calling interface:

```
double cosd(double x);
long double cosdl(long double x);
float cosdf(float x);
```

COT

Description: The `cot` function returns the cotangent of x measured in radians.

errno: ERANGE, for overflow conditions at $x = 0$.

Calling interface:

```
double cot(double x);
long double cotl(long double x);
float cotf(float x);
```

COTD

Description: The `cotd` function returns the cotangent of x measured in degrees.

errno: ERANGE, for overflow conditions at $x = 0$.

Calling interface:

```
double cotd(double x);
long double cotdl(long double x);
float cotdf(float x);
```

SIN

Description: The `sin` function returns the sine of `x` measured in radians. This function may be inlined with the Itanium® compiler.

Calling interface:

```
double sin(double x);
long double sinl(long double x);
float sinf(float x);
```

SINCOS

Description: The `sincos` function returns both the sine and cosine of `x` measured in radians. This function may be inlined with the Itanium® compiler.

Calling interface:

```
void sincos(double x, double *sinval, double *cosval);
void sincosl(long double x, long double *sinval, long double *cosval);
void sincosf(float x, float *sinval, float *cosval);
```

SINCOSD

Description: The `sincosd` function returns both the sine and cosine of `x` measured in degrees.

Calling interface:

```
void sincosd(double x, double *sinval, double *cosval);
void sincosdl(long double x, long double *sinval, long double *cosval);
void sincosdf(float x, float *sinval, float *cosval);
```

SIND

Description: The `sind` function computes the sine of `x` measured in degrees.

Calling interface:

```
double sind(double x);
long double sindl(long double x);
float sindf(float x);
```

TAN

Description: The `tan` function returns the tangent of `x` measured in radians.

Calling interface:

```
double tan(double x);
long double tanl(long double x);
float tanf(float x);
```

TAND

Description: The `tand` function returns the tangent of `x` measured in degrees.

errno: ERANGE, for overflow conditions

Calling interface:

```
double tand(double x);
long double tandl(long double x);
float tandf(float x);
```

Hyperbolic Functions

The Intel Math library supports the following hyperbolic functions:

ACOSH

Description: The `acosh` function returns the inverse hyperbolic cosine of `x`.

errno: EDOM, for $x < 1$

Calling interface:

```
double acosh(double x);
long double acoshl(long double x);
float acoshf(float x);
```

ASINH

Description: The `asinh` function returns the inverse hyperbolic sine of `x`.

Calling interface:

```
double asinh(double x);
long double asinhl(long double x);
float asinhf(float x);
```

ATANH

Description: The `atanh` function returns the inverse hyperbolic tangent of `x`.

errno: EDOM, for $x < 1$

errno: ERANGE, for $x = 1$

Calling interface:

```
double atanh(double x);
long double atanhl(long double x);
float atanhf(float x);
```

COSH

Description: The `cosh` function returns the hyperbolic cosine of x , $(e^x + e^{-x})/2$.

errno: ERANGE, for overflow conditions

Calling interface:

```
double cosh(double x);
long double coshl(long double x);
float coshf(float x);
```

SINH

Description: The `sinh` function returns the hyperbolic sine of x , $(e^x - e^{-x})/2$.

errno: ERANGE, for overflow conditions

Calling interface:

```
double sinh(double x);
long double sinhl(long double x);
float sinhlf(float x);
```

SINHCOSH

Description: The `sinhcosh` function returns both the hyperbolic sine and hyperbolic cosine of x .

errno: ERANGE, for overflow conditions

Calling interface:

```
void sinhcosh(double x, float *sinval, float *cosval);
void sinhcoshl(long double x, long double *sinval, long double *cosval);
void sinhcoshf(float x, float *sinval, float *cosval);
```

TANH

Description: The `tanh` function returns the hyperbolic tangent of x , $(e^x - e^{-x}) / (e^x + e^{-x})$.

Calling interface:

```
double tanh(double x);
long double tanhl(long double x);
float tanhf(float x);
```

Exponential Functions

The Intel Math library supports the following exponential functions:

CBRT

Description: The `cbrt` function returns the cube root of x .

Calling interface:

```
double cbrt(double x);
long double cbtrl(long double x);
float cbrtf(float x);
```

EXP

Description: The `exp` function returns e raised to the x power, e^x . This function may be inlined by the Itanium® compiler.

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double exp(double x);
long double expl(long double x);
float expf(float x);
```

EXP10

Description: The `exp10` function returns 10 raised to the x power, 10^x .

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double exp10(double x);
long double exp10l(long double x);
float exp10f(float x);
```

EXP2

Description: The `exp2` function returns 2 raised to the x power, 2^x .

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double exp2(double x);
long double exp2l(long double x);
float exp2f(float x);
```

EXPM1

Description: The `expm1` function returns e raised to the x power minus 1, $e^x - 1$.

errno: ERANGE, for overflow conditions

Calling interface:

```
double expm1(double x);
long double expm1l(long double x);
float expm1f(float x);
```

FREXP

Description: The `frexp` function converts a floating-point number x into signed normalized fraction in $[1/2, 1)$ multiplied by an integral power of two. The signed normalized fraction is returned, and the integer exponent stored at location `exp`.

Calling interface:

```
double frexp(double x, int *exp);
long double frexp(long double x, int *exp);
float frexpf(float x, int *exp);
```

HYPOT

Description: The `hypot` function returns the square root of $(x^2 + y^2)$.

errno: ERANGE, for overflow conditions

Calling interface:

```
double hypot(double x, double y);
long double hypotl(long double x, long double y);
float hypotf(float x, float y);
```

ILOGB

Description: The `ilogb` function returns the exponent of x base two as a signed `int` value.

errno: ERANGE, for $x = 0$

Calling interface:

```
int ilogb(double x);
int ilogbl(long double x);
int ilogbf(float x);
```

INVSQRT

Description: The `invsqrt` function returns the inverse square root.

Calling interface:

```
double invsqrt(double x);
long double invsqrtl(long double x);
float invsqrtf(float x);
```

LDEXP

Description: The `ldexp` function returns $x * 2^{\text{exp}}$, where `exp` is an integer value.

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double ldexp(double x, int exp);
long double ldexpl(long double x, int exp);
float ldexpf(float x, int exp);
```

LOG

Description: The `log` function returns the natural log of `x`, $\ln(x)$. This function may be inlined by the Itanium® compiler.

errno: EDOM, for $x < 0$

errno: ERANGE, for $x = 0$

Calling interface:

```
double log(double x);
long double logl(long double x);
float logf(float x);
```

LOG10

Description: The `log10` function returns the base-10 log of `x`, $\log_{10}(x)$. This function may be inlined by the Itanium® compiler.

errno: EDOM, for $x < 0$

errno: ERANGE, for $x = 0$

Calling interface:

```
double log10(double x);
long double log10l(long double x);
float log10f(float x);
```

LOG1P

Description: The `log1p` function returns the natural log of $(x+1)$, $\ln(x + 1)$.

errno: EDOM, for $x < -1$

errno: ERANGE, for $x = -1$

Calling interface:

```
double log1p(double x);
long double log1pl(long double x);
float log1pf(float x);
```

LOG2

Description: The `log2` function returns the base-2 log of `x`, $\log_2(x)$.

errno: EDOM, for $x < 0$

errno: ERANGE, for $x = 0$

Calling interface:

```
double log2(double x);
long double log2l(long double x);
float log2f(float x);
```

LOGB

Description: The `logb` function returns the signed exponent of `x`.

errno: EDOM, for $x = 0$

Calling interface:

```
double logb(double x);
long double logbl(long double x);
float logbf(float x);
```

POW

Description: The `pow` function returns `x` raised to the power of `y`, x^y .

Calling interface:

errno: EDOM, for $x = 0$ and $y < 0$

errno: EDOM, for $x < 0$ and `y` is a non-integer

errno: ERANGE, for overflow and underflow conditions

```
double pow(double x, double y);
long double powl(double x, double y);
float powf(float x, float y);
```

SCALB

Description: The `scalb` function returns $x * 2^y$, where `y` is a floating-point value.

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double scalb(double x, double y);
long double scalbl(long double x, long double y);
float scalbf(float x, float y);
```

SCALBN

Description: The `scalbn` function returns $x * 2^n$, where n is an integer value.

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double scalbn(double x, int n);
long double scalbnl (long double x, int n);
float scalbnf(float x, int n);
```

SCALBLN

Description: The `scalbln` function returns $x * 2^n$, where n is a long integer value.

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double scalbln(double x, long int n);
long double scalblnl (long double x, long int n);
float scalblnf(float x, long int n);
```

SQRT

Description: The `sqrt` function returns the correctly rounded square root.

errno: EDOM, for $x < 0$

Calling interface:

```
double sqrt(double x);
long double sqrtl(long double x);
float sqrtf(float x);
```

Special Functions

The Intel Math library supports the following special functions:

ANNUITY

Description: The `annuity` function computes the present value factor for an annuity, $(1 - (1+x)^{-y}) / x$, where x is a rate and y is a period.

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double annuity(double x, double y);
long double annuity(double x, double y);
float annuityf(float x, double y);
```

COMPOUND

Description: The `compound` function computes the compound interest factor, $(1+x)^y$, where `x` is a rate and `y` is a period.

errno: ERANGE, for underflow and overflow conditions

Calling interface:

```
double compound(double x, double y);
long double compound(double x, double y);
float compoundf(float x, double y);
```

ERF

Description: The `erf` function returns the error function value.

Calling interface:

```
double erf(double x);
long double erfl(long double x);
float erff(float x);
```

ERFC

Description: The `erfc` function returns the complementary error function value.

errno: ERANGE, for underflow conditions

Calling interface:

```
double erfc(double x);
long double erfcl(long double x);
float erfcf(float x);
```

GAMMA

Description: The `gamma` function returns the value of the logarithm of the absolute value of gamma.

errno: ERANGE, for overflow conditions when `x` is a negative integer.

Calling interface:

```
double gamma(double x);
long double gammal(long double x);
float gammaf(float x);
```

GAMMA_R

Description: The `gamma_r` function returns the value of the logarithm of the absolute value of gamma. The sign of the gamma function is returned in the integer `signgam`.

Calling interface:

```
double gamma_r(double x, int *signgam);
long double gammal_r(long double x, int *signgam);
float gammaf_r(float x, int *signgam);
```

J0

Description: Computes the Bessel function (of the first kind) of x with order 0.

Calling interface:

```
double j0(double x);
double j0l(long double x);
float j0f(float x);
```

J1

Description: Computes the Bessel function (of the first kind) of x with order 1.

Calling interface:

```
double j1(double x);
double j1l(long double x);
float j1f(float x);
```

JN

Description: Computes the Bessel function (of the first kind) of x with order n .

Calling interface:

```
double jn(int n, double x);
double jnl(int n, long double x);
float jnf(int n, float x);
```

LGAMMA

Description: The `lgamma` function returns the value of the logarithm of the absolute value of gamma.

errno: ERANGE, for overflow conditions, $x=0$ or negative integers.

Calling interface:

```
double lgamma(double x);
long double lgammal(long double x);
float lgammaf(float x);
```

LGAMMA_R

Description: The `lgamma_r` function returns the value of the logarithm of the absolute value of gamma. The sign of the gamma function is returned in the integer `signgam`.

errno: ERANGE, for overflow conditions, $x=0$ or negative integers.

Calling interface:

```
double lgamma_r(double x, int *signgam);
long double lgammal_r(double x, int *signgam);
float lgammaf_r(float x, int *signgam);
```

TGAMMA

Description: The `tgamma` function computes the gamma function of `x`.

errno: EDOM, for `x=0` or negative integers.

Calling interface:

```
double tgamma(double x);
long double tgamma1(long double x);
float tgammaf(float x);
```

Y0

Description: Computes the Bessel function (of the second kind) of `x` with order 0.

errno: EDOM, for `x <= 0`

Calling interface:

```
double y0(double x);
double y0l(long double x);
float y0f(float x);
```

Y1

Description: Computes the Bessel function (of the second kind) of `x` with order 1.

errno: EDOM, for `x <= 0`

Calling interface:

```
double y1(double x);
double y1l(long double x);
float y1f(float x);
```

YN

Description: Computes the Bessel function (of the second kind) of `x` with order `n`.

errno: EDOM, for `x <= 0`

Calling interface:

```
double yn(int n, double x);
double ynl(int n, long double x);
float ynf(int n, float x);
```

Nearest Integer Functions

The Intel Math library supports the following nearest integer functions:

CEIL

Description: The `ceil` function returns the smallest integral value not less than `x` as a floating-point number. This function may be inlined with the Itanium® compiler.

Calling interface:

```
double ceil(double x);
long double ceill(long double x);
float ceilf(float x);
```

FLOOR

Description: The `floor` function returns the largest integral value not greater than `x` as a floating-point value. This function may be inlined with the Itanium® compiler.

Calling interface:

```
double floor(double x);
long double floorl(long double x);
float floorf(float x);
```

LLRINT

Description: The `llrint` function returns the rounded integer value (according to the current rounding direction) as a `long long int`.

errno: ERANGE, for values too large

Calling interface:

```
long long int llrint(double x);
long long int llrintl(long double x);
long long int llrintf(float x);
```

LLROUND

Description: The `llround` function returns the rounded integer value as a `long long int`.

errno: ERANGE, for values too large

Calling interface:

```
long long int llround(double x);
long long int llroundl(long double x);
long long int llroundf(float x);
```

LRINT

Description: The `lrint` function returns the rounded integer value (according to the current rounding direction) as a `long int`.

Calling interface:

```
long int lrint(double x);
long int lrintl(long double x);
long int lrintf(float x);
```

LROUND

Description: The `lround` function returns the rounded integer value as a `long int`. Halfway cases are rounded away from zero.

errno: ERANGE, for values too large

Calling interface:

```
long int lround(double x);
long int lroundl(long double x);
long int lroundf(float x);
```

MODF

Description: The `modf` function returns the value of the signed fractional part of `x` and stores the integral part at `x *iptr` as a floating-point number.

Calling interface:

```
double modf(double x, double *iptr);
long double modfl(long double x, long double *iptr);
float modff(float x, float *iptr);
```

NEARBYINT

Description: The `nearbyint` function returns the rounded integer value as a floating-point number, using the current rounding direction.

Calling interface:

```
double nearbyint(double x);
long double nearbyintl(long double x);
float nearbyintf(float x);
```

RINT

Description: The `rint` function returns the rounded integer value as a floating-point number, using the current rounding direction.

Calling interface:

```
double rint(double x);
long double rintl(long double x);
float rintf(float x);
```

ROUND

Description: The `round` function returns the nearest integral value as a floating-point number. Halfway cases are rounded away from zero.

Calling interface:

```
double round(double x);
long double roundl(long double x);
float roundf(float x);
```

TRUNC

Description: The `trunc` function returns the truncated integral value as a floating-point number.

Calling interface:

```
double trunc(double x);
long double trunc1(long double x);
float truncf(float x);
```

Remainder Functions

The Intel Math library supports the following remainder functions:

FMOD

Description: The `fmod` function returns the value $x - n * y$ for integer n such that if y is nonzero, the result has the same sign as x and magnitude less than the magnitude of y .

errno: EDOM, for $x = 0$

Calling interface:

```
double fmod(double x, double y);
long double fmod1(long double x, long double y);
float fmodf(float x, float y);
```

REMAINDER

Description: The `remainder` function returns the value of $x \text{ REM } y$ as required by the IEEE standard.

Calling interface:

```
double remainder(double x, double y);
long double remainder1(long double x, long double y);
float remainderf(float x, float y);
```

REMQUO

Description: The `remquo` function returns the value of $x \text{ REM } y$. In the object pointed to by `quo` the function stores a value whose sign is the sign of x/y and whose magnitude is congruent modulo 2^{24} of the integral quotient of x/y , where n is an implementation-defined integer greater than or equal to 3.

Calling interface:

```
double remquo(double x, double y, int *quo);
long double remquol(long double x, long double y, int *quo);
float remquof(float x, float y, int *quo);
```

Miscellaneous Functions

The Intel Math library supports the following miscellaneous functions:

COPYSIGN

Description: The `copysign` function returns the value with the magnitude of x and the sign of y .

Calling interface:

```
double copysign(double x, double y);
long double copysignl(long double x, long double y);
float copysignf(float x, float y);
```

FABS

Description: The `fabs` function returns the absolute value of x .

Calling interface:

```
double fabs(double x);
long double fabsl(long double x);
float fabsf(float x);
```

FDIM

Description: The `fdim` function returns the positive difference value, $x-y$ (for $x > y$) or $+0$ (for $x \leq y$).

errno: ERANGE, for values too large

Calling interface:

```
double fdim(double x, double y);
long double fdiml(long double x, long double y);
float fdimf(float x, float y);
```

FINITE

Description: The `finite` function returns 1 if x is not a NaN or \pm infinity. Otherwise 0 is returned..

Calling interface:

```
int finite(double x);
int finitel(long double x);
int finitef(float x);
```

FMA

Description: The `fma` functions return $(x*y)+z$.

Calling interface:

```
double fma(double x, double y, long double z);
long double fmal(long double x, long double y, long double z);
float fmaf(float x, float y, long double z);
```

FMAX

Description: The `fmax` function returns the maximum numeric value of its arguments.

Calling interface:

```
double fmax(double x, double y);
long double fmaxl(long double x, long double y);
float fmaxf(float x, float y);
```

FMIN

Description: The `fmin` function returns the minimum numeric value of its arguments.

Calling interface:

```
double fmin(double x, double y);
long double fminl(long double x, long double y);
float fminf(float x, float y);
```

FPCLASSIFY

Description: The `fpclassify` function returns the value of the number classification macro appropriate to the value of its argument.

Calling interface:

```
double fpclassify(double x);
long double fpclassifyl(long double x);
float fpclassifyf(float x);
```

ISFINITE

Description: The `isfinite` function returns 1 if `x` is not a NaN or +/- infinity. Otherwise 0 is returned..

Calling interface:

```
int isfinite(double x);
int isfinitel(long double x);
int isfinitef(float x);
```

ISGREATER

Description: The `isgreater` function returns 1 if `x` is greater than `y`. This function does not raise the invalid floating-point exception.

Calling interface:

```
int isgreater(double x, double y);
int isgreaterl(long double x, long double y);
int isgreaterf(float x, float y);
```

ISGREATEREQUAL

Description: The `isgreaterequal` function returns 1 if `x` is greater than or equal to `y`. This function does not raise the invalid floating-point exception.

Calling interface:

```
int isgreaterequal(double x, double y);
int isgreaterequal1(long double x, long double y);
int isgreaterequalf(float x, float y);
```

ISINF

Description: The `isinf` function returns a non-zero value if and only if its argument has an infinite value.

Calling interface:

```
int isinf(double x);
int isinfl(long double x);
int isinff(float x);
```

ISLESS

Description: The `isless` function returns 1 if `x` is less than `y`. This function does not raise the invalid floating-point exception.

Calling interface:

```
int isless(double x, double y);
int islessl(long double x, long double y);
int islessf(float x, float y);
```

ISLESSEQUAL

Description: The `islessequal` function returns 1 if `x` is less than or equal to `y`. This function does not raise the invalid floating-point exception.

Calling interface:

```
int islessequal(double x, double y);
int islessequal1(long double x, long double y);
int islessequalf(float x, float y);
```

ISLESSGREATER

Description: The `islessgreater` function returns 1 if `x` is less than or greater than `y`. This function does not raise the invalid floating-point exception.

Calling interface:

```
int islessgreater(double x, double y);
int islessgreaterl(long double x, long double y);
int islessgreaterf(float x, float y);
```

ISNAN

Description: The `isnan` function returns a non-zero value if and only if `x` has a NaN value.

Calling interface:

```
int isnan(double x);
int isnanl(long double x);
int isnanf(float x);
```

ISNORMAL

Description: The `isnormal` function returns a non-zero value if and only if `x` is normal.

Calling interface:

```
int isnormal(double x);
int isnormal1(long double x);
int isnormalf(float x);
```

ISUNORDERED

Description: The `isunordered` function returns 1 if either `x` or `y` is a NaN. This function does not raise the invalid floating-point exception.

Calling interface:

```
int isunordered(double x, double y);
int isunorderedl(long double x, long double y);
int isunorderedf(float x, float y);
```

NEXTAFTER

Description: The `nextafter` function returns the next representable value in the specified format after `x` in the direction of `y`.

errno: ERANGE, for values too large

Calling interface:

```
double nextafter(double x, double y);
long double nextafterl(long double x, long double y);
float nextafterf(float x, float y);
```

NEXTTOWARD

Description: The `nexttoward` function returns the next representable value in the specified format after `x` in the direction of `y`. If `x` equals `y`, then the function returns `y` converted to the type of the function.

errno: ERANGE, for values too large

Calling interface:

```
double nexttoward(double x, double y);
long double nexttowardl(long double x, long double y);
float nexttowardf(float x, float y);
```

SIGNBIT

Description: The `signbit` function returns a non-zero value if and only if the sign of `x` is negative.

Calling interface:

```
int signbit(double x);
int signbitl(long double x);
int signbitf(float x);
```

SIGNIFICAND

Description: The `significand` function returns the significand of `x` in the interval [1,2). For `x` equal to zero, NaN, or +/- infinity, the original `x` is returned.

Calling interface:

```
double significand(double x);
long double significandl(long double x);
float significandf(float x);
```

Complex Functions

The Intel Math library supports the following complex functions:

CABS

Description: The `cabs` function returns the complex absolute value of `z`.

Calling interface:

```
double _cabs(double _Complex z);
long double _cabsl(long double _Complex z);
float _cabsf(float _Complex z);
```

CACOS

Description: The `cacos` function returns the complex inverse cosine of `z`.

Calling interface:

```
double _Complex cacos(double _Complex z);
long double _Complex cacosl(long double _Complex z);
float _Complex cacosf(float _Complex z);
```

CACOSH

Description: The `cacosh` function returns the complex inverse hyperbolic cosine of z .

Calling interface:

```
double _Complex cacosh(double _Complex z);
long double _Complex cacoshl(long double _Complex z);
float _Complex cacoshf(float _Complex z);
```

CARG

Description: The `carg` function returns the value of the argument in the interval $[-\pi, +\pi]$.

Calling interface:

```
double carg(double _Complex z);
long double cargl(long double _Complex z);
float cargf(float _Complex z);
```

CASIN

Description: The `casin` function returns the complex inverse sine of z .

Calling interface:

```
double _Complex casin(double _Complex z);
long double _Complex casinl(long double _Complex z);
float _Complex casinf(float _Complex z);
```

CASINH

Description: The `casinh` function returns the complex inverse hyperbolic sine of z .

Calling interface:

```
double _Complex casinh(double _Complex z);
long double _Complex casinhl(long double _Complex z);
float _Complex casinhf(float _Complex z);
```

CATAN

Description: The `catan` function returns the complex inverse tangent of z .

Calling interface:

```
double _Complex catan(double _Complex z);
long double _Complex catanl(long double _Complex z);
float _Complex catanf(float _Complex z);
```

CATANH

Description: The `catanh` function returns the complex inverse hyperbolic tangent of z .

Calling interface:

```
double _Complex catanh(double _Complex z);
long double _Complex catanh1(long double _Complex z);
float _Complex catanhf(float _Complex z);
```

CCOS

Description: The `ccos` function returns the complex cosine of z .

Calling interface:

```
double _Complex ccos(double _Complex z);  
long double _Complex ccosl(long double _Complex z);  
float _Complex ccosf(float _Complex z);
```

CCOSH

Description: The `ccosh` function returns the complex hyperbolic cosine of z .

Calling interface:

```
double _Complex ccosh(double _Complex z);  
long double _Complex ccoshl(long double _Complex z);  
float _Complex ccoshf(float _Complex z);
```

CEXP

Description: The `cexp` function computes e^z .

Calling interface:

```
double _Complex cexp(double _Complex z);  
long double _Complex cexpl(long double _Complex z);  
float _Complex cexpf(float _Complex z);
```

CEXP2

Description: The `cexp2` function computes 2^z .

Calling interface:

```
double _Complex cexp2(double _Complex z);  
long double _Complex cexp2l(long double _Complex z);  
float _Complex cexp2f(float _Complex z);
```

CEXP10

Description: The `cexp10` function computes 10^z .

Calling interface:

```
double _Complex cexp10(double _Complex z);  
long double _Complex cexp10l(long double _Complex z);  
float _Complex cexp10f(float _Complex z);
```

CIMAG

Description: The `cimag` function returns the imaginary part value of z .

Calling interface:

```
double cimag(double _Complex z);  
long double cimagl(long double _Complex z);  
float cimagf(float _Complex z);
```

CIS

Description: The `cis` function returns the cosine and sine (as a complex value) of `z` measured in radians.

Calling interface:

```
double _Complex cis(double z);
long double _Complex cisl(long double z);
float _Complex cisf(float z);
```

CISD

Description: The `cis` function returns the cosine and sine (as a complex value) of `z` measured in degrees.

Calling interface:

```
double _Complex cis(double z);
long double _Complex cisl(long double z);
float _Complex cisf(float z);
```

CLOG

Description: The `clog` function returns the complex natural logarithm of `z`.

Calling interface:

```
double _Complex clog(double _Complex z);
long double _Complex clogl(long double _Complex z);
float _Complex clogf(float _Complex z);
```

CLOG2

Description: The `clog2` function returns the complex logarithm base 2 of `z`.

Calling interface:

```
double _Complex clog2(double _Complex z);
long double _Complex clog2l(long double _Complex z);
float _Complex clog2f(float _Complex z);
```

CLOG10

Description: The `clog10` function returns the complex logarithm base 10 of `z`.

Calling interface:

```
double _Complex clog10(double _Complex z);
long double _Complex clog10l(long double _Complex z);
float _Complex clog10f(float _Complex z);
```

CONJ

Description: The `conj` function returns the complex conjugate of `z`, by reversing the sign of its imaginary part.

Calling interface:

```
double _Complex conj(double _Complex z);  
long double _Complex conjl(long double _Complex z);  
float _Complex conjf(float _Complex z);
```

CPOW

Description: The `cpow` function returns the complex power function, x^y .

Calling interface:

```
double _Complex cpow(double _Complex x, double _Complex y);  
long double _Complex cpowl(long double _Complex x, double _Complex y);  
float _Complex cpowf(float _Complex x, float _Complex y);
```

CPROJ

Description: The `cproj` function returns a projection of `z` onto the Riemann sphere.

Calling interface:

```
double _Complex cproj(double _Complex z);  
long double _Complex cprojl(long double _Complex z);  
float _Complex cprojf(float _Complex z);
```

CREAL

Description: The `creal` function returns the real part value of `z`.

Calling interface:

```
double creal(double _Complex z);  
long double creall(long double _Complex z);  
float crealf(float _Complex z);
```

CSIN

Description: The `csin` function returns the complex sine of `z`.

Calling interface:

```
double _Complex csin(double _Complex z);  
long double _Complex csinl(long double _Complex z);  
float _Complex csinf(float _Complex z);
```

CSINH

Description: The `csinh` function returns the complex hyperbolic sine of z .

Calling interface:

```
double _Complex csinh(double _Complex z);
long double _Complex csinh1(long double _Complex z);
float _Complex csinhf(float _Complex z);
```

CSQRT

Description: The `csqrt` function returns the complex square root of z .

Calling interface:

```
double _Complex csqrt(double _Complex z);
long double _Complex csqrt1(long double _Complex z);
float _Complex csqrtf(float _Complex z);
```

CTAN

Description: The `ctan` function returns the complex tangent of z .

Calling interface:

```
double _Complex ctan(double _Complex z);
long double _Complex ctan1(long double _Complex z);
float _Complex ctanf(float _Complex z);
```

CTANH

Description: The `ctanh` function returns the complex hyperbolic tangent of z .

Calling interface:

```
double _Complex ctanh(double _Complex z);
long double _Complex ctanh1(long double _Complex z);
float _Complex ctanhf(float _Complex z);
```

C99 Macros

The Intel Math library and `mathimf.h` header file support the following C99 macros:

```
int fpclassify(x);
```

```
int isfinite(x);
```

```
int isgreater(x, y);
```

```
int isgreaterequal(x, y);
```

```
int isinf(x);
```

```
int isless(x, y);
```

```
int islessequal(x, y);
```

```
int islessgreater(x, y);
```

```
int isnan(x);
```

```
int isnormal(x);
```

```
int isunordered(x, y);
```

```
int signbit(x);
```

See also, Miscellaneous Functions.

Intel® C++ Intrinsic Reference

The Intel® Pentium® 4 processor and other Intel processors have instructions to enable development of optimized multimedia applications. The instructions are implemented through extensions to previously implemented instructions. This technology uses the single instruction, multiple data (SIMD) technique. By processing data elements in parallel, applications with media-rich bit streams are able to significantly improve performance using SIMD instructions. The Intel® Itanium® processor also supports these instructions.

The most direct way to use these instructions is to inline the assembly language instructions into your source code. However, this can be time-consuming and tedious, and assembly language inline programming is not supported on all compilers. Instead, Intel provides easy implementation through the use of API extension sets referred to as intrinsics.

Intrinsics are special coding extensions that allow using the syntax of C function calls and C variables instead of hardware registers. Using these intrinsics frees programmers from having to program in assembly language and manage registers. In addition, the compiler optimizes the instruction scheduling so that executables run faster.

In addition, the native intrinsics for the Itanium processor give programmers access to Itanium instructions that cannot be generated using the standard constructs of the C and C++ languages. The Intel® C++ Compiler also supports general purpose intrinsics that work across all IA-32 and Itanium-based platforms.

For more information on intrinsics, please refer to the following publications:

Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference Manual, Intel Corporation, doc. number 243191

Intrinsics Availability on Intel Processors

Processors:	MMX(TM) Technology Ininsics	Streaming SIMD Extensions	Streaming SIMD Extensions 2	Itanium Processor Instructions
Itanium Processor	X	X	N/A	X
Pentium 4 Processor	X	X	X	N/A
Pentium III Processor	X	X	N/A	N/A
Pentium II Processor	X	N/A	N/A	N/A
Pentium with MMX Technology	X	N/A	N/A	N/A
Pentium Pro Processor	N/A	N/A	N/A	N/A
Pentium Processor	N/A	N/A	N/A	N/A

Benefits of Using Intrinsics

The major benefit of using intrinsics is that you now have access to key features that are not available using conventional coding practices. Intrinsics enable you to code with the syntax of C function calls and variables instead of assembly language. Most MMX(TM) technology, Streaming SIMD Extensions, and Streaming SIMD Extensions 2 intrinsics have a corresponding C intrinsic that implements that instruction directly. This frees you from managing registers and enables the compiler to optimize the instruction scheduling.

The MMX technology and Streaming SIMD Extension instructions use the following new features:

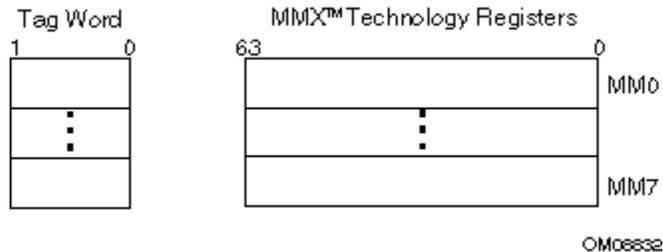
- new Registers--Enable packed data of up to 128 bits in length for optimal SIMD processing
- new Data Types--Enable packing of up to 16 elements of data in one register

The Streaming SIMD Extensions 2 intrinsics are defined only for IA-32, not for Itanium®-based systems. Streaming SIMD Extensions 2 operate on 128 bit quantities - 2 64-bit double precision floating point values. The Itanium architecture does not support parallel double precision computation, so Streaming SIMD Extensions 2 are not implemented on Itanium-based systems.

New Registers

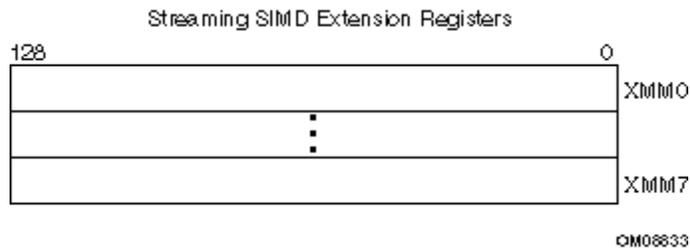
A key feature provided by the architecture of the processors are new register sets. The MMX instructions use eight 64-bit registers (mm0 to mm7) which are aliased on the floating-point stack registers.

MMX(TM) Technology Registers



Streaming SIMD Extensions Registers

The Streaming SIMD Extensions use eight 128-bit registers (xmm0 to xmm7).



These new data registers enable the processing of data elements in parallel. Because each register can hold more than one data element, the processor can process more than one data element simultaneously. This processing capability is also known as single-instruction multiple data processing (SIMD).

For each computational and data manipulation instruction in the new extension sets, there is a corresponding C intrinsic that implements that instruction directly. This frees you from managing registers and assembly programming. Further, the compiler optimizes the instruction scheduling so that your executable runs faster.



Note

The MM and XMM registers are the SIMD registers used by the IA-32 platforms to implement MMX technology and Streaming SIMD Extensions/Streaming SIMD Extensions 2 intrinsics. On the Itanium-based platforms, the MMX and Streaming SIMD Extension intrinsics use the 64-bit general registers and the 64-bit significant of the 80-bit floating-point register.

Data Types

Intrinsic functions use four new C data types as operands, representing the new registers that are used as the operands to these intrinsic functions. The following table shows the data type availability marked with "X".

New Data Types Available

New Data Type	MMX(TM) Technology	Streaming SIMD Extensions	Streaming SIMD Extensions 2	Itanium® Processor
<code>__m64</code>	X	X	X	X
<code>__m128</code>	N/A	X	X	X
<code>__m128d</code>	N/A	N/A	X	X
<code>__m128i</code>	N/A	N/A	X	X

`__m64` Data Type

The `__m64` data type is used to represent the contents of an MMX register, which is the register that is used by the MMX technology intrinsics. The `__m64` data type can hold eight 8-bit values, four 16-bit values, two 32-bit values, or one 64-bit value.

`__m128` Data Types

The `__m128` data type is used to represent the contents of a Streaming SIMD Extension register used by the Streaming SIMD Extension intrinsics. The `__m128` data type can hold four 32-bit floating values.

The `__m128d` data type can hold two 64-bit floating-point values.

The `__m128i` data type can hold sixteen 8-bit, eight 16-bit, four 32-bit, or two 64-bit integer values.

The compiler aligns `__m128` local and global data to 16-byte boundaries on the stack. To align `integer`, `float`, or `double` arrays, you can use the `declspec` statement.

New Data Types Usage Guidelines

Since these new data types are not basic ANSI C data types, you must observe the following usage restrictions:

- Use new data types only on either side of an assignment, as a return value, or as a parameter. You cannot use it with other arithmetic expressions (+, -, etc).
- Use new data types as objects in aggregates, such as unions to access the byte elements and structures.
- Use new data types only with the respective intrinsics described in this documentation. The new data types are supported on both sides of an assignment statement: as parameters to a function call, and as a return value from a function call.

Naming and Usage Syntax

Most of the intrinsic names use a notational convention as follows:

```
__mm_<intrin_op>_<suffix>
```

<code><intrin_op></code>	Indicates the intrinsics basic operation; for example, <code>add</code> for addition and <code>sub</code> for subtraction.
<code><suffix></code>	Denotes the type of data operated on by the instruction. The first one or two letters of each suffix denotes whether the data is packed (<code>p</code>), extended packed (<code>ep</code>), or scalar (<code>s</code>). The remaining letters denote the type: <ul style="list-style-type: none"> • <code>s</code> single-precision floating point • <code>d</code> double-precision floating point • <code>i128</code> signed 128-bit integer • <code>i64</code> signed 64-bit integer • <code>u64</code> unsigned 64-bit integer • <code>i32</code> signed 32-bit integer • <code>u32</code> unsigned 32-bit integer • <code>i16</code> signed 16-bit integer • <code>u16</code> unsigned 16-bit integer • <code>i8</code> signed 8-bit integer • <code>u8</code> unsigned 8-bit integer

A number appended to a variable name indicates the element of a packed object. For example, `r0` is the lowest word of `r`. Some intrinsics are "composites" because they require more than one instruction to implement them.

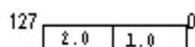
The packed values are represented in right-to-left order, with the lowest value being used for scalar operations. Consider the following example operation:

```
double a[2] = {1.0, 2.0};
__m128d t = __mm_load_pd(a);
```

The result is the same as either of the following:

```
__m128d t = __mm_set_pd(2.0, 1.0);
__m128d t = __mm_setr_pd(1.0, 2.0);
```

In other words, the `xmm` register that holds the value `t` will look as follows:



The "scalar" element is 1.0. Due to the nature of the instruction, some intrinsics require their arguments to be immediates (constant integer literals).

Intrinsic Syntax

To use an intrinsic in your code, insert a line with the following syntax:

```
data_type intrinsic_name (parameters)
```

Where,

<code>data_type</code>	Is the return data type, which can be either <code>void</code> , <code>int</code> , <code>__m64</code> , <code>__m128</code> , <code>__m128d</code> , <code>__m128i</code> , <code>__int64</code> . Intrinsic that can be implemented across all IA may return other data types as well, as indicated in the intrinsic syntax definitions.
<code>intrinsic_name</code>	Is the name of the intrinsic, which behaves like a function that you can use in your C++ code instead of inlining the actual instruction.
<code>parameters</code>	Represents the parameters required by each intrinsic.

Intrinsics For All IA

The intrinsics in this section function across all IA-32 and Itanium®-based platforms. They are offered as a convenience to the programmer. They are grouped as follows:

- Integer Arithmetic Related
- Floating-Point Related
- String and Block Copy Related
- Miscellaneous

Integer Arithmetic Related

Intrinsic	Description
<code>int abs(int)</code>	Returns the absolute value of an integer.
<code>long labs(long)</code>	Returns the absolute value of a long integer.
<code>unsigned long _lrotl(unsigned long value, int shift)</code>	Rotates bits left for an unsigned long integer.
<code>unsigned long _lrotr(unsigned long value, int shift)</code>	Rotates bits right for an unsigned long integer.
<code>unsigned int __rotl(unsigned int value, int shift)</code>	Rotates bits left for an unsigned integer.
<code>unsigned int __rotr(unsigned int value, int shift)</code>	Rotates bits right for an unsigned integer.

**Note**

Passing a constant shift value in the rotate intrinsics results in higher performance.

Floating-point Related

Intrinsic	Description
<code>double fabs(double)</code>	Returns the absolute value of a floating-point value.
<code>double log(double)</code>	Returns the natural logarithm $\ln(x)$, $x > 0$, with double precision.
<code>float logf(float)</code>	Returns the natural logarithm $\ln(x)$, $x > 0$, with single precision.
<code>double log10(double)</code>	Returns the base 10 logarithm $\log_{10}(x)$, $x > 0$, with double precision.
<code>float log10f(float)</code>	Returns the base 10 logarithm $\log_{10}(x)$, $x > 0$, with single precision.
<code>double exp(double)</code>	Returns the exponential function with double precision.
<code>float expf(float)</code>	Returns the exponential function with single precision.
<code>double pow(double, double)</code>	Returns the value of x to the power y with double precision.
<code>float powf(float, float)</code>	Returns the value of x to the power y with single precision.
<code>double sin(double)</code>	Returns the sine of x with double precision.
<code>float sinf(float)</code>	Returns the sine of x with single precision.
<code>double cos(double)</code>	Returns the cosine of x with double precision.
<code>float cosf(float)</code>	Returns the cosine of x with single precision.
<code>double tan(double)</code>	Returns the tangent of x with double precision.
<code>float tanf(float)</code>	Returns the tangent of x with single precision.
<code>double acos(double)</code>	Returns the arccosine of x with double precision.
<code>float acosf(float)</code>	Returns the arccosine of x with single precision.
<code>double acosh(double)</code>	Compute the inverse hyperbolic cosine of the argument with double precision.
<code>float acoshf(float)</code>	Compute the inverse hyperbolic cosine of the argument with single precision.
<code>double asin(double)</code>	Compute arc sine of the argument with double precision.

Intrinsic	Description
float asinf(float)	Compute arc sine of the argument with single precision.
double asinh(double)	Compute inverse hyperbolic sine of the argument with double precision.
float asinhf(float)	Compute inverse hyperbolic sine of the argument with single precision.
double atan(double)	Compute arc tangent of the argument with double precision.
float atanf(float)	Compute arc tangent of the argument with single precision.
double atanh(double)	Compute inverse hyperbolic tangent of the argument with double precision.
float atanhf(float)	Compute inverse hyperbolic tangent of the argument with single precision.
float cabs(double)**	Computes absolute value of complex number.
double ceil(double)	Computes smallest integral value of double precision argument not less than the argument.
float ceilf(float)	Computes smallest integral value of single precision argument not less than the argument.
double cosh(double)	Computes the hyperbolic cosine of double precision argument.
float coshf(float)	Computes the hyperbolic cosine of single precision argument.
float fabsf(float)	Computes absolute value of single precision argument.
double floor(double)	Computes the largest integral value of the double precision argument not greater than the argument.
float floorf(float)	Computes the largest integral value of the single precision argument not greater than the argument.
double fmod(double)	Computes the floating-point remainder of the division of the first argument by the second argument with double precision.
float fmodf(float)	Computes the floating-point remainder of the division of the first argument by the second argument with single precision.
double hypot(double, double)	Computes the length of the hypotenuse of a right angled triangle with double precision.
float hypotf(float)	Computes the length of the hypotenuse of a right angled triangle with single precision.

Intrinsic	Description
<code>double rint(double)</code>	Computes the integral value represented as double using the IEEE rounding mode.
<code>float rintf(float)</code>	Computes the integral value represented with single precision using the IEEE rounding mode.
<code>double sinh(double)</code>	Computes the hyperbolic sine of the double precision argument.
<code>float sinhf(float)</code>	Computes the hyperbolic sine of the single precision argument.
<code>float sqrtf(float)</code>	Computes the square root of the single precision argument.
<code>double tanh(double)</code>	Computes the hyperbolic tangent of the double precision argument.
<code>float tanhf(float)</code>	Computes the hyperbolic tangent of the single precision argument.

* Not implemented on Itanium®-based systems.

** `double` in this case is a complex number made up of two single precision (32-bit floating point) elements (real and imaginary parts).

String and Block Copy Related

The following are not implemented as intrinsics on Itanium®-based platforms.

Intrinsic	Description
<code>char *_strset(char *, _int32)</code>	Sets all characters in a string to a fixed value.
<code>void *memcmp(const void *cs, const void *ct, size_t n)</code>	Compares two regions of memory. Return <0 if <code>cs<ct</code> , 0 if <code>cs=ct</code> , or >0 if <code>cs>ct</code> .
<code>void *memcpy(void *s, const void *ct, size_t n)</code>	Copies from memory. Returns <code>s</code> .
<code>void *memset(void *s, int c, size_t n)</code>	Sets memory to a fixed value. Returns <code>s</code> .

Intrinsic	Description
<code>char *strcat(char * s, const char * ct)</code>	Appends to a string. Returns <i>s</i> .
<code>int *strcmp(const char *, const char *)</code>	Compares two strings. Return <0 if <i>cs</i> < <i>ct</i> , 0 if <i>cs</i> = <i>ct</i> , or >0 if <i>cs</i> > <i>ct</i> .
<code>char *strcpy(char * s, const char * ct)</code>	Copies a string. Returns <i>s</i> .
<code>size_t strlen(const char * cs)</code>	Returns the length of string <i>cs</i> .
<code>int strncmp(char *, char *, int)</code>	Compare two strings, but only specified number of characters.
<code>int strncpy(char *, char *, int)</code>	Copies a string, but only specified number of characters.

Miscellaneous Intrinsics

The intrinsic functions listed here are common to IA-32 and the Itanium® architecture.

Intrinsic	Description
<code>_abnormal_termination(void)</code>	Can be invoked only by termination handlers. Returns TRUE if the termination handler is invoked as a result of a premature exit of the corresponding try-finally region.
<code>void *_alloca(int)</code>	Allocates the buffers.
<code>extern int _bit_scan_forward(int x)</code>	Returns the bit index of the least significant set bit of <i>x</i> . If <i>x</i> is 0, the result is undefined.

Intrinsic	Description
extern int <code>_bit_scan_reverse(int)</code>	Returns the bit index of the most significant set bit of x. If x is 0, the result is undefined.
extern int <code>_bswap(int)</code>	Reverses the byte order of x. Bits 0-7 are swapped with bits 24-31, and bits 8-15 are swapped with bits 16-23.
<code>_exception_code(void)</code>	Returns the exception code.
<code>_exception_info(void)</code>	Returns the exception information.
<code>void _enable()</code>	Enables the interrupt.
<code>void _disable()</code>	Disables the interrupt.
int <code>_in_byte(int)</code>	Intrinsic that maps to the IA-32 instruction IN. Transfer data byte from port specified by argument.
int <code>_in_dword(int)</code>	Intrinsic that maps to the IA-32 instruction IN. Transfer double word from port specified by argument.
int <code>_in_word(int)</code>	Intrinsic that maps to the IA-32 instruction IN. Transfer word from port specified by argument.
int <code>_inp(int)</code>	Same as <code>_in_byte</code>
int <code>_inpd(int)</code>	Same as <code>_in_dword</code>
int <code>_inpw(int)</code>	Same as <code>_in_word</code>
int <code>_out_byte(int, int)</code>	Intrinsic that maps to the IA-32 instruction OUT. Transfer data byte in second argument to port specified by first argument.
int <code>_out_dword(int, int)</code>	Intrinsic that maps to the IA-32 instruction OUT. Transfer double word in second argument to port specified by first argument.
int <code>_out_word(int, int)</code>	Intrinsic that maps to the IA-32 instruction OUT. Transfer word in second argument to port specified by first argument.
int <code>_outp(int, int)</code>	Same as <code>_out_byte</code>
int <code>_outpd(int, int)</code>	Same as <code>_out_dword</code>
int <code>_outpw(int, int)</code>	Same as <code>_out_word</code>
extern int <code>_popcnt32(int x)</code>	Returns the number of set bits in x.

Intrinsic	Description
<code>extern __int64 _rdtsc(void)</code>	Returns the current value of the processor's 64-bit time stamp counter.
<code>extern __int64 _rdpmc(int p)</code>	Returns the current value of the 40-bit performance monitoring counter specified by <code>p</code> .
<code>int _setjmp(jmp_buf)*</code>	A fast version of <code>setjmp()</code> , which bypasses the termination handling. Saves the callee-save registers, stack pointer and return address.

MMX(TM) Technology Intrinsics

Support for MMX(TM) Technology

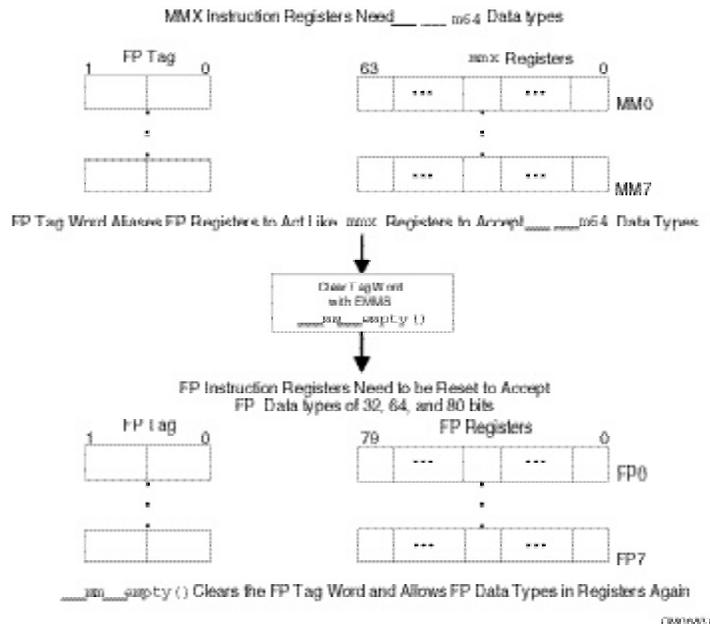
MMX(TM) technology is an extension to the Intel architecture (IA) instruction set. The MMX instruction set adds 57 opcodes and a 64-bit quadword data type, and eight 64-bit registers. Each of the eight registers can be directly addressed using the register names `mm0` to `mm7`.

The prototypes for MMX technology intrinsics are in the `mmintrin.h` header file.

The EMMS Instruction: Why You Need It

Using `EMMS` is like emptying a container to accommodate new content. For instance, MMX(TM) instructions automatically enable an `FP` tag word in the register to enable use of the `__m64` data type. This resets the `FP` register set to alias it as the MMX register set. To enable the `FP` register set again, reset the register state with the `EMMS` instruction or via the `_mm_empty()` intrinsic.

Why You Need EMMS to Reset After an MMX(TM) Instruction



Caution

Failure to empty the multimedia state after using an MMX instruction and before using a floating-point instruction can result in unexpected execution or poor performance.

EMMS Usage Guidelines

The guidelines when to use EMMS are:

- Do not use on Itanium®-based systems. There are no special registers (or overlay) for the MMX(TM) instructions or Streaming SIMD Extensions on Itanium-based systems even though the intrinsics are supported.
- Use `_mm_empty()` after an MMX instruction if the next instruction is a floating-point (FP) instruction -- for example, before calculations on `float`, `double` or `long double`. You must be aware of all situations when your code generates an MMX instruction with the Intel® C++ Compiler, i.e.:
 - when using an MMX technology intrinsic
 - when using Streaming SIMD Extension integer intrinsics that use the `__m64` data type
 - when referencing an `__m64` data type variable
 - when using an MMX instruction through inline assembly
- Do not use `_mm_empty()` before an MMX instruction, since using `_mm_empty()` before an MMX instruction incurs an operation with no benefit (no-op).
- Use different functions for operations that use FP instructions and those that use MMX instructions. This eliminates the need to empty the multimedia state within the body of a critical loop.
- Use `_mm_empty()` during runtime initialization of `__m64` and FP data types. This ensures resetting the register between data type transitions.
- See the "Correct Usage" coding example.

Incorrect Usage	Correct Usage
<code>__m64 x = _m_paddd(y, z); float f = init();</code>	<code>__m64 x = _m_paddd(y, z); float f = (_mm_empty(), init());</code>

For more documentation on EMMS, visit the <http://developer.intel.com> Web site.

MMX(TM) Technology General Support Intrinsics

The prototypes for MMX(TM) technology intrinsics are in the `mmintrin.h` header file.

Intrinsic Name	Alternate Name	Corresponding Instruction	Operation	Signed	Saturation
<code>_m_empty</code>	<code>_mm_empty</code>	EMMS	Empty MM state	--	--
<code>_m_from_int</code>	<code>_mm_cvtsi32_si64</code>	MOVD	Convert from int	--	--
<code>_m_to_int</code>	<code>_mm_cvtsi64_si32</code>	MOVD	Convert from int	--	--
<code>_m_packsswb</code>	<code>_mm_packs_pi16</code>	PACKSSWB	Pack	Yes	Yes
<code>_m_packssdw</code>	<code>_mm_packs_pi32</code>	PACKSSDW	Pack	Yes	Yes
<code>_m_packuswb</code>	<code>_mm_packs_pu16</code>	PACKUSWB	Pack	No	Yes
<code>_m_punpckhbw</code>	<code>_mm_unpackhi_pi8</code>	PUNPCKHBW	Interleave	--	--
<code>_m_punpckhwd</code>	<code>_mm_unpackhi_pi16</code>	PUNPCKHWD	Interleave	--	--
<code>_m_punpckhdq</code>	<code>_mm_unpackhi_pi32</code>	PUNPCKHDQ	Interleave	--	--
<code>_m_punpcklbw</code>	<code>_mm_unpacklo_pi8</code>	PUNPCKLBW	Interleave	--	--
<code>_m_punpcklwd</code>	<code>_mm_unpacklo_pi16</code>	PUNPCKLWD	Interleave	--	--
<code>_m_punpckldq</code>	<code>_mm_unpacklo_pi32</code>	PUNPCKLDQ	Interleave	--	--

```
void _m_empty(void)
```

Empty the multimedia state.

```
__m64 _m_from_int(int i)
```

Convert the integer object `i` to a 64-bit `__m64` object. The integer value is zero-extended to 64 bits.

```
int _m_to_int(__m64 m)
```

Convert the lower 32 bits of the `__m64` object `m` to an integer.

```
__m64 _m_packsswb(__m64 m1, __m64 m2)
```

Pack the four 16-bit values from m1 into the lower four 8-bit values of the result with signed saturation, and pack the four 16-bit values from m2 into the upper four 8-bit values of the result with signed saturation.

```
__m64 _m_packssdw(__m64 m1, __m64 m2)
```

Pack the two 32-bit values from m1 into the lower two 16-bit values of the result with signed saturation, and pack the two 32-bit values from m2 into the upper two 16-bit values of the result with signed saturation.

```
__m64 _m_packuswb(__m64 m1, __m64 m2)
```

Pack the four 16-bit values from m1 into the lower four 8-bit values of the result with unsigned saturation, and pack the four 16-bit values from m2 into the upper four 8-bit values of the result with unsigned saturation.

```
__m64 _m_punpckhbw(__m64 m1, __m64 m2)
```

Interleave the four 8-bit values from the high half of m1 with the four values from the high half of m2. The interleaving begins with the data from m1.

```
__m64 _m_punpckhwd(__m64 m1, __m64 m2)
```

Interleave the two 16-bit values from the high half of m1 with the two values from the high half of m2. The interleaving begins with the data from m1.

```
__m64 _m_punpckhdq(__m64 m1, __m64 m2)
```

Interleave the 32-bit value from the high half of m1 with the 32-bit value from the high half of m2. The interleaving begins with the data from m1.

```
__m64 _m_punpcklbw(__m64 m1, __m64 m2)
```

Interleave the four 8-bit values from the low half of m1 with the four values from the low half of m2. The interleaving begins with the data from m1.

```
__m64 _m_punpcklwd(__m64 m1, __m64 m2)
```

Interleave the two 16-bit values from the low half of m1 with the two values from the low half of m2. The interleaving begins with the data from m1.

```
__m64 _m_punpckldq(__m64 m1, __m64 m2)
```

Interleave the 32-bit value from the low half of m1 with the 32-bit value from the low half of m2. The interleaving begins with the data from m1.

MMX(TM) Technology Packed Arithmetic Intrinsics

The prototypes for MMX(TM) technology intrinsics are in the `mmintrin.h` header file.

Packed Arithmetic Intrinsics, Part 1

Intrinsic Name	Alternate Name	Corresponding Instruction	Operation	Signed
<code>_m_paddb</code>	<code>_mm_add_pi8</code>	PADDB	Addition	--
<code>_m_paddw</code>	<code>_mm_add_pi16</code>	PADDW	Addition	--
<code>_m_paddd</code>	<code>_mm_add_pi32</code>	PADDD	Addition	--
<code>_m_paddsb</code>	<code>_mm_adds_pi8</code>	PADDSB	Addition	Yes
<code>_m_paddsw</code>	<code>_mm_adds_pi16</code>	PADDSW	Addition	Yes
<code>_m_paddusb</code>	<code>_mm_adds_pu8</code>	PADDUSB	Addition	No
<code>_m_paddusw</code>	<code>_mm_adds_pu16</code>	PADDUSW	Addition	No
<code>_m_psubb</code>	<code>_mm_sub_pi8</code>	PSUBB	Subtraction	--
<code>_m_psubw</code>	<code>_mm_sub_pi16</code>	PSUBW	Subtraction	--
<code>_m_psubd</code>	<code>_mm_sub_pi32</code>	PSUBD	Subtraction	--
<code>_m_psubsb</code>	<code>_mm_subs_pi8</code>	PSUBSB	Subtraction	Yes
<code>_m_psubsw</code>	<code>_mm_subs_pi16</code>	PSUBSW	Subtraction	Yes
<code>_m_psubusb</code>	<code>_mm_subs_pu8</code>	PSUBUSB	Subtraction	No
<code>_m_psubusw</code>	<code>_mm_subs_pu16</code>	PSUBUSW	Subtraction	No
<code>_m_pmaddwd</code>	<code>_mm_madd_pi16</code>	PMADDWD	Multiplication	--
<code>_m_pmulhw</code>	<code>_mm_mulhi_pi16</code>	PMULHW	Multiplication	Yes
<code>_m_pmullw</code>	<code>_mm_mullo_pi16</code>	PMULLW	Multiplication	--

Packed Arithmetic Intrinsics, Part 2

Intrinsic Name	Alternate Name	Corresponding Instruction	Argument Values/Bits	Result Values/Bits
<code>_m_paddb</code>	<code>_mm_add_pi8</code>	PADDB	8/8	8/8
<code>_m_paddw</code>	<code>_mm_add_pi16</code>	PADDW	4/16	4/16
<code>_m_paddd</code>	<code>_mm_add_pi32</code>	PADDD	2/32	2/32
<code>_m_paddsb</code>	<code>_mm_adds_pi8</code>	PADDSB	8/8	8/8
<code>_m_paddsw</code>	<code>_mm_adds_pi16</code>	PADDSW	4/16	4/16

Intrinsic Name	Alternate Name	Corresponding Instruction	Argument Values/Bits	Result Values/Bits
<code>_m_paddusb</code>	<code>_mm_adds_pu8</code>	PADDUSB	8/8	8/8
<code>_m_paddusw</code>	<code>_mm_adds_pu16</code>	PADDUSW	4/16	4/16
<code>_m_psubb</code>	<code>_mm_sub_pi8</code>	PSUBB	8/8	8/8
<code>_m_psubw</code>	<code>_mm_sub_pi16</code>	PSUBW	4/16	4/16
<code>_m_psubd</code>	<code>_mm_sub_pi32</code>	PSUBD	2/32	2/32
<code>_m_psubsb</code>	<code>_mm_subs_pi8</code>	PSUBSB	8/8	8/8
<code>_m_psubsw</code>	<code>_mm_subs_pi16</code>	PSUBSW	4/16	4/16
<code>_m_psubusb</code>	<code>_mm_subs_pu8</code>	PSUBUSB	8/8	8/8
<code>_m_psubusw</code>	<code>_mm_subs_pu16</code>	PSUBUSW	4/16	4/16
<code>_m_pmaddwd</code>	<code>_mm_madd_pi16</code>	PMADDWD	4/16	2/32
<code>_m_pmulhw</code>	<code>_mm_mulhi_pi16</code>	PMULHW	4/16	4/16 (high)
<code>_m_pmullw</code>	<code>_mm_mullo_pi16</code>	PMULLW	4/16	4/16 (low)

```
__m64 _m_paddb(__m64 m1, __m64 m2)
```

Add the eight 8-bit values in m1 to the eight 8-bit values in m2.

```
__m64 _m_paddw(__m64 m1, __m64 m2)
```

Add the four 16-bit values in m1 to the four 16-bit values in m2.

```
__m64 _m_paddd(__m64 m1, __m64 m2)
```

Add the two 32-bit values in m1 to the two 32-bit values in m2.

```
__m64 _m_paddsb(__m64 m1, __m64 m2)
```

Add the eight signed 8-bit values in m1 to the eight signed 8-bit values in m2 using saturating arithmetic.

```
__m64 _m_paddsw(__m64 m1, __m64 m2)
```

Add the four signed 16-bit values in m1 to the four signed 16-bit values in m2 using saturating arithmetic.

```
__m64 _m_paddusb(__m64 m1, __m64 m2)
```

Add the eight unsigned 8-bit values in m1 to the eight unsigned 8-bit values in m2 and using saturating arithmetic.

```
__m64 _m_paddusw(__m64 m1, __m64 m2)
```

Add the four unsigned 16-bit values in m1 to the four unsigned 16-bit values in m2 using saturating arithmetic.

```
__m64 _m_psubb(__m64 m1, __m64 m2)
```

Subtract the eight 8-bit values in m2 from the eight 8-bit values in m1.

```
__m64 _m_psubw(__m64 m1, __m64 m2)
```

Subtract the four 16-bit values in m2 from the four 16-bit values in m1.

```
__m64 _m_psubd(__m64 m1, __m64 m2)
```

Subtract the two 32-bit values in m2 from the two 32-bit values in m1.

```
__m64 _m_psubsb(__m64 m1, __m64 m2)
```

Subtract the eight signed 8-bit values in m2 from the eight signed 8-bit values in m1 using saturating arithmetic.

```
__m64 _m_psubsw(__m64 m1, __m64 m2)
```

Subtract the four signed 16-bit values in m2 from the four signed 16-bit values in m1 using saturating arithmetic.

```
__m64 _m_psubusb(__m64 m1, __m64 m2)
```

Subtract the eight unsigned 8-bit values in m2 from the eight unsigned 8-bit values in m1 using saturating arithmetic.

```
__m64 _m_psubusw(__m64 m1, __m64 m2)
```

Subtract the four unsigned 16-bit values in m2 from the four unsigned 16-bit values in m1 using saturating arithmetic.

```
__m64 _m_pmaddwd(__m64 m1, __m64 m2)
```

Multiply four 16-bit values in m1 by four 16-bit values in m2 producing four 32-bit intermediate results, which are then summed by pairs to produce two 32-bit results.

```
__m64 _m_pmulhw(__m64 m1, __m64 m2)
```

Multiply four signed 16-bit values in m1 by four signed 16-bit values in m2 and produce the high 16 bits of the four results.

```
__m64 _m_pmulw(__m64 m1, __m64 m2)
```

Multiply four 16-bit values in m1 by four 16-bit values in m2 and produce the low 16 bits of the four results.

MMX(TM) Technology Shift Intrinsics

The prototypes for MMX(TM) technology intrinsics are in the `mmintrin.h` header file.

Intrinsic Name	Alternate Name	Shift Direction	Shift Type	Corresponding Instruction
<code>_m_psllw</code>	<code>_mm_sll_pi16</code>	left	Logical	PSLLW
<code>_m_psllwi</code>	<code>_mm_slli_pi16</code>	left	Logical	PSLLWI
<code>_m_psllld</code>	<code>_mm_sll_pi32</code>	left	Logical	PSLLD
<code>_m_psllldi</code>	<code>_mm_slli_pi32</code>	left	Logical	PSLLDI
<code>_m_psllq</code>	<code>_mm_sll_si64</code>	left	Logical	PSLLQ
<code>_m_psllqi</code>	<code>_mm_slli_si64</code>	left	Logical	PSLLQI
<code>_m_psraw</code>	<code>_mm_sra_pi16</code>	right	Arithmetic	PSRAW
<code>_m_psrawi</code>	<code>_mm_srai_pi16</code>	right	Arithmetic	PSRAWI
<code>_m_psrads</code>	<code>_mm_sra_pi32</code>	right	Arithmetic	PSRAD
<code>_m_psradi</code>	<code>_mm_srai_pi32</code>	right	Arithmetic	PSRADI
<code>_m_psrlw</code>	<code>_mm_srl_pi16</code>	right	Logical	PSRLW
<code>_m_psrlwi</code>	<code>_mm_srli_pi16</code>	right	Logical	PSRLWI
<code>_m_psrld</code>	<code>_mm_srl_pi32</code>	right	Logical	PSRLD
<code>_m_psrldi</code>	<code>_mm_srli_pi32</code>	right	Logical	PSRLDI
<code>_m_psrlq</code>	<code>_mm_srl_si64</code>	right	Logical	PSRLQ
<code>_m_psrlqi</code>	<code>_mm_srli_si64</code>	right	Logical	PSRLQI

```
__m64 _m_psllw(__m64 m, __m64 count)
```

Shift four 16-bit values in `m` left the amount specified by `count` while shifting in zeros.

```
__m64 _m_psllwi(__m64 m, int count)
```

Shift four 16-bit values in `m` left the amount specified by `count` while shifting in zeros. For the best performance, `count` should be a constant.

```
__m64 _m_psllld(__m64 m, __m64 count)
```

Shift two 32-bit values in `m` left the amount specified by `count` while shifting in zeros.

```
__m64 _m_psllldi(__m64 m, int count)
```

Shift two 32-bit values in `m` left the amount specified by `count` while shifting in zeros. For the best performance, `count` should be a constant.

```
__m64 _m_psllq(__m64 m, __m64 count)
```

Shift the 64-bit value in `m` left the amount specified by `count` while shifting in zeros.

```
__m64 _m_psllqi(__m64 m, int count)
```

Shift the 64-bit value in `m` left the amount specified by `count` while shifting in zeros. For the best performance, `count` should be a constant.

```
__m64 _m_psraw(__m64 m, __m64 count)
```

Shift four 16-bit values in `m` right the amount specified by `count` while shifting in the sign bit.

```
__m64 _m_psrawi(__m64 m, int count)
```

Shift four 16-bit values in `m` right the amount specified by `count` while shifting in the sign bit. For the best performance, `count` should be a constant.

```
__m64 _m_psradi(__m64 m, __m64 count)
```

Shift two 32-bit values in `m` right the amount specified by `count` while shifting in the sign bit.

```
__m64 _m_psradi(__m64 m, int count)
```

Shift two 32-bit values in `m` right the amount specified by `count` while shifting in the sign bit. For the best performance, `count` should be a constant.

```
__m64 _m_psrld(__m64 m, __m64 count)
```

Shift four 16-bit values in `m` right the amount specified by `count` while shifting in zeros.

```
__m64 _m_psrldi(__m64 m, int count)
```

Shift four 16-bit values in `m` right the amount specified by `count` while shifting in zeros. For the best performance, `count` should be a constant.

```
__m64 _m_psrld(__m64 m, __m64 count)
```

Shift two 32-bit values in `m` right the amount specified by `count` while shifting in zeros.

```
__m64 _m_psrldi(__m64 m, int count)
```

Shift two 32-bit values in `m` right the amount specified by `count` while shifting in zeros. For the best performance, `count` should be a constant.

```
__m64 _m_psrldq(__m64 m, __m64 count)
```

Shift the 64-bit value in `m` right the amount specified by `count` while shifting in zeros.

```
__m64 _m_psrlqi(__m64 m, int count)
```

Shift the 64-bit value in `m` right the amount specified by `count` while shifting in zeros. For the best performance, `count` should be a constant.

MMX(TM) Technology Logical Intrinsics

The prototypes for MMX(TM) technology intrinsics are in the `mmintrin.h` header file.

Intrinsic Name	Alternate Name	Operation	Corresponding Instruction
<code>_m_pand</code>	<code>_mm_and_si64</code>	Bitwise AND	PAND
<code>_m_pandn</code>	<code>_mm_andnot_si64</code>	Logical NOT	PANDN
<code>_m_por</code>	<code>_mm_or_si64</code>	Bitwise OR	POR
<code>_m_pxor</code>	<code>_mm_xor_si64</code>	Bitwise Exclusive OR	PXOR

```
__m64 _m_pand(__m64 m1, __m64 m2)
```

Perform a bitwise AND of the 64-bit value in `m1` with the 64-bit value in `m2`.

```
__m64 _m_pandn(__m64 m1, __m64 m2)
```

Perform a logical NOT on the 64-bit value in `m1` and use the result in a bitwise AND with the 64-bit value in `m2`.

```
__m64 _m_por(__m64 m1, __m64 m2)
```

Perform a bitwise OR of the 64-bit value in `m1` with the 64-bit value in `m2`.

```
__m64 _m_pxor(__m64 m1, __m64 m2)
```

Perform a bitwise XOR of the 64-bit value in `m1` with the 64-bit value in `m2`.

MMX(TM) Technology Compare Intrinsics

The prototypes for MMX(TM) technology intrinsics are in the `mmintrin.h` header file.

Intrinsic Name	Alternate Name	Comparison	Number of Elements	Element Bit Size	Corresponding Instruction
<code>_m_pcmpeqb</code>	<code>_mm_cmpeq_pi8</code>	Equal	8	8	PCMPEQB
<code>_m_pcmpeqw</code>	<code>_mm_cmpeq_pi16</code>	Equal	4	16	PCMPEQW
<code>_m_pcmpeqd</code>	<code>_mm_cmpeq_pi32</code>	Equal	2	32	PCMPEQD
<code>_m_pcmpgtb</code>	<code>_mm_cmpgt_pi8</code>	Greater Than	8	8	PCMPGTB

Intrinsic Name	Alternate Name	Comparison	Number of Elements	Element Bit Size	Corresponding Instruction
<code>_m_pcmpgtw</code>	<code>_mm_cmpgt_pi16</code>	Greater Than	4	16	PCMPGTW
<code>_m_pcmpgtd</code>	<code>_mm_cmpgt_pi32</code>	Greater Than	2	32	PCMPGTD

```
__m64 _m_pcmpeqb(__m64 m1, __m64 m2)
```

If the respective 8-bit values in m1 are equal to the respective 8-bit values in m2 set the respective 8-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _m_pcmpeqw(__m64 m1, __m64 m2)
```

If the respective 16-bit values in m1 are equal to the respective 16-bit values in m2 set the respective 16-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _m_pcmpeqd(__m64 m1, __m64 m2)
```

If the respective 32-bit values in m1 are equal to the respective 32-bit values in m2 set the respective 32-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _m_pcmpgtb(__m64 m1, __m64 m2)
```

If the respective 8-bit values in m1 are greater than the respective 8-bit values in m2 set the respective 8-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _m_pcmpgtw(__m64 m1, __m64 m2)
```

If the respective 16-bit values in m1 are greater than the respective 16-bit values in m2 set the respective 16-bit resulting values to all ones, otherwise set them to all zeros.

```
__m64 _m_pcmpgtd(__m64 m1, __m64 m2)
```

If the respective 32-bit values in m1 are greater than the respective 32-bit values in m2 set the respective 32-bit resulting values to all ones, otherwise set them all to zeros.

MMX(TM) Technology Set Intrinsics

The prototypes for MMX(TM) technology intrinsics are in the `mmintrin.h` header file.

Intrinsic Name	Operation	Number of Elements	Element Bit Size	Signed	Reverse Order
<code>_mm_setzero_si64</code>	set to zero	1	64	No	No
<code>_mm_set_pi32</code>	set integer values	2	32	No	No
<code>_mm_set_pi16</code>	set integer values	4	16	No	No
<code>_mm_set_pi8</code>	set integer values	8	8	No	No

Intrinsic Name	Operation	Number of Elements	Element Bit Size	Signed	Reverse Order
<code>__mm_set1_pi32</code>	set integer values	2	32	Yes	No
<code>__mm_set1_pi16</code>	set integer values	4	16	Yes	No
<code>__mm_set1_pi8</code>	set integer values	8	8	Yes	No
<code>__mm_setr_pi32</code>	set integer values	2	32	No	Yes
<code>__mm_setr_pi16</code>	set integer values	4	16	No	Yes
<code>__mm_setr_pi8</code>	set integer values	8	8	No	Yes

 **Note**

In the following descriptions regarding the bits of the MMX register, bit 0 is the least significant and bit 63 is the most significant.

```
__m64 __mm_setzero_si64()
```

PXOR

Sets the 64-bit value to zero.

```
r := 0x0
```

```
__m64 __mm_set_pi32(int i1, int i0)
```

(composite) Sets the 2 signed 32-bit integer values.

```
r0 := i0
```

```
r1 := i1
```

```
__m64 __mm_set_pi16(short s3, short s2, short s1, short s0)
```

(composite) Sets the 4 signed 16-bit integer values.

```
r0 := w0
```

```
r1 := w1
```

```
r2 := w2
```

```
r3 := w3
```

```
__m64 __mm_set_pi8(char b7, char b6, char b5, char b4, char b3, char b2,
char b1, char b0)
```

(composite) Sets the 8 signed 8-bit integer values.

```
r0 := b0
```

```
r1 := b1
```

```
...
```

```
r7 := b7
```

```
__m64 __mm_set1_pi32(int i)
```

Sets the 2 signed 32-bit integer values to *i*.

```
r0 := i
```

```
r1 := i
```

```
__m64 _mm_set1_pi16(short s)
```

(composite) Sets the 4 signed 16-bit integer values to w.

```
r0 := w
r1 := w
r2 := w
r3 := w
```

```
__m64 _mm_set1_pi8(char b)
```

(composite) Sets the 8 signed 8-bit integer values to b

```
r0 := b
r1 := b
...
r7 := b
```

```
__m64 _mm_setr_pi32(int i1, int i0)
```

(composite) Sets the 2 signed 32-bit integer values in reverse order.

```
r0 := i0
r1 := i1
```

```
__m64 _mm_setr_pi16(short s3, short s2, short s1, short s0)
```

(composite) Sets the 4 signed 16-bit integer values in reverse order.

```
r0 := w0
r1 := w1
r2 := w2
r3 := w3
```

```
__m64 _mm_setr_pi8(char b7, char b6, char b5, char b4, char b3, char b2,
char b1, char b0)
```

(composite) Sets the 8 signed 8-bit integer values in reverse order.

```
r0 := b0
r1 := b1
...
r7 := b7
```

MMX(TM) Technology Intrinsics on Itanium® Architecture

MMX(TM) technology intrinsics provide access to the MMX technology instruction set on Itanium®-based systems. To provide source compatibility with the IA-32 architecture, these intrinsics are equivalent both in name and functionality to the set of IA-32-based MMX intrinsics.

Some intrinsics have more than one name. When one intrinsic has two names, both names generate the same instructions, but the first is preferred as it conforms to a newer naming standard.

The prototypes for MMX technology intrinsics are in the `mmintrin.h` header file.

Data Types

The C data type `__m64` is used when using MMX technology intrinsics. It can hold eight 8-bit values, four 16-bit values, two 32-bit values, or one 64-bit value.

The `__m64` data type is not a basic ANSI C data type. Therefore, observe the following usage restrictions:

- Use the new data type only on the left-hand side of an assignment, as a return value, or as a parameter. You cannot use it with other arithmetic expressions (" + ", " - ", and so on).
- Use the new data type as objects in aggregates, such as unions, to access the byte elements and structures; the address of an `__m64` object may be taken.
- Use new data types only with the respective intrinsics described in this documentation.

For complete details of the hardware instructions, see the Intel® Architecture MMX Technology Programmer's Reference Manual. For descriptions of data types, see the Intel® Architecture Software Developer's Manual, Volume 2.

Streaming SIMD Extensions

This section describes the C++ language-level features supporting the Streaming SIMD Extensions (SSE) in the Intel® C++ Compiler. These topics explain the following features of the intrinsics:

- Floating Point Intrinsics
- Arithmetic Operation Intrinsics
- Logical Operation Intrinsics
- Comparison Intrinsics
- Conversion Intrinsics
- Load Operations
- Set Operations
- Store Operations
- Cacheability Support
- Integer Intrinsics
- Memory and Initialization Intrinsics
- Miscellaneous Intrinsics
- Using Streaming SIMD Extensions on Itanium® Architecture

The prototypes for SSE intrinsics are in the `xmmintrin.h` header file.



Note

You can also use the single `ia32intrin.h` header file for any IA-32 intrinsics.

Floating-point Intrinsics for Streaming SIMD Extensions

You should be familiar with the hardware features provided by the Streaming SIMD Extensions (SSE) when writing programs with the intrinsics. The following are four important issues to keep in mind:

- Certain intrinsics, such as `_mm_loadr_ps` and `_mm_cmpgt_ss`, are not directly supported by the instruction set. While these intrinsics are convenient programming aids, be mindful that they may consist of more than one machine-language instruction.
- Floating-point data loaded or stored as `__m128` objects must be generally 16-byte-aligned.

- Some intrinsics require that their argument be immediates, that is, constant integers (literals), due to the nature of the instruction.
- The result of arithmetic operations acting on two NaN (Not a Number) arguments is undefined. Therefore, FP operations using NaN arguments will not match the expected behavior of the corresponding assembly instructions.

Arithmetic Operations for Streaming SIMD Extensions

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

Intrinsic	Instruction	Operation	R0	R1	R2	R3
<code>_mm_add_ss</code>	ADDSS	Addition	a0 [op] b0	a1	a2	a3
<code>_mm_add_ps</code>	ADDPS	Addition	a0 [op] b0	a1 [op] b1	a2 [op] b2	a3 [op] b3
<code>_mm_sub_ss</code>	SUBSS	Subtraction	a0 [op] b0	a1	a2	a3
<code>_mm_sub_ps</code>	SUBPS	Subtraction	a0 [op] b0	a1 [op] b1	a2 [op] b2	a3 [op] b3
<code>_mm_mul_ss</code>	MULSS	Multiplication	a0 [op] b0	a1	a2	a3
<code>_mm_mul_ps</code>	MULPS	Multiplication	a0 [op] b0	a1 [op] b1	a2 [op] b2	a3 [op] b3
<code>_mm_div_ss</code>	DIVSS	Division	a0 [op] b0	a1	a2	a3
<code>_mm_div_ps</code>	DIVPS	Division	a0 [op] b0	a1 [op] b1	a2 [op] b2	a3 [op] b3
<code>_mm_sqrt_ss</code>	SQRTSS	Squared Root	[op] a0	a1	a2	a3
<code>_mm_sqrt_ps</code>	SQRTPS	Squared Root	[op] a0	[op] b1	[op] b2	[op] b3
<code>_mm_rcp_ss</code>	RCPSS	Reciprocal	[op] a0	a1	a2	a3
<code>_mm_rcp_ps</code>	RCPPS	Reciprocal	[op] a0	[op] b1	[op] b2	[op] b3
<code>_mm_rsqrt_ss</code>	RSQRTSS	Reciprocal Square Root	[op] a0	a1	a2	a3
<code>_mm_rsqrt_ps</code>	RSQRTPS	Reciprocal Squared Root	[op] a0	[op] b1	[op] b2	[op] b3

Intrinsic	Instruction	Operation	R0	R1	R2	R3
<code>__mm_min_ss</code>	MINSS	Computes Minimum	<code>[op](a0,b0)</code>	a1	a2	a3
<code>__mm_min_ps</code>	MINPS	Computes Minimum	<code>[op](a0,b0)</code>	<code>[op](a1,b1)</code>	<code>[op](a2,b2)</code>	<code>[op](a3,b3)</code>
<code>__mm_max_ss</code>	MAXSS	Computes Maximum	<code>[op](a0,b0)</code>	a1	a2	a3
<code>__mm_max_ps</code>	MAXPS	Computes Maximum	<code>[op](a0,b0)</code>	<code>[op](a1,b1)</code>	<code>[op](a2,b2)</code>	<code>[op](a3,b3)</code>

```
__m128 __mm_add_ss(__m128 a, __m128 b)
```

Adds the lower SP FP (single-precision, floating-point) values of a and b; the upper 3 SP FP values are passed through from a.

```
r0 := a0 + b0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_add_ps(__m128 a, __m128 b)
```

Adds the four SP FP values of a and b.

```
r0 := a0 + b0
r1 := a1 + b1
r2 := a2 + b2
r3 := a3 + b3
```

```
__m128 __mm_sub_ss(__m128 a, __m128 b)
```

Subtracts the lower SP FP values of a and b. The upper 3 SP FP values are passed through from a.

```
r0 := a0 - b0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_sub_ps(__m128 a, __m128 b)
```

Subtracts the four SP FP values of a and b.

```
r0 := a0 - b0
r1 := a1 - b1
r2 := a2 - b2
r3 := a3 - b3
```

```
__m128 __mm_mul_ss(__m128 a, __m128 b)
```

Multiplies the lower SP FP values of a and b; the upper 3 SP FP values are passed through from a.

```
r0 := a0 * b0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_mul_ps(__m128 a, __m128 b)
```

Multiplies the four SP FP values of a and b.

```
r0 := a0 * b0
r1 := a1 * b1
r2 := a2 * b2
r3 := a3 * b3
```

```
__m128 _mm_div_ss(__m128 a, __m128 b )
```

Divides the lower SP FP values of a and b; the upper 3 SP FP values are passed through from a.

```
r0 := a0 / b0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_div_ps(__m128 a, __m128 b)
```

Divides the four SP FP values of a and b.

```
r0 := a0 / b0
r1 := a1 / b1
r2 := a2 / b2
r3 := a3 / b3
```

```
__m128 _mm_sqrt_ss(__m128 a)
```

Computes the square root of the lower SP FP value of a ; the upper 3 SP FP values are passed through.

```
r0 := sqrt(a0)
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_sqrt_ps(__m128 a)
```

Computes the square roots of the four SP FP values of a.

```
r0 := sqrt(a0)
r1 := sqrt(a1)
r2 := sqrt(a2)
r3 := sqrt(a3)
```

```
__m128 _mm_rcp_ss(__m128 a)
```

Computes the approximation of the reciprocal of the lower SP FP value of a; the upper 3 SP FP values are passed through.

```
r0 := recip(a0)
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_rcp_ps(__m128 a)
```

Computes the approximations of reciprocals of the four SP FP values of a.

```
r0 := recip(a0)
r1 := recip(a1)
r2 := recip(a2)
r3 := recip(a3)
```

```
__m128 _mm_rsqrt_ss(__m128 a)
```

Computes the approximation of the reciprocal of the square root of the lower SP FP value of a; the upper 3 SP FP values are passed through.

```
r0 := recip(sqrt(a0))
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_rsqrt_ps(__m128 a)
```

Computes the approximations of the reciprocals of the square roots of the four SP FP values of a.

```
r0 := recip(sqrt(a0))
r1 := recip(sqrt(a1))
r2 := recip(sqrt(a2))
r3 := recip(sqrt(a3))
```

```
__m128 _mm_min_ss(__m128 a, __m128 b)
```

Computes the minimum of the lower SP FP values of a and b; the upper 3 SP FP values are passed through from a.

```
r0 := min(a0, b0)
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_min_ps(__m128 a, __m128 b)
```

Computes the minimum of the four SP FP values of a and b.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
r2 := min(a2, b2)
r3 := min(a3, b3)
```

```
__m128 _mm_max_ss(__m128 a, __m128 b)
```

Computes the maximum of the lower SP FP values of a and b; the upper 3 SP FP values are passed through from a.

```
r0 := max(a0, b0)
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_max_ps(__m128 a, __m128 b)
```

Computes the maximum of the four SP FP values of a and b.

```
r0 := max(a0, b0)
r1 := max(a1, b1)
r2 := max(a2, b2)
r3 := max(a3, b3)
```

Logical Operations for Streaming SIMD Extensions

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

Intrinsic Name	Operation	Corresponding Instruction
<code>_mm_and_ps</code>	Bitwise AND	ANDPS
<code>_mm_andnot_ps</code>	Logical NOT	ANDNPS
<code>_mm_or_ps</code>	Bitwise OR	ORPS
<code>_mm_xor_ps</code>	Bitwise Exclusive OR	XORPS

```
__m128 _mm_and_ps(__m128 a, __m128 b)
```

Computes the bitwise And of the four SP FP values of a and b.

```
r0 := a0 & b0
r1 := a1 & b1
r2 := a2 & b2
r3 := a3 & b3
```

```
__m128 _mm_andnot_ps(__m128 a, __m128 b)
```

Computes the bitwise AND-NOT of the four SP FP values of a and b.

```
r0 := ~a0 & b0
r1 := ~a1 & b1
r2 := ~a2 & b2
r3 := ~a3 & b3
```

```
__m128 _mm_or_ps(__m128 a, __m128 b)
```

Computes the bitwise OR of the four SP FP values of a and b.

```
r0 := a0 | b0
r1 := a1 | b1
r2 := a2 | b2
r3 := a3 | b3
```

```
__m128 _mm_xor_ps(__m128 a, __m128 b)
```

Computes bitwise XOR (exclusive-or) of the four SP FP values of a and b.

```
r0 := a0 ^ b0
r1 := a1 ^ b1
r2 := a2 ^ b2
r3 := a3 ^ b3
```

Comparisons for Streaming SIMD Extensions

Each comparison intrinsic performs a comparison of a and b. For the packed form, the four SP FP values of a and b are compared, and a 128-bit mask is returned. For the scalar form, the lower SP FP values of a and b are compared, and a 32-bit mask is returned; the upper three SP FP values are passed through from a. The mask is set to 0xffffffff for each element where the comparison is true and 0x0 where the comparison is false.

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

Intrinsic Name	Comparison	Corresponding Instruction
<code>_mm_cmpeq_ss</code>	Equal	CMPEQSS
<code>_mm_cmpeq_ps</code>	Equal	CMPEQPS
<code>_mm_cmplt_ss</code>	Less Than	CMPLTSS
<code>_mm_cmplt_ps</code>	Less Than	CMPLTPS
<code>_mm_cmple_ss</code>	Less Than or Equal	CMPLSS
<code>_mm_cmple_ps</code>	Less Than or Equal	CMPLPS
<code>_mm_cmpgt_ss</code>	Greater Than	CMPLTSS
<code>_mm_cmpgt_ps</code>	Greater Than	CMPLTPS
<code>_mm_cmpge_ss</code>	Greater Than or Equal	CMPLSS
<code>_mm_cmpge_ps</code>	Greater Than or Equal	CMPLPS
<code>_mm_cmpneq_ss</code>	Not Equal	CMPNEQSS
<code>_mm_cmpneq_ps</code>	Not Equal	CMPNEQPS
<code>_mm_cmpnlt_ss</code>	Not Less Than	CMPNLTSS
<code>_mm_cmpnlt_ps</code>	Not Less Than	CMPNLTPS
<code>_mm_cmpnle_ss</code>	Not Less Than or Equal	CMPNLESS
<code>_mm_cmpnle_ps</code>	Not Less Than or Equal	CMPNLEPS
<code>_mm_cmpngt_ss</code>	Not Greater Than	CMPNLTSS
<code>_mm_cmpngt_ps</code>	Not Greater Than	CMPNLTPS
<code>_mm_cmpnge_ss</code>	Not Greater Than or Equal	CMPNLESS
<code>_mm_cmpnge_ps</code>	Not Greater Than or Equal	CMPNLEPS
<code>_mm_cmpord_ss</code>	Ordered	CMPORDSS
<code>_mm_cmpord_ps</code>	Ordered	CMPORDPS
<code>_mm_cmpunord_ss</code>	Unordered	CMPUNORDSS

Intrinsic Name	Comparison	Corresponding Instruction
<code>__mm_cmpunord_ps</code>	Unordered	CMPUNORDPS
<code>__mm_comieq_ss</code>	Equal	COMISS
<code>__mm_comilt_ps</code>	Less Than	COMISS
<code>__mm_comile_ss</code>	Less Than or Equal	COMISS
<code>__mm_comigt_ss</code>	Greater Than	COMISS
<code>__mm_comige_ss</code>	Greater Than or Equal	COMISS
<code>__mm_comineq_ss</code>	Not Equal	COMISS
<code>__mm_ucomieq_ss</code>	Equal	UCOMISS
<code>__mm_ucomilt_ss</code>	Less Than	UCOMISS
<code>__mm_ucomile_ss</code>	Less Than or Equal	UCOMISS
<code>__mm_ucomigt_ss</code>	Greater Than	UCOMISS
<code>__mm_ucomige_ss</code>	Greater Than or Equal	UCOMISS
<code>__mm_ucomineq_ss</code>	Not Equal	UCOMISS

```
__m128 __mm_cmpeq_ss(__m128 a, __m128 b)
```

Compare for equality.

```
r0 := (a0 == b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpeq_ps(__m128 a, __m128 b)
```

Compare for equality.

```
r0 := (a0 == b0) ? 0xffffffff : 0x0
r1 := (a1 == b1) ? 0xffffffff : 0x0
r2 := (a2 == b2) ? 0xffffffff : 0x0
r3 := (a3 == b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmplt_ss(__m128 a, __m128 b)
```

Compare for less-than.

```
r0 := (a0 < b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmplt_ps(__m128 a, __m128 b)
```

Compare for less-than.

```
r0 := (a0 < b0) ? 0xffffffff : 0x0
r1 := (a1 < b1) ? 0xffffffff : 0x0
r2 := (a2 < b2) ? 0xffffffff : 0x0
r3 := (a3 < b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmples_s(__m128 a, __m128 b)
```

Compare for less-than-or-equal.

```
r0 := (a0 <= b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmples_ps(__m128 a, __m128 b)
```

Compare for less-than-or-equal.

```
r0 := (a0 <= b0) ? 0xffffffff : 0x0
r1 := (a1 <= b1) ? 0xffffffff : 0x0
r2 := (a2 <= b2) ? 0xffffffff : 0x0
r3 := (a3 <= b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpgts_s(__m128 a, __m128 b)
```

Compare for greater-than.

```
r0 := (a0 > b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpgts_ps(__m128 a, __m128 b)
```

Compare for greater-than.

```
r0 := (a0 > b0) ? 0xffffffff : 0x0
r1 := (a1 > b1) ? 0xffffffff : 0x0
r2 := (a2 > b2) ? 0xffffffff : 0x0
r3 := (a3 > b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpgts_s(__m128 a, __m128 b)
```

Compare for greater-than-or-equal.

```
r0 := (a0 >= b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpgts_ps(__m128 a, __m128 b)
```

Compare for greater-than-or-equal.

```
r0 := (a0 >= b0) ? 0xffffffff : 0x0
r1 := (a1 >= b1) ? 0xffffffff : 0x0
r2 := (a2 >= b2) ? 0xffffffff : 0x0
r3 := (a3 >= b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpneqs_s(__m128 a, __m128 b)
```

Compare for inequality.

```
r0 := (a0 != b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpneqs_ps(__m128 a, __m128 b)
```

Compare for inequality.

```
r0 := (a0 != b0) ? 0xffffffff : 0x0
r1 := (a1 != b1) ? 0xffffffff : 0x0
r2 := (a2 != b2) ? 0xffffffff : 0x0
r3 := (a3 != b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpnlt_ss(__m128 a, __m128 b)
```

Compare for not-less-than.

```
r0 := !(a0 < b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpnlt_ps(__m128 a, __m128 b)
```

Compare for not-less-than.

```
r0 := !(a0 < b0) ? 0xffffffff : 0x0
r1 := !(a1 < b1) ? 0xffffffff : 0x0
r2 := !(a2 < b2) ? 0xffffffff : 0x0
r3 := !(a3 < b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpnle_ss(__m128 a, __m128 b)
```

Compare for not-less-than-or-equal.

```
r0 := !(a0 <= b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpnle_ps(__m128 a, __m128 b)
```

Compare for not-less-than-or-equal.

```
r0 := !(a0 <= b0) ? 0xffffffff : 0x0
r1 := !(a1 <= b1) ? 0xffffffff : 0x0
r2 := !(a2 <= b2) ? 0xffffffff : 0x0
r3 := !(a3 <= b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpngt_ss(__m128 a, __m128 b)
```

Compare for not-greater-than.

```
r0 := !(a0 > b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpngt_ps(__m128 a, __m128 b)
```

Compare for not-greater-than.

```
r0 := !(a0 > b0) ? 0xffffffff : 0x0
r1 := !(a1 > b1) ? 0xffffffff : 0x0
r2 := !(a2 > b2) ? 0xffffffff : 0x0
r3 := !(a3 > b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpnge_ss(__m128 a, __m128 b)
```

Compare for not-greater-than-or-equal.

```
r0 := !(a0 >= b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpnge_ps(__m128 a, __m128 b)
```

Compare for not-greater-than-or-equal.

```
r0 := !(a0 >= b0) ? 0xffffffff : 0x0
r1 := !(a1 >= b1) ? 0xffffffff : 0x0
r2 := !(a2 >= b2) ? 0xffffffff : 0x0
r3 := !(a3 >= b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpord_ss(__m128 a, __m128 b)
```

Compare for ordered.

```
r0 := (a0 ord? b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpord_ps(__m128 a, __m128 b)
```

Compare for ordered.

```
r0 := (a0 ord? b0) ? 0xffffffff : 0x0
r1 := (a1 ord? b1) ? 0xffffffff : 0x0
r2 := (a2 ord? b2) ? 0xffffffff : 0x0
r3 := (a3 ord? b3) ? 0xffffffff : 0x0
```

```
__m128 __mm_cmpunord_ss(__m128 a, __m128 b)
```

Compare for unordered.

```
r0 := (a0 unord? b0) ? 0xffffffff : 0x0
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 __mm_cmpunord_ps(__m128 a, __m128 b)
```

Compare for unordered.

```
r0 := (a0 unord? b0) ? 0xffffffff : 0x0
r1 := (a1 unord? b1) ? 0xffffffff : 0x0
r2 := (a2 unord? b2) ? 0xffffffff : 0x0
r3 := (a3 unord? b3) ? 0xffffffff : 0x0
```

```
int __mm_comieq_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a equal to b. If a and b are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 == b0) ? 0x1 : 0x0
```

```
int __mm_comilt_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a less than b. If a is less than b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 < b0) ? 0x1 : 0x0
```

```
int __mm_comile_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a less than or equal to b. If a is less than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 <= b0) ? 0x1 : 0x0
```

```
int __mm_comigt_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a greater than b. If a is greater than b are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 > b0) ? 0x1 : 0x0
```

```
int __mm_comige_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a greater than or equal to b. If a is greater than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 >= b0) ? 0x1 : 0x0
```

```
int __mm_comineq_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a not equal to b. If a and b are not equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 != b0) ? 0x1 : 0x0
```

```
int __mm_ucomieq_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a equal to b. If a and b are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 == b0) ? 0x1 : 0x0
```

```
int __mm_ucomilt_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a less than b. If a is less than b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 < b0) ? 0x1 : 0x0
```

```
int __mm_ucomile_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a less than or equal to b. If a is less than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 <= b0) ? 0x1 : 0x0
```

```
int __mm_ucomigt_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a greater than b. If a is greater than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 > b0) ? 0x1 : 0x0
```

```
int __mm_ucomige_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a greater than or equal to b. If a is greater than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 >= b0) ? 0x1 : 0x0
```

```
int __mm_ucomineq_ss(__m128 a, __m128 b)
```

Compares the lower SP FP value of a and b for a not equal to b. If a and b are not equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 != b0) ? 0x1 : 0x0
```

Conversion Operations for Streaming SIMD Extensions

The conversions operations are listed in the following table followed by a description of each intrinsic with the most recent mnemonic naming convention. The alternate name is provided in case you have used these intrinsics before.

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

Intrinsic Name	Alternate Name	Corresponding Instruction
<code>_mm_cvt_ss2si</code>	<code>_mm_cvtss_si32</code>	CVTSS2SI
<code>_mm_cvt_ps2pi</code>	<code>_mm_cvtps_pi32</code>	CVTPS2PI
<code>_mm_cvtt_ss2si</code>	<code>_mm_cvttss_si32</code>	CVTTSS2SI
<code>_mm_cvtt_ps2pi</code>	<code>_mm_cvttps_pi32</code>	CVTTPS2PI
<code>_mm_cvt_si2ss</code>	<code>_mm_cvtsi32_ss</code>	CVTSI2SS
<code>_mm_cvt_pi2ps</code>	<code>_mm_cvtpi32_ps</code>	CVTTPS2PI
<code>_mm_cvtpi16_ps</code>		composite
<code>_mm_cvtpu16_ps</code>		composite
<code>_mm_cvtpi8_ps</code>		composite
<code>_mm_cvtpu8_ps</code>		composite
<code>_mm_cvtpi32x2_ps</code>		composite
<code>_mm_cvtps_pi16</code>		composite
<code>_mm_cvtps_pi8</code>		composite

```
int _mm_cvt_ss2si(__m128 a)
```

Convert the lower SP FP value of `a` to a 32-bit integer according to the current rounding mode.

```
r := (int)a0
```

```
__m64 _mm_cvt_ps2pi(__m128 a)
```

Convert the two lower SP FP values of `a` to two 32-bit integers according to the current rounding mode, returning the integers in packed form.

```
r0 := (int)a0
```

```
r1 := (int)a1
```

```
int _mm_cvtt_ss2si(__m128 a)
```

Convert the lower SP FP value of `a` to a 32-bit integer with truncation.

```
r := (int)a0
```

```
__m64 _mm_cvtt_ps2pi(__m128 a)
```

Convert the two lower SP FP values of *a* to two 32-bit integer with truncation, returning the integers in packed form.

```
r0 := (int)a0
r1 := (int)a1
```

```
__m128 _mm_cvt_si2ss(__m128, int)
```

Convert the 32-bit integer value *b* to an SP FP value; the upper three SP FP values are passed through from *a*.

```
r0 := (float)b
r1 := a1 ; r2 := a2 ; r3 := a3
```

```
__m128 _mm_cvt_pi2ps(__m128, __m64)
```

Convert the two 32-bit integer values in packed form in *b* to two SP FP values; the upper two SP FP values are passed through from *a*.

```
r0 := (float)b0
r1 := (float)b1
r2 := a2
r3 := a3
```

```
__inline __m128 _mm_cvtpi16_ps(__m64 a)
```

Convert the four 16-bit signed integer values in *a* to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)a2
r3 := (float)a3
```

```
__inline __m128 _mm_cvtpu16_ps(__m64 a)
```

Convert the four 16-bit unsigned integer values in *a* to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)a2
r3 := (float)a3
```

```
__inline __m128 _mm_cvtpi8_ps(__m64 a)
```

Convert the lower four 8-bit signed integer values in *a* to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)a2
r3 := (float)a3
```

```
__inline __m128 _mm_cvtpu8_ps(__m64 a)
```

Convert the lower four 8-bit unsigned integer values in *a* to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)a2
r3 := (float)a3
```

```
__inline __m128 _mm_cvtpi32x2_ps(__m64 a, __m64 b)
```

Convert the two 32-bit signed integer values in *a* and the two 32-bit signed integer values in *b* to four single precision FP values.

```
r0 := (float)a0
r1 := (float)a1
r2 := (float)b0
r3 := (float)b1
```

```
__inline __m64 _mm_cvtps_pi16(__m128 a)
```

Convert the four single precision FP values in *a* to four signed 16-bit integer values.

```
r0 := (short)a0
r1 := (short)a1
r2 := (short)a2
r3 := (short)a3
```

```
__inline __m64 _mm_cvtps_pi8(__m128 a)
```

Convert the four single precision FP values in *a* to the lower four signed 8-bit integer values of the result.

```
r0 := (char)a0
r1 := (char)a1
r2 := (char)a2
r3 := (char)a3
```

Load Operations for Streaming SIMD Extensions

See summary table in Summary of Memory and Initialization topic.

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

```
__m128 _mm_load_ss(float * p )
```

Loads an SP FP value into the low word and clears the upper three words.

```
r0 := *p
r1 := 0.0 ; r2 := 0.0 ; r3 := 0.0
```

```
__m128 _mm_load_ps1(float * p )
```

Loads a single SP FP value, copying it into all four words.

```
r0 := *p
r1 := *p
r2 := *p
r3 := *p
```

```
__m128 _mm_load_ps(float * p )
```

Loads four SP FP values. The address must be 16-byte-aligned.

```
r0 := p[0]
r1 := p[1]
r2 := p[2]
r3 := p[3]
```

```
__m128 _mm_loadu_ps(float * p)
```

Loads four SP FP values. The address need not be 16-byte-aligned.

```
r0 := p[0]
r1 := p[1]
r2 := p[2]
r3 := p[3]
```

```
__m128 _mm_loadr_ps(float * p)
```

Loads four SP FP values in reverse order. The address must be 16-byte-aligned.

```
r0 := p[3]
r1 := p[2]
r2 := p[1]
r3 := p[0]
```

Set Operations for Streaming SIMD Extensions

See summary table in Summary of Memory and Initialization topic.

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

```
__m128 _mm_set_ss(float w )
```

Sets the low word of an SP FP value to `w` and clears the upper three words.

```
r0 := w
r1 := r2 := r3 := 0.0
```

```
__m128 _mm_set_ps1(float w )
```

Sets the four SP FP values to `w`.

```
r0 := r1 := r2 := r3 := w
```

```
__m128 _mm_set_ps(float z, float y, float x, float w )
```

Sets the four SP FP values to the four inputs.

```
r0 := w
r1 := x
r2 := y
r3 := z
```

```
__m128 _mm_setr_ps(float z, float y, float x, float w )
```

Sets the four SP FP values to the four inputs in reverse order.

```
r0 := z
r1 := y
r2 := x
r3 := w
```

```
__m128 _mm_setzero_ps(void)
```

Clears the four SP FP values.

```
r0 := r1 := r2 := r3 := 0.0
```

Store Operations for Streaming SIMD Extensions

See summary table in Summary of Memory and Initialization topic.

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

```
void _mm_store_ss(float * p, __m128 a)
```

Stores the lower SP FP value.

```
*p := a0
```

```
void _mm_store_ps1(float * p, __m128 a )
```

Stores the lower SP FP value across four words.

```
p[0] := a0
```

```
p[1] := a0
```

```
p[2] := a0
```

```
p[3] := a0
```

```
void _mm_store_ps(float *p, __m128 a)
```

Stores four SP FP values. The address must be 16-byte-aligned.

```
p[0] := a0
```

```
p[1] := a1
```

```
p[2] := a2
```

```
p[3] := a3
```

```
void _mm_storeu_ps(float *p, __m128 a)
```

Stores four SP FP values. The address need not be 16-byte-aligned.

```
p[0] := a0
```

```
p[1] := a1
```

```
p[2] := a2
```

```
p[3] := a3
```

```
void _mm_storer_ps(float * p, __m128 a )
```

Stores four SP FP values in reverse order. The address must be 16-byte-aligned.

```
p[0] := a3
```

```
p[1] := a2
```

```
p[2] := a1
```

```
p[3] := a0
```

```
__m128 _mm_move_ss( __m128 a, __m128 b)
```

Sets the low word to the SP FP value of `b`. The upper 3 SP FP values are passed through from `a`.

```
r0 := b0
```

```
r1 := a1
```

```
r2 := a2
```

```
r3 := a3
```

Cacheability Support Using Streaming SIMD Extensions

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

```
void _mm_pause(void)
```

The execution of the next instruction is delayed an implementation specific amount of time. The instruction does not modify the architectural state. This intrinsic provides especially significant performance gain.

PAUSE Intrinsic

The PAUSE intrinsic is used in spin-wait loops with the processors implementing dynamic execution (especially out-of-order execution). In the spin-wait loop, PAUSE improves the speed at which the code detects the release of the lock. For dynamic scheduling, the PAUSE instruction reduces the penalty of exiting from the spin-loop.

Example of loop with the PAUSE instruction:

```
spin_loop:pause
cmp eax, A
jne spin_loop
```

In this example, the program spins until memory location A matches the value in register `eax`. The code sequence that follows shows a test-and-test-and-set. In this example, the spin occurs only after the attempt to get a lock has failed.

```
get_lock: mov eax, 1
xchg eax, A ; Try to get lock
cmp eax, 0 ; Test if successful
jne spin_loop
```

Critical Section

```
// critical_section code
mov A, 0 ; Release lock
jmp continue
spin_loop: pause;
// spin-loop hint
cmp 0, A ;
// check lock availability
jne spin_loop
jmp get_lock
// continue: other code
```

Note that the first branch is predicted to fall-through to the critical section in anticipation of successfully gaining access to the lock. It is highly recommended that all spin-wait loops include the PAUSE instruction. Since PAUSE is backwards compatible to all existing IA-32 processor generations, a test for processor type (a CPUID test) is not needed. All legacy processors will execute PAUSE as a NOP, but in processors which use the PAUSE as a hint there can be significant performance benefit.

Integer Intrinsics Using Streaming SIMD Extensions

The integer intrinsics are listed in the following table followed by a description of each intrinsic with the most recent mnemonic naming convention.

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

Intrinsic Name	Alternate Name	Operation	Corresponding Instruction
<code>_m_pextrw</code>	<code>_mm_extract_pi16</code>	Extract one of four words	PEXTRW
<code>_m_pinsrw</code>	<code>_mm_insert_pi16</code>	Insert a word	PINSRW
<code>_m_pmaxsw</code>	<code>_mm_max_pi16</code>	Compute the maximum	PMAXSW
<code>_m_pmaxub</code>	<code>_mm_max_pu8</code>	Compute the maximum, unsigned	PMAXUB
<code>_m_pminsw</code>	<code>_mm_min_pi16</code>	Compute the minimum	PMINSW
<code>_m_pminub</code>	<code>_mm_min_pu8</code>	Compute the minimum, unsigned	PMINUB
<code>_m_pmovmskb</code>	<code>_mm_movemask_pi8</code>	Create an eight-bit mask	PMOVMASKB
<code>_m_pmulhuw</code>	<code>_mm_mulhi_pu16</code>	Multiply, return high bits	PMULHUW
<code>_m_pshufw</code>	<code>_mm_shuffle_pi16</code>	Return a combination of four words	PSHUFW
<code>_m_maskmovq</code>	<code>_mm_maskmove_si64</code>	Conditional Store	MASKMOVQ
<code>_m_pavgb</code>	<code>_mm_avg_pu8</code>	Compute rounded average	PAVGB
<code>_m_pavgw</code>	<code>_mm_avg_pu16</code>	Compute rounded average	PAVGW
<code>_m_psadbw</code>	<code>_mm_sad_pu8</code>	Compute sum of absolute differences	PSADBW

For these intrinsics you need to empty the multimedia state for the mmx register. See The EMMS Instruction: Why You Need It and When to Use It topic for more details.

```
int _m_pextrw(__m64 a, int n)
```

Extracts one of the four words of `a`. The selector `n` must be an immediate.

```
r := (n==0) ? a0 : ( (n==1) ? a1 : ( (n==2) ? a2 : a3 ) )
```

```
__m64 _m_pinsrw(__m64 a, int d, int n)
```

Inserts word `d` into one of four words of `a`. The selector `n` must be an immediate.

```
r0 := (n==0) ? d : a0;
r1 := (n==1) ? d : a1;
r2 := (n==2) ? d : a2;
r3 := (n==3) ? d : a3;
```

```
__m64 _m_pmaxsw(__m64 a, __m64 b)
```

Computes the element-wise maximum of the words in a and b.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
r2 := min(a2, b2)
r3 := min(a3, b3)
```

```
__m64 _m_pmaxub(__m64 a, __m64 b)
```

Computes the element-wise maximum of the unsigned bytes in a and b.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
...
r7 := min(a7, b7)
```

```
__m64 _m_pminsw(__m64 a, __m64 b)
```

Computes the element-wise minimum of the words in a and b.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
r2 := min(a2, b2)
r3 := min(a3, b3)
```

```
__m64 _m_pminub(__m64 a, __m64 b)
```

Computes the element-wise minimum of the unsigned bytes in a and b.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
...
r7 := min(a7, b7)
```

```
int _m_pmovmskb(__m64 a)
```

Creates an 8-bit mask from the most significant bits of the bytes in a.

```
r := sign(a7)<<7 | sign(a6)<<6 | ... | sign(a0)
```

```
__m64 _m_pmulhuw(__m64 a, __m64 b)
```

Multiplies the unsigned words in a and b, returning the upper 16 bits of the 32-bit intermediate results.

```
r0 := hiword(a0 * b0)
r1 := hiword(a1 * b1)
r2 := hiword(a2 * b2)
r3 := hiword(a3 * b3)
```

```
__m64 _m_pshufw(__m64 a, int n)
```

Returns a combination of the four words of a. The selector n must be an immediate.

```
r0 := word (n&0x3) of a
r1 := word ((n>>2)&0x3) of a
r2 := word ((n>>4)&0x3) of a
r3 := word ((n>>6)&0x3) of a
```

```
void _m_maskmovq(__m64 d, __m64 n, char *p)
```

Conditionally store byte elements of `d` to address `p`. The high bit of each byte in the selector `n` determines whether the corresponding byte in `d` will be stored.

```
if (sign(n0)) p[0] := d0
if (sign(n1)) p[1] := d1
...
if (sign(n7)) p[7] := d7
```

```
__m64 _m_pavgb(__m64 a, __m64 b)
```

Computes the (rounded) averages of the unsigned bytes in `a` and `b`.

```
t = (unsigned short)a0 + (unsigned short)b0
r0 = (t >> 1) | (t & 0x01)
...
t = (unsigned short)a7 + (unsigned short)b7
r7 = (unsigned char)((t >> 1) | (t & 0x01))
```

```
__m64 _m_pavgw(__m64 a, __m64 b)
```

Computes the (rounded) averages of the unsigned words in `a` and `b`.

```
t = (unsigned int)a0 + (unsigned int)b0
r0 = (t >> 1) | (t & 0x01)
...
t = (unsigned word)a7 + (unsigned word)b7
r7 = (unsigned short)((t >> 1) | (t & 0x01))
```

```
__m64 _m_psadbw(__m64 a, __m64 b)
```

Computes the sum of the absolute differences of the unsigned bytes in `a` and `b`, returning the value in the lower word. The upper three words are cleared.

```
r0 = abs(a0-b0) + ... + abs(a7-b7)
r1 = r2 = r3 = 0
```

Memory and Initialization Using Streaming SIMD Extensions

This section describes the `load`, `set`, and `store` operations, which let you load and store data into memory. The `load` and `set` operations are similar in that both initialize `__m128` data. However, the `set` operations take a float argument and are intended for initialization with constants, whereas the `load` operations take a floating point argument and are intended to mimic the instructions for loading data from memory. The `store` operation assigns the initialized data to the address.

The intrinsics are listed in the following table. Syntax and a brief description are contained the following topics.

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

Intrinsic Name	Alternate Name	Operation	Corresponding Instruction
<code>_mm_load_ss</code>		Load the low value and clear the three high values	MOVSS
<code>_mm_load_ps1</code>	<code>_mm_load1_ps</code>	Load one value into all four words	MOVSS + Shuffling
<code>_mm_load_ps</code>		Load four values, address aligned	MOVAPS
<code>_mm_loadu_ps</code>		Load four values, address unaligned	MOVUPS
<code>_mm_loadr_ps</code>		Load four values, in reverse order	MOVAPS + Shuffling
<code>_mm_set_ss</code>		Set the low value and clear the three high values	Composite
<code>_mm_set_ps1</code>	<code>_mm_set1_ps</code>	Set all four words with the same value	Composite
<code>_mm_set_ps</code>		Set four values, address aligned	Composite
<code>_mm_setr_ps</code>		Set four values, in reverse order	Composite
<code>_mm_setzero_ps</code>		Clear all four values	Composite
<code>_mm_store_ss</code>		Store the low value	MOVSS
<code>_mm_store_ps1</code>	<code>_mm_store1_ps</code>	Store the low value across all four words. The address must be 16-byte aligned.	Shuffling + MOVSS
<code>_mm_store_ps</code>		Store four values, address aligned	MOVAPS
<code>_mm_storeu_ps</code>		Store four values, address unaligned	MOVUPS
<code>_mm_storer_ps</code>		Store four values, in reverse order	MOVAPS + Shuffling
<code>_mm_move_ss</code>		Set the low word, and pass in three high values	MOVSS
<code>_mm_getcsr</code>		Return register contents	STMXCSR
<code>_mm_setcsr</code>		Control Register	LDMXCSR

Intrinsic Name	Alternate Name	Operation	Corresponding Instruction
<code>_mm_prefetch</code>			
<code>_mm_stream_pi</code>			
<code>_mm_stream_ps</code>			
<code>_mm_sfence</code>			
<code>_mm_cvtss_f32</code>			

```
__m128 _mm_load_ss(float const*a)
```

Loads an SP FP value into the low word and clears the upper three words.

```
r0 := *a
r1 := 0.0 ; r2 := 0.0 ; r3 := 0.0
```

```
__m128 _mm_load_ps1(float const*a)
```

Loads a single SP FP value, copying it into all four words.

```
r0 := *a
r1 := *a
r2 := *a
r3 := *a
```

```
__m128 _mm_load_ps(float const*a)
```

Loads four SP FP values. The address must be 16-byte-aligned.

```
r0 := a[0]
r1 := a[1]
r2 := a[2]
r3 := a[3]
```

```
__m128 _mm_loadu_ps(float const*a)
```

Loads four SP FP values. The address need not be 16-byte-aligned.

```
r0 := a[0]
r1 := a[1]
r2 := a[2]
r3 := a[3]
```

```
__m128 _mm_loadr_ps(float const*a)
```

Loads four SP FP values in reverse order. The address must be 16-byte-aligned.

```
r0 := a[3]
r1 := a[2]
r2 := a[1]
r3 := a[0]
```

```
__m128 _mm_set_ss(float a)
```

Sets the low word of an SP FP value to a and clears the upper three words.

```
r0 := c
r1 := r2 := r3 := 0.0
```

```
__m128 _mm_set_ps1(float a)
```

Sets the four SP FP values to a.

```
r0 := r1 := r2 := r3 := a
```

```
__m128 _mm_set_ps(float a, float b, float c, float d)
```

Sets the four SP FP values to the four inputs.

```
r0 := a
r1 := b
r2 := c
r3 := d
```

```
__m128 _mm_setr_ps(float a, float b, float c, float d)
```

Sets the four SP FP values to the four inputs in reverse order.

```
r0 := d
r1 := c
r2 := b
r3 := a
```

```
__m128 _mm_setzero_ps(void)
```

Clears the four SP FP values.

```
r0 := r1 := r2 := r3 := 0.0
```

```
void _mm_store_ss(float *v, __m128 a)
```

Stores the lower SP FP value.

```
*v := a0
```

```
void _mm_store_ps1(float *v, __m128 a)
```

Stores the lower SP FP value across four words.

```
v[0] := a0
v[1] := a0
v[2] := a0
v[3] := a0
```

```
void _mm_store_ps(float *v, __m128 a)
```

Stores four SP FP values. The address must be 16-byte-aligned.

```
v[0] := a0
v[1] := a1
v[2] := a2
v[3] := a3
```

```
void _mm_storeu_ps(float *v, __m128 a)
```

Stores four SP FP values. The address need not be 16-byte-aligned.

```
v[0] := a0
v[1] := a1
v[2] := a2
v[3] := a3
```

```
void _mm_storer_ps(float *v, __m128 a)
```

Stores four SP FP values in reverse order. The address must be 16-byte-aligned.

```
v[0] := a3
v[1] := a2
v[2] := a1
v[3] := a0
```

```
__m128 _mm_move_ss(__m128 a, __m128 b)
```

Sets the low word to the SP FP value of *b*. The upper 3 SP FP values are passed through from *a*.

```
r0 := b0
r1 := a1
r2 := a2
r3 := a3
```

```
unsigned int _mm_getcsr(void)
```

Returns the contents of the control register.

```
void _mm_setcsr(unsigned int i)
```

Sets the control register to the value specified.

```
void _mm_prefetch(char const*a, int sel)
```

(uses PREFETCH) Loads one cache line of data from address *a* to a location "closer" to the processor. The value *sel* specifies the type of prefetch operation: the constants `_MM_HINT_T0`, `_MM_HINT_T1`, `_MM_HINT_T2`, and `_MM_HINT_NTA` should be used for IA-32, corresponding to the type of prefetch instruction. The constants `_MM_HINT_T1`, `_MM_HINT_NT1`, `_MM_HINT_NT2`, and `_MM_HINT_NTA` should be used for Itanium®-based systems.

```
void _mm_stream_pi(__m64 *p, __m64 a)
```

(uses MOVNTQ) Stores the data in *a* to the address *p* without polluting the caches. This intrinsic requires you to empty the multimedia state for the mmx register. See The EMMS Instruction: Why You Need It and When to Use It topic.

```
void _mm_stream_ps(float *p, __m128 a)
```

(see MOVNTPS) Stores the data in *a* to the address *p* without polluting the caches. The address must be 16-byte-aligned.

```
void _mm_sfence(void)
```

(uses SFENCE) Guarantees that every preceding store is globally visible before any subsequent store.

```
float _mm_cvtss_f32(__m128 a)
```

This intrinsic extracts a single precision floating point value from the first vector element of an `__m128`. It does so in the most efficient manner possible in the context used. This intrinsic doesn't map to any specific SSE instruction.

Miscellaneous Intrinsics Using Streaming SIMD Extensions

The prototypes for Streaming SIMD Extensions (SSE) intrinsics are in the `xmmintrin.h` header file.

Intrinsic Name	Operation	Corresponding Instruction
<code>_mm_shuffle_ps</code>	Shuffle	SHUFPS
<code>_mm_unpackhi_ps</code>	Unpack High	UNPCKHPS
<code>_mm_unpacklo_ps</code>	Unpack Low	UNPCKLPS
<code>_mm_loadh_pi</code>	Load High	MOVHPS <i>reg</i> , <i>mem</i>
<code>_mm_storeh_pi</code>	Store High	MOVHPS <i>mem</i> , <i>reg</i>
<code>_mm_movehl_ps</code>	Move High to Low	MOVHLPS
<code>_mm_movelh_ps</code>	Move Low to High	MOVLHPS
<code>_mm_loadl_pi</code>	Load Low	MOVLPS <i>reg</i> , <i>mem</i>
<code>_mm_storel_pi</code>	Store Low	MOVLPS <i>mem</i> , <i>reg</i>
<code>_mm_movemask_ps</code>	Create four-bit mask	MOVMSKPS

```
__m128 _mm_shuffle_ps(__m128 a, __m128 b, unsigned int imm8)
```

Selects four specific SP FP values from *a* and *b*, based on the mask *imm8*. The mask must be an immediate. See Macro Function for Shuffle Using Streaming SIMD Extensions for a description of the shuffle semantics.

```
__m128 _mm_unpackhi_ps(__m128 a, __m128 b)
```

Selects and interleaves the upper two SP FP values from *a* and *b*.

```
r0 := a2
r1 := b2
r2 := a3
r3 := b3
```

```
__m128 _mm_unpacklo_ps(__m128 a, __m128 b)
```

Selects and interleaves the lower two SP FP values from a and b.

```
r0 := a0
r1 := b0
r2 := a1
r3 := b1
```

```
__m128 _mm_loadh_pi(__m128, __m64 const *p)
```

Sets the upper two SP FP values with 64 bits of data loaded from the address p.

```
r0 := a0
r1 := a1
r2 := *p0
r3 := *p1
```

```
void _mm_storeh_pi(__m64 *p, __m128 a)
```

Stores the upper two SP FP values to the address p.

```
*p0 := a2
*p1 := a3
```

```
__m128 _mm_movehl_ps(__m128 a, __m128 b)
```

Moves the upper 2 SP FP values of b to the lower 2 SP FP values of the result. The upper 2 SP FP values of a are passed through to the result.

```
r3 := a3
r2 := a2
r1 := b3
r0 := b2
```

```
__m128 _mm_movelh_ps(__m128 a, __m128 b)
```

Moves the lower 2 SP FP values of b to the upper 2 SP FP values of the result. The lower 2 SP FP values of a are passed through to the result.

```
r3 := b1
r2 := b0
r1 := a1
r0 := a0
```

```
__m128 _mm_loadl_pi(__m128 a, __m64 const *p)
```

Sets the lower two SP FP values with 64 bits of data loaded from the address p; the upper two values are passed through from a.

```
r0 := *p0
r1 := *p1
r2 := a2
r3 := a3
```

```
void _mm_storel_pi(__m64 *p, __m128 a)
```

Stores the lower two SP FP values of a to the address p.

```
*p0 := a0
*p1 := a1
```

```
int __mm_movemask_ps(__m128 a)
```

Creates a 4-bit mask from the most significant bits of the four SP FP values.

```
r := sign(a3)<<3 | sign(a2)<<2 | sign(a1)<<1 | sign(a0)
```

Using Streaming SIMD Extensions on Itanium® Architecture

The Streaming SIMD Extensions (SSE) intrinsics provide access to Itanium® instructions for Streaming SIMD Extensions. To provide source compatibility with the IA-32 architecture, these intrinsics are equivalent both in name and functionality to the set of IA-32-based SSE intrinsics.

To write programs with the intrinsics, you should be familiar with the hardware features provided by SSE. Keep the following issues in mind:

- Certain intrinsics are provided only for compatibility with previously-defined IA-32 intrinsics. Using them on Itanium-based systems probably leads to performance degradation.
- Floating-point (FP) data loaded stored as `__m128` objects must be 16-byte-aligned.
- Some intrinsics require that their arguments be immediates -- that is, constant integers (literals), due to the nature of the instruction.

Data Types

The new data type `__m128` is used with the SSE intrinsics. It represents a 128-bit quantity composed of four single-precision FP values. This corresponds to the 128-bit IA-32 Streaming SIMD Extensions register.

The compiler aligns `__m128` local data to 16-byte boundaries on the stack. Global data of these types is also 16 byte-aligned. To align `integer`, `float`, or `double` arrays, you can use the `declspec` alignment.

Because Itanium instructions treat the SSE registers in the same way whether you are using packed or scalar data, there is no `__m32` data type to represent scalar data. For scalar operations, use the `__m128` objects and the "scalar" forms of the intrinsics; the compiler and the processor implement these operations with 32-bit memory references. But, for better performance the packed form should be substituting for the scalar form whenever possible.

The address of a `__m128` object may be taken.

For more information, see Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference Manual, Intel Corporation, doc. number 243191.

Implementation on Itanium-based systems

SSE intrinsics are defined for the `__m128` data type, a 128-bit quantity consisting of four single-precision FP values. SIMD instructions for Itanium-based systems operate on 64-bit FP register quantities containing two single-precision floating-point values. Thus, each `__m128` operand is actually a pair of FP registers and therefore each intrinsic corresponds to at least one pair of Itanium instructions operating on the pair of FP register operands.

Compatibility versus Performance

Many of the SSE intrinsics for Itanium-based systems were created for compatibility with existing IA-32 intrinsics and not for performance. In some situations, intrinsic usage that improved performance on IA-32 will not do so on Itanium-based systems. One reason for this is that some intrinsics map nicely into the IA-32 instruction set but not into the Itanium instruction set. Thus, it is important to differentiate between intrinsics which were implemented for a performance advantage on Itanium-based systems, and those implemented simply to provide compatibility with existing IA-32 code.

The following intrinsics are likely to reduce performance and should only be used to initially port legacy code or in non-critical code sections:

- Any SSE scalar intrinsic (`_ss variety`) - use packed (`_ps`) version if possible
- `comi` and `ucomi` SSE comparisons - these correspond to IA-32 `COMISS` and `UCOMISS` instructions only. A sequence of Itanium instructions are required to implement these.
- Conversions in general are multi-instruction operations. These are particularly expensive:
`_mm_cvtpi16_ps`, `_mm_cvtpul6_ps`, `_mm_cvtpi8_ps`, `_mm_cvtpu8_ps`,
`_mm_cvtpi32x2_ps`, `_mm_cvtps_pi16`, `_mm_cvtps_pi8`
- SSE utility intrinsic `_mm_movemask_ps`

If the inaccuracy is acceptable, the SIMD reciprocal and reciprocal square root approximation intrinsics (`rcp` and `rsqrt`) are much faster than the true `div` and `sqrt` intrinsics.

Macro Function for Shuffle Using Streaming SIMD Extensions

The Streaming SIMD Extensions (SSE) provide a macro function to help create constants that describe shuffle operations. The macro takes four small integers (in the range of 0 to 3) and combines them into an 8-bit immediate value used by the `SHUFPD` instruction.

Shuffle Function Macro

```
_MM_SHUFFLE(z,y,x,w)
/* expands to the following value */
(z<<6) | (y<<4) | (x<<2) | w
```

You can view the four integers as selectors for choosing which two words from the first input operand and which two words from the second are to be put into the result word.

View of Original and Result Words with Shuffle Function Macro

```
      127      0
; m1 = [ a | b | c | d ]
      127      0
; m2 = [ e | f | g | h ]
m3 = _mm_shuffle_ps(m1, m2,
  _MM_SHUFFLE(1,0,3,2))
      127      0
; m3 = [ g | h | a | b ]
```

Macro Functions to Read and Write the Control Registers

The following macro functions enable you to read and write bits to and from the control register. For details, see Set Operations. For Itanium®-based systems, these macros do not allow you to access all of the bits of the FPSR. See the descriptions for the `getfpsr()` and `setfpsr()` intrinsics in the Native Intrinsics for Itanium Instructions topic.

Exception State Macros	Macro Arguments
<code>_MM_SET_EXCEPTION_STATE(x)</code>	<code>_MM_EXCEPT_INVALID</code>
<code>_MM_GET_EXCEPTION_STATE()</code>	<code>_MM_EXCEPT_DIV_ZERO</code>
	<code>_MM_EXCEPT_DENORM</code>
Macro Definitions Write to and read from the sixth-least significant control register bit, respectively.	<code>_MM_EXCEPT_OVERFLOW</code>
	<code>_MM_EXCEPT_UNDERFLOW</code>
	<code>_MM_EXCEPT_INEXACT</code>

The following example tests for a divide-by-zero exception.

Exception State Macros with `_MM_EXCEPT_DIV_ZERO`

```
if (_MM_GET_EXCEPTION_STATE(x) & _MM_EXCEPT_DIV_ZERO) {
    /* Exception has occurred */
}
```

Exception Mask Macros	Macro Arguments
<code>_MM_SET_EXCEPTION_MASK(x)</code>	<code>_MM_MASK_INVALID</code>
<code>_MM_GET_EXCEPTION_MASK()</code>	<code>_MM_MASK_DIV_ZERO</code>
	<code>_MM_MASK_DENORM</code>
Macro Definitions Write to and read from the seventh through twelfth control register bits, respectively. Note: All six exception mask bits are always affected. Bits not set explicitly are cleared.	<code>_MM_MASK_OVERFLOW</code>
	<code>_MM_MASK_UNDERFLOW</code>
	<code>_MM_MASK_INEXACT</code>

The following example masks the overflow and underflow exceptions and unmask all other exceptions.

Exception Mask with <code>_MM_MASK_OVERFLOW</code> and <code>_MM_MASK_UNDERFLOW</code>	
<code>_MM_SET_EXCEPTION_MASK(MM_MASK_OVERFLOW _MM_MASK_UNDERFLOW)</code>	
Rounding Mode	Macro Arguments
<code>_MM_SET_ROUNDING_MODE(x)</code>	<code>_MM_ROUND_NEAREST</code>
<code>_MM_GET_ROUNDING_MODE()</code>	<code>_MM_ROUND_DOWN</code>
Macro Definition Write to and read from bits thirteen and fourteen of the control register.	<code>_MM_ROUND_UP</code>
	<code>_MM_ROUND_TOWARD_ZERO</code>

The following example tests the rounding mode for round toward zero.

Rounding Mode with <code>_MM_ROUND_TOWARD_ZERO</code>	
<pre>if (_MM_GET_ROUNDING_MODE() == _MM_ROUND_TOWARD_ZERO) { /* Rounding mode is round toward zero */ }</pre>	
Flush-to-Zero Mode	Macro Arguments
<code>_MM_SET_FLUSH_ZERO_MODE(x)</code>	<code>_MM_FLUSH_ZERO_ON</code>
<code>_MM_GET_FLUSH_ZERO_MODE()</code>	<code>_MM_FLUSH_ZERO_OFF</code>
Macro Definition Write to and read from bit fifteen of the control register.	

The following example disables flush-to-zero mode.

Flush-to-Zero Mode with <code>_MM_FLUSH_ZERO_OFF</code>
<code>_MM_SET_FLUSH_ZERO_MODE(_MM_FLUSH_ZERO_OFF)</code>

Macro Function for Matrix Transposition

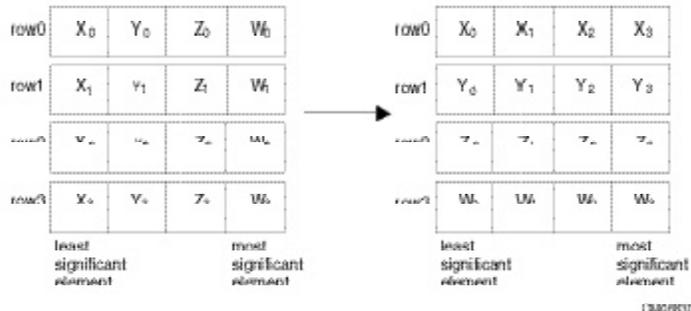
The Streaming SIMD Extensions (SSE) also provide the following macro function to transpose a 4 by 4 matrix of single precision floating point values.

```
_MM_TRANSPOSE4_PS(row0, row1, row2, row3)
```

The arguments `row0`, `row1`, `row2`, and `row3` are `__m128` values whose elements form the corresponding rows of a 4 by 4 matrix. The matrix transposition is returned in arguments `row0`, `row1`, `row2`, and `row3` where `row0` now holds column 0 of the original matrix, `row1` now holds column 1 of the original matrix, and so on.

The transposition function of this macro is illustrated in the "Matrix Transposition Using the `_MM_TRANSPOSE4_PS`" figure.

Matrix Transposition Using `_MM_TRANSPOSE4_PS` Macro



Streaming SIMD Extensions 2

This section describes the C++ language-level features supporting the Intel® Pentium® 4 processor Streaming SIMD Extensions 2 (SSE2) in the Intel® C++ Compiler, which are divided into two categories:

- Floating-Point Intrinsics -- describes the arithmetic, logical, compare, conversion, memory, and initialization intrinsics for the double-precision floating-point data type (`__m128d`).
- Integer Intrinsics -- describes the arithmetic, logical, compare, conversion, memory, and initialization intrinsics for the extended-precision integer data type (`__m128i`).



Note

The Pentium 4 processor SSE2 intrinsics are defined only for IA-32 platforms, not Itanium®-based platforms. Pentium 4 processor SSE2 operate on 128 bit quantities -- 2 64-bit double precision floating point values. The Itanium processor does not support parallel double precision computation, so Pentium 4 processor SSE2 are not implemented on Itanium-based systems.

For more details, refer to the *Pentium® 4 processor Streaming SIMD Extensions 2 External Architecture Specification (EAS)* and other Pentium 4 processor manuals available for download from the developer.intel.com web site. You should be familiar with the hardware features provided by the SSE2 when writing programs with the intrinsics. The following are three important issues to keep in mind:

- Certain intrinsics, such as `_mm_loadr_pd` and `_mm_cmpgt_sd`, are not directly supported by the instruction set. While these intrinsics are convenient programming aids, be mindful of their implementation cost.
- Data loaded or stored as `__m128d` objects must be generally 16-byte-aligned.
- Some intrinsics require that their argument be immediates, that is, constant integers (literals), due to the nature of the instruction.

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

**Note**

You can also use the single `ia32intrin.h` header file for any IA-32 intrinsics.

Floating-point Arithmetic Operations for Streaming SIMD Extensions 2

The arithmetic operations for the Streaming SIMD Extensions 2 (SSE2) are listed in the following table. The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

Intrinsic Name	Corresponding Instruction	Operation	R0 Value	R1 Value
<code>_mm_add_sd</code>	ADDSD	Addition	a0 [op] b0	a1
<code>_mm_add_pd</code>	ADDPD	Addition	a0 [op] b0	a1 [op] b1
<code>_mm_sub_sd</code>	SUBSD	Subtraction	a0 [op] b0	a1
<code>_mm_sub_pd</code>	SUBPD	Subtraction	a0 [op] b0	a1 [op] b1
<code>_mm_mul_sd</code>	MULSD	Multiplication	a0 [op] b0	a1
<code>_mm_mul_pd</code>	MULPD	Multiplication	a0 [op] b0	a1 [op] b1
<code>_mm_div_sd</code>	DIVSD	Division	a0 [op] b0	a1
<code>_mm_div_pd</code>	DIVPD	Division	a0 [op] b0	a1 [op] b1
<code>_mm_sqrt_sd</code>	SQRTSD	Computes Square Root	a0 [op] b0	a1
<code>_mm_sqrt_pd</code>	SQRTPD	Computes Square Root	a0 [op] b0	a1 [op] b1
<code>_mm_min_sd</code>	MINSD	Computes Minimum	a0 [op] b0	a1
<code>_mm_min_pd</code>	MINPD	Computes Minimum	a0 [op] b0	a1 [op] b1
<code>_mm_max_sd</code>	MAXSD	Computes Maximum	a0 [op] b0	a1
<code>_mm_max_pd</code>	MAXPD	Computes Maximum	a0 [op] b0	a1 [op] b1

```
__m128d __mm_add_sd(__m128d a, __m128d b)
```

Adds the lower DP FP (double-precision, floating-point) values of a and b ; the upper DP FP value is passed through from a.

```
r0 := a0 + b0
r1 := a1
```

```
__m128d __mm_add_pd(__m128d a, __m128d b)
```

Adds the two DP FP values of a and b.

```
r0 := a0 + b0
r1 := a1 + b1
```

```
__m128d __mm_sub_sd(__m128d a, __m128d b)
```

Subtracts the lower DP FP value of b from a. The upper DP FP value is passed through from a.

```
r0 := a0 - b0
r1 := a1
```

```
__m128d __mm_sub_pd(__m128d a, __m128d b)
```

Subtracts the two DP FP values of b from a.

```
r0 := a0 - b0
r1 := a1 - b1
```

```
__m128d __mm_mul_sd(__m128d a, __m128d b)
```

Multiplies the lower DP FP values of a and b. The upper DP FP is passed through from a.

```
r0 := a0 * b0
r1 := a1
```

```
__m128d __mm_mul_pd(__m128d a, __m128d b)
```

Multiplies the two DP FP values of a and b.

```
r0 := a0 * b0
r1 := a1 * b1
```

```
__m128d __mm_div_sd(__m128d a, __m128d b)
```

Divides the lower DP FP values of a and b. The upper DP FP value is passed through from a.

```
r0 := a0 / b0
r1 := a1
```

```
__m128d __mm_div_pd(__m128d a, __m128d b)
```

Divides the two DP FP values of a and b.

```
r0 := a0 / b0
r1 := a1 / b1
```

```
__m128d __mm_sqrt_sd(__m128d a, __m128d b)
```

Computes the square root of the lower DP FP value of b. The upper DP FP value is passed through from a.

```
r0 := sqrt(b0)
r1 := a1
```

```
__m128d _mm_sqrt_pd(__m128d a)
```

Computes the square roots of the two DP FP values of a.

```
r0 := sqrt(a0)
r1 := sqrt(a1)
```

```
__m128d _mm_min_sd(__m128d a, __m128d b)
```

Computes the minimum of the lower DP FP values of a and b. The upper DP FP value is passed through from a.

```
r0 := min (a0, b0)
r1 := a1
```

```
__m128d _mm_min_pd(__m128d a, __m128d b)
```

Computes the minima of the two DP FP values of a and b.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
```

```
__m128d _mm_max_sd(__m128d a, __m128d b)
```

Computes the maximum of the lower DP FP values of a and b. The upper DP FP value is passed through from a.

```
r0 := max (a0, b0)
r1 := a1
```

```
__m128d _mm_max_pd(__m128d a, __m128d b)
```

Computes the maxima of the two DP FP values of a and b.

```
r0 := max(a0, b0)
r1 := max(a1, b1)
```

Floating-point Logical Operations for Streaming SIMD Extensions 2

The prototypes for Streaming SIMD Extensions 2 (SSE2) intrinsics are in the `emmintrin.h` header file.

```
__m128d _mm_and_pd(__m128d a, __m128d b)
```

(uses `ANDPD`) Computes the bitwise AND of the two DP FP values of a and b.

```
r0 := a0 & b0
r1 := a1 & b1
```

```
__m128d _mm_andnot_pd(__m128d a, __m128d b)
```

(uses `ANDNPD`) Computes the bitwise AND of the 128-bit value in b and the bitwise NOT of the 128-bit value in a.

```
r0 := (~a0) & b0
r1 := (~a1) & b1
```

```
__m128d _mm_or_pd(__m128d a, __m128d b)
```

(uses `ORPD`) Computes the bitwise OR of the two DP FP values of a and b.

```
r0 := a0 | b0
r1 := a1 | b1
```

```
__m128d _mm_xor_pd(__m128d a, __m128d b)
```

(uses XORPD) Computes the bitwise XOR of the two DP FP values of a and b.

```
r0 := a0 ^ b0
```

```
r1 := a1 ^ b1
```

Floating-point Comparison Operations for Streaming SIMD Extensions 2

Each comparison intrinsic performs a comparison of a and b. For the packed form, the two DP FP values of a and b are compared, and a 128-bit mask is returned. For the scalar form, the lower DP FP values of a and b are compared, and a 64-bit mask is returned; the upper DP FP value is passed through from a. The mask is set to 0xffffffffffffffff for each element where the comparison is true and 0x0 where the comparison is false. The *r* following the instruction name indicates that the operands to the instruction are reversed in the actual implementation. The comparison intrinsics for the Streaming SIMD Extensions 2 (SSE2) are listed in the following table followed by detailed descriptions.

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

Intrinsic Name	Corresponding Instruction	Compare For:
<code>_mm_cmpeq_pd</code>	CMPEQPD	Equality
<code>_mm_cmplt_pd</code>	CMPLTPD	Less Than
<code>_mm_cmple_pd</code>	CMPLEPD	Less Than or Equal
<code>_mm_cmpgt_pd</code>	CMPLTPDr	Greater Than
<code>_mm_cmpge_pd</code>	CMPLEPDr	Greater Than or Equal
<code>_mm_cmpord_pd</code>	CMPOrdPD	Ordered
<code>_mm_cmpunord_pd</code>	CMPUNORDPD	Unordered
<code>_mm_cmpneq_pd</code>	CMPNEQPD	Inequality
<code>_mm_cmpnlt_pd</code>	CMPNLTPD	Not Less Than
<code>_mm_cmpnle_pd</code>	CMPNLEPD	Not Less Than or Equal
<code>_mm_cmpngt_pd</code>	CMPNLTPDr	Not Greater Than
<code>_mm_cmpnge_pd</code>	CMPLEPDr	Not Greater Than or Equal
<code>_mm_cmpeq_sd</code>	CMPEQSD	Equality
<code>_mm_cmplt_sd</code>	CMPLTSD	Less Than
<code>_mm_cmple_sd</code>	CMPLESD	Less Than or Equal
<code>_mm_cmpgt_sd</code>	CMPLTSDr	Greater Than
<code>_mm_cmpge_sd</code>	CMPLESDr	Greater Than or Equal
<code>_mm_cmpord_sd</code>	CMPOrdsD	Ordered
<code>_mm_cmpunord_sd</code>	CMPUNORDSD	Unordered

Intrinsic Name	Corresponding Instruction	Compare For:
<code>_mm_cmpneq_sd</code>	CMPNEQSD	Inequality
<code>_mm_cmpnlt_sd</code>	CMPNLTSD	Not Less Than
<code>_mm_cmpnle_sd</code>	CMPNLESD	Not Less Than or Equal
<code>_mm_cmpngt_sd</code>	CMPNLTSDr	Not Greater Than
<code>_mm_cmpnge_sd</code>	CMPNLESDR	Not Greater Than or Equal
<code>_mm_comieq_sd</code>	COMISD	Equality
<code>_mm_comilt_sd</code>	COMISD	Less Than
<code>_mm_comile_sd</code>	COMISD	Less Than or Equal
<code>_mm_comigt_sd</code>	COMISD	Greater Than
<code>_mm_comige_sd</code>	COMISD	Greater Than or Equal
<code>_mm_comineq_sd</code>	COMISD	Not Equal
<code>_mm_ucomieq_sd</code>	UCOMISD	Equality
<code>_mm_ucomilt_sd</code>	UCOMISD	Less Than
<code>_mm_ucomile_sd</code>	UCOMISD	Less Than or Equal
<code>_mm_ucomigt_sd</code>	UCOMISD	Greater Than
<code>_mm_ucomige_sd</code>	UCOMISD	Greater Than or Equal
<code>_mm_ucomineq_sd</code>	UCOMISD	Not Equal

```
__m128d _mm_cmpeq_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for equality.

```
r0 := (a0 == b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 == b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmplt_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a less than b.

```
r0 := (a0 < b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 < b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmple_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a less than or equal to b.

```
r0 := (a0 <= b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 <= b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpgt_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a greater than b.

```
r0 := (a0 > b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 > b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpge_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a greater than or equal to b.

```
r0 := (a0 >= b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 >= b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpord_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for ordered.

```
r0 := (a0 ord b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 ord b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpunord_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for unordered.

```
r0 := (a0 unord b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 unord b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpneq_pd ( __m128d a, __m128d b)
```

Compares the two DP FP values of a and b for inequality.

```
r0 := (a0 != b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := (a1 != b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpnlt_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a not less than b.

```
r0 := !(a0 < b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := !(a1 < b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpnle_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a not less than or equal to b.

```
r0 := !(a0 <= b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := !(a1 <= b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpngt_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a not greater than b.

```
r0 := !(a0 > b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := !(a1 > b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d _mm_cmpnge_pd(__m128d a, __m128d b)
```

Compares the two DP FP values of a and b for a not greater than or equal to b.

```
r0 := !(a0 >= b0) ? 0xffffffffffffffff : 0x0
```

```
r1 := !(a1 >= b1) ? 0xffffffffffffffff : 0x0
```

```
__m128d __mm_cmpeq_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for equality. The upper DP FP value is passed through from a.
 r0 := (a0 == b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmplt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a less than b. The upper DP FP value is passed through from a.
 r0 := (a0 < b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmple_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a less than or equal to b. The upper DP FP value is passed through from a.
 r0 := (a0 <= b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmpgt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a greater than b. The upper DP FP value is passed through from a.
 r0 := (a0 > b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmpge_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a greater than or equal to b. The upper DP FP value is passed through from a.
 r0 := (a0 >= b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmpord_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for ordered. The upper DP FP value is passed through from a.
 r0 := (a0 ord b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmpunord_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for unordered. The upper DP FP value is passed through from a.
 r0 := (a0 unord b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmpneq_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for inequality. The upper DP FP value is passed through from a.
 r0 := (a0 != b0) ? 0xffffffffffffffff : 0x0
 r1 := a1

```
__m128d __mm_cmpnlt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a not less than b. The upper DP FP value is passed through from a.

```
r0 := !(a0 < b0) ? 0xffffffffffffffff : 0x0
r1 := a1
```

```
__m128d __mm_cmpnle_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a not less than or equal to b. The upper DP FP value is passed through from a.

```
r0 := !(a0 <= b0) ? 0xffffffffffffffff : 0x0
r1 := a1
```

```
__m128d __mm_cmpngt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a not greater than b. The upper DP FP value is passed through from a.

```
r0 := !(a0 > b0) ? 0xffffffffffffffff : 0x0
r1 := a1
```

```
__m128d __mm_cmpnge_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a not greater than or equal to b. The upper DP FP value is passed through from a.

```
r0 := !(a0 >= b0) ? 0xffffffffffffffff : 0x0
r1 := a1
```

```
int __mm_comieq_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a equal to b. If a and b are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 == b0) ? 0x1 : 0x0
```

```
int __mm_comilt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a less than b. If a is less than b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 < b0) ? 0x1 : 0x0
```

```
int __mm_comile_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a less than or equal to b. If a is less than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 <= b0) ? 0x1 : 0x0
```

```
int __mm_comigt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a greater than b. If a is greater than b are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 > b0) ? 0x1 : 0x0
```

```
int _mm_comige_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a greater than or equal to b. If a is greater than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 >= b0) ? 0x1 : 0x0
```

```
int _mm_comineq_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a not equal to b. If a and b are not equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 != b0) ? 0x1 : 0x0
```

```
int _mm_ucomieq_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a equal to b. If a and b are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 == b0) ? 0x1 : 0x0
```

```
int _mm_ucomilt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a less than b. If a is less than b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 < b0) ? 0x1 : 0x0
```

```
int _mm_ucomile_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a less than or equal to b. If a is less than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 <= b0) ? 0x1 : 0x0
```

```
int _mm_ucomigt_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a greater than b. If a is greater than b are equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 > b0) ? 0x1 : 0x0
```

```
int _mm_ucomige_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a greater than or equal to b. If a is greater than or equal to b, 1 is returned. Otherwise 0 is returned.

```
r := (a0 >= b0) ? 0x1 : 0x0
```

```
int _mm_ucomineq_sd(__m128d a, __m128d b)
```

Compares the lower DP FP value of a and b for a not equal to b. If a and b are not equal, 1 is returned. Otherwise 0 is returned.

```
r := (a0 != b0) ? 0x1 : 0x0
```

Floating-point Conversion Operations for Streaming SIMD Extensions 2

Each conversion intrinsic takes one data type and performs a conversion to a different type. Some conversions such as `_mm_cvtpd_ps` result in a loss of precision. The rounding mode used in such cases is determined by the value in the MXCSR register. The default rounding mode is round-to-nearest. Note that the rounding mode used by the C and C++ languages when performing a type conversion is to truncate.

The `_mm_cvttpd_epi32` and `_mm_cvttssd_si32` intrinsics use the truncate rounding mode regardless of the mode specified by the MXCSR register.

The conversion-operation intrinsics for Streaming SIMD Extensions 2 (SSE2) are listed in the following table followed by detailed descriptions.

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

Intrinsic Name	Corresponding Instruction	Return Type	Parameters
<code>_mm_cvtpd_ps</code>	CVTPD2PS	<code>__m128</code>	(<code>__m128d a</code>)
<code>_mm_cvtps_pd</code>	CVTPS2PD	<code>__m128d</code>	(<code>__m128 a</code>)
<code>_mm_cvtepi32_pd</code>	CVTDQ2PD	<code>__m128d</code>	(<code>__m128i a</code>)
<code>_mm_cvtpd_epi32</code>	CVTPD2DQ	<code>__m128i</code>	(<code>__m128d a</code>)
<code>_mm_cvtsd_si32</code>	CVTSD2SI	<code>int</code>	(<code>__m128d a</code>)
<code>_mm_cvtsd_ss</code>	CVTSD2SS	<code>__m128</code>	(<code>__m128 a</code> , <code>__m128d b</code>)
<code>_mm_cvtsi32_sd</code>	CVTSI2SD	<code>__m128d</code>	(<code>__m128d a</code> , <code>int b</code>)
<code>_mm_cvtss_sd</code>	CVTSS2SD	<code>__m128d</code>	(<code>__m128d a</code> , <code>__m128 b</code>)
<code>_mm_cvttpd_epi32</code>	CVTTPD2DQ	<code>__m128i</code>	(<code>__m128d a</code>)
<code>_mm_cvttssd_si32</code>	CVTTSD2SI	<code>int</code>	(<code>__m128d a</code>)
<code>_mm_cvtpd_pi32</code>	CVTPD2PI	<code>__m64</code>	(<code>__m128d a</code>)
<code>_mm_cvttpd_pi32</code>	CVTTPD2PI	<code>__m64</code>	(<code>__m128d a</code>)
<code>_mm_cvtpi32_pd</code>	CVTPI2PD	<code>__m128d</code>	(<code>__m64 a</code>)
<code>_mm_cvtsd_f64</code>	None	<code>double</code>	(<code>__m128d a</code>)

```
__m128 _mm_cvtpd_ps(__m128d a)
```

Converts the two DP FP values of `a` to SP FP values.

```
r0 := (float) a0
r1 := (float) a1
r2 := 0.0 ; r3 := 0.0
```

```
__m128d _mm_cvtps_pd(__m128 a)
```

Converts the lower two SP FP values of `a` to DP FP values.

```
r0 := (double) a0
r1 := (double) a1
```

```
__m128d _mm_cvtepi32_pd(__m128i a)
```

Converts the lower two signed 32-bit integer values of `a` to DP FP values.

```
r0 := (double) a0
r1 := (double) a1
```

```
__m128i _mm_cvtpd_epi32(__m128d a)
```

Converts the two DP FP values of a to 32-bit signed integer values.

```
r0 := (int) a0
r1 := (int) a1
r2 := 0x0 ; r3 := 0x0
```

```
int _mm_cvtsd_si32(__m128d a)
```

Converts the lower DP FP value of a to a 32-bit signed integer value.

```
r := (int) a0
```

```
__m128 _mm_cvtsd_ss(__m128 a, __m128d b)
```

Converts the lower DP FP value of b to an SP FP value. The upper SP FP values in a are passed through.

```
r0 := (float) b0
r1 := a1; r2 := a2 ; r3 := a3
```

```
__m128d _mm_cvtsi32_sd(__m128d a, int b)
```

Converts the signed integer value in b to a DP FP value. The upper DP FP value in a is passed through.

```
r0 := (double) b
r1 := a1
```

```
__m128d _mm_cvtss_sd(__m128d a, __m128 b)
```

Converts the lower SP FP value of b to a DP FP value. The upper value DP FP value in a is passed through.

```
r0 := (double) b0
r1 := a1
```

```
__m128i _mm_cvttpd_epi32(__m128d a)
```

Converts the two DP FP values of a to 32-bit signed integers using truncate.

```
r0 := (int) a0
r1 := (int) a1
r2 := 0x0 ; r3 := 0x0
```

```
int _mm_cvtttsd_si32(__m128d a)
```

Converts the lower DP FP value of a to a 32-bit signed integer using truncate.

```
r := (int) a0
```

```
__m64 _mm_cvtpd_pi32(__m128d a)
```

Converts the two DP FP values of a to 32-bit signed integer values.

```
r0 := (int) a0
r1 := (int) a1
```

```
__m64 _mm_cvttpd_pi32(__m128d a)
```

Converts the two DP FP values of a to 32-bit signed integer values using truncate.

```
r0 := (int) a0
r1 := (int) a1
```

```
__m128d _mm_cvtpi32_pd(__m64 a)
```

Converts the two 32-bit signed integer values of `a` to DP FP values.

```
r0 := (double) a0
r1 := (double) a1
```

```
_mm_cvtsd_f64(__m128d a)
```

This intrinsic extracts a double precision floating point value from the first vector element of an `__m128d`. It does so in the most efficient manner possible in the context used. This intrinsic does not map to any specific SSE2 instruction.

Floating-point Memory and Initialization Operations for Streaming SIMD Extensions 2

This section describes the `load`, `set`, and `store` operations, which let you load and store data into memory. The `load` and `set` operations are similar in that both initialize `__m128d` data. However, the `set` operations take a double argument and are intended for initialization with constants, while the `load` operations take a double pointer argument and are intended to mimic the instructions for loading data from memory. The `store` operation assigns the initialized data to the address.



Note

There is no intrinsic for move operations. To move data from one register to another, a simple assignment, `A = B`, suffices, where `A` and `B` are the source and target registers for the move operation.

The prototypes for Streaming SIMD Extensions 2 (SSE2) intrinsics are in the `emmintrin.h` header file.

Floating-point Load Operations for Streaming SIMD Extensions 2

The following `load` operation intrinsics and their respective instructions are functional in the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
__m128d _mm_load_pd(double const*dp)
```

(uses `MOVAPD`) Loads two DP FP values. The address `p` must be 16-byte aligned.

```
r0 := p[0]
r1 := p[1]
```

```
__m128d _mm_load1_pd(double const*dp)
```

(uses `MOVSD` + shuffling) Loads a single DP FP value, copying to both elements. The address `p` need not be 16-byte aligned.

```
r0 := *p
r1 := *p
```

```
__m128d _mm_loadr_pd(double const*dp)
```

(uses MOVAPD + shuffling) Loads two DP FP values in reverse order. The address `p` must be 16-byte aligned.

```
r0 := p[1]
r1 := p[0]
```

```
__m128d _mm_loadu_pd(double const*dp)
```

(uses MOVUPD) Loads two DP FP values. The address `p` need not be 16-byte aligned.

```
r0 := p[0]
r1 := p[1]
```

```
__m128d _mm_load_sd(double const*dp)
```

(uses MOVSD) Loads a DP FP value. The upper DP FP is set to zero. The address `p` need not be 16-byte aligned.

```
r0 := *p
r1 := 0.0
```

```
__m128d _mm_loadh_pd(__m128d a, double const*dp)
```

(uses MOVHPD) Loads a DP FP value as the upper DP FP value of the result. The lower DP FP value is passed through from `a`. The address `p` need not be 16-byte aligned.

```
r0 := a0
r1 := *p
```

```
__m128d _mm_loadl_pd(__m128d a, double const*dp)
```

(uses MOVLPD) Loads a DP FP value as the lower DP FP value of the result. The upper DP FP value is passed through from `a`. The address `p` need not be 16-byte aligned.

```
r0 := *p
r1 := a1
```

Floating-point Set Operations for Streaming SIMD Extensions 2

The following `set` operation intrinsics and their respective instructions are functional in the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
__m128d _mm_set_sd(double w)
```

(composite) Sets the lower DP FP value to `w` and sets the upper DP FP value to zero.

```
r0 := w
r1 := 0.0
```

```
__m128d _mm_set1_pd(double w)
```

(composite) Sets the 2 DP FP values to `w`.

```
r0 := w
r1 := w
```

```
__m128d _mm_set_pd(double w, double x)
```

(composite) Sets the lower DP FP value to *x* and sets the upper DP FP value to *w*.

```
r0 := x
r1 := w
```

```
__m128d _mm_setr_pd(double w, double x)
```

(composite) Sets the lower DP FP value to *w* and sets the upper DP FP value to *x*.

```
r0 := w
r1 := x
```

```
__m128d _mm_setzero_pd(void)
```

(uses XORPD) Sets the 2 DP FP values to zero.

```
r0 := 0.0
r1 := 0.0
```

```
__m128d _mm_move_sd( __m128d a, __m128d b)
```

(uses MOVSD) Sets the lower DP FP value to the lower DP FP value of *b*. The upper DP FP value is passed through from *a*.

```
r0 := b0
r1 := a1
```

Floating-point Store Operations for Streaming SIMD Extensions 2

The following store operation intrinsics and their respective instructions are functional in the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
void _mm_store_sd(double *dp, __m128d a)
```

(uses MOVSD) Stores the lower DP FP value of *a*. The address *dp* need not be 16-byte aligned.

```
*dp := a0
```

```
void _mm_storel_pd(double *dp, __m128d a)
```

(uses MOVAPD + shuffling) Stores the lower DP FP value of *a* twice. The address *dp* must be 16-byte aligned.

```
dp[0] := a0
dp[1] := a0
```

```
void _mm_store_pd(double *dp, __m128d a)
```

(uses MOVAPD) Stores two DP FP values. The address *dp* must be 16-byte aligned.

```
dp[0] := a0
dp[1] := a1
```

```
void _mm_storeu_pd(double *dp, __m128d a)
```

(uses MOVUPD) Stores two DP FP values. The address *dp* need not be 16-byte aligned.

```
dp[0] := a0
dp[1] := a1
```

```
void _mm_storer_pd(double *dp, __m128d a)
```

(uses MOVAPD + shuffling) Stores two DP FP values in reverse order. The address `dp` must be 16-byte aligned.

```
dp[0] := a1
dp[1] := a0
```

```
void _mm_storeh_pd(double *dp, __m128d a)
```

(uses MOVHPD) Stores the upper DP FP value of `a`.

```
*dp := a1
```

```
void _mm_storel_pd(double *dp, __m128d a)
```

(uses MOVLPD) Stores the lower DP FP value of `a`.

```
*dp := a0
```

Integer Arithmetic Operations for Streaming SIMD Extensions 2

The integer arithmetic operations for Streaming SIMD Extensions 2 (SSE2) are listed in the following table followed by their descriptions. The packed arithmetic intrinsics for SSE2 are listed in the Floating-point Arithmetic Operations topic.

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

Intrinsic	Instruction	Operation
<code>_mm_add_epi8</code>	PADDB	Addition
<code>_mm_add_epi16</code>	PADDW	Addition
<code>_mm_add_epi32</code>	PADDQ	Addition
<code>_mm_add_si64</code>	PADDQ	Addition
<code>_mm_add_epi64</code>	PADDQ	Addition
<code>_mm_adds_epi8</code>	PADDSB	Addition
<code>_mm_adds_epi16</code>	PADDSW	Addition
<code>_mm_adds_epu8</code>	PADDUSB	Addition
<code>_mm_adds_epu16</code>	PADDUSW	Addition
<code>_mm_avg_epu8</code>	PAVGB	Computes Average
<code>_mm_avg_epu16</code>	PAVGW	Computes Average
<code>_mm_madd_epi16</code>	PMADDWD	Multiplication/Addition
<code>_mm_max_epi16</code>	PMAXSW	Computes Maxima
<code>_mm_max_epu8</code>	PMAXUB	Computes Maxima
<code>_mm_min_epi16</code>	PMINSW	Computes Minima

Intrinsic	Instruction	Operation
<code>__mm_min_epu8</code>	PMINUB	Computes Minima
<code>__mm_mulhi_epi16</code>	PMULHW	Multiplication
<code>__mm_mulhi_epu16</code>	PMULHUW	Multiplication
<code>__mm_mullo_epi16</code>	PMULLW	Multiplication
<code>__mm_mul_su32</code>	PMULUDQ	Multiplication
<code>__mm_mul_epu32</code>	PMULUDQ	Multiplication
<code>__mm_sad_epu8</code>	PSADBW	Computes Difference/Adds
<code>__mm_sub_epi8</code>	PSUBB	Subtraction
<code>__mm_sub_epi16</code>	PSUBW	Subtraction
<code>__mm_sub_epi32</code>	PSUBD	Subtraction
<code>__mm_sub_si64</code>	PSUBQ	Subtraction
<code>__mm_sub_epi64</code>	PSUBQ	Subtraction
<code>__mm_subs_epi8</code>	PSUBSB	Subtraction
<code>__mm_subs_epi16</code>	PSUBSW	Subtraction
<code>__mm_subs_epu8</code>	PSUBUSB	Subtraction
<code>__mm_subs_epu16</code>	PSUBUSW	Subtraction

```
__mm128i __mm_add_epi8(__m128i a, __m128i b)
```

Adds the 16 signed or unsigned 8-bit integers in a to the 16 signed or unsigned 8-bit integers in b.

```
r0 := a0 + b0
r1 := a1 + b1
...
r15 := a15 + b15
```

```
__mm128i __mm_add_epi16(__m128i a, __m128i b)
```

Adds the 8 signed or unsigned 16-bit integers in a to the 8 signed or unsigned 16-bit integers in b.

```
r0 := a0 + b0
r1 := a1 + b1
...
r7 := a7 + b7
```

```
__m128i __mm_add_epi32(__m128i a, __m128i b)
```

Adds the 4 signed or unsigned 32-bit integers in a to the 4 signed or unsigned 32-bit integers in b.

```
r0 := a0 + b0
r1 := a1 + b1
r2 := a2 + b2
r3 := a3 + b3
```

```
__m64 _mm_add_si64(__m64 a, __m64 b)
```

Adds the signed or unsigned 64-bit integer a to the signed or unsigned 64-bit integer b.
 $r := a + b$

```
__m128i _mm_add_epi64(__m128i a, __m128i b)
```

Adds the 2 signed or unsigned 64-bit integers in a to the 2 signed or unsigned 64-bit integers in b.
 $r0 := a0 + b0$
 $r1 := a1 + b1$

```
__m128i _mm_adds_epi8(__m128i a, __m128i b)
```

Adds the 16 signed 8-bit integers in a to the 16 signed 8-bit integers in b using saturating arithmetic.
 $r0 := \text{SignedSaturate}(a0 + b0)$
 $r1 := \text{SignedSaturate}(a1 + b1)$
 \dots
 $r15 := \text{SignedSaturate}(a15 + b15)$

```
__m128i _mm_adds_epi16(__m128i a, __m128i b)
```

Adds the 8 signed 16-bit integers in a to the 8 signed 16-bit integers in b using saturating arithmetic.
 $r0 := \text{SignedSaturate}(a0 + b0)$
 $r1 := \text{SignedSaturate}(a1 + b1)$
 \dots
 $r7 := \text{SignedSaturate}(a7 + b7)$

```
__m128i _mm_adds_epu8(__m128i a, __m128i b)
```

Adds the 16 unsigned 8-bit integers in a to the 16 unsigned 8-bit integers in b using saturating arithmetic.
 $r0 := \text{UnsignedSaturate}(a0 + b0)$
 $r1 := \text{UnsignedSaturate}(a1 + b1)$
 \dots
 $r15 := \text{UnsignedSaturate}(a15 + b15)$

```
__m128i _mm_adds_epu16(__m128i a, __m128i b)
```

Adds the 8 unsigned 16-bit integers in a to the 8 unsigned 16-bit integers in b using saturating arithmetic.
 $r0 := \text{UnsignedSaturate}(a0 + b0)$
 $r1 := \text{UnsignedSaturate}(a1 + b1)$
 \dots
 $r15 := \text{UnsignedSaturate}(a7 + b7)$

```
__m128i _mm_avg_epu8(__m128i a, __m128i b)
```

Computes the average of the 16 unsigned 8-bit integers in a and the 16 unsigned 8-bit integers in b and rounds.

```
r0 := (a0 + b0) / 2  
r1 := (a1 + b1) / 2  
...  
r15 := (a15 + b15) / 2
```

```
__m128i _mm_avg_epu16(__m128i a, __m128i b)
```

Computes the average of the 8 unsigned 16-bit integers in *a* and the 8 unsigned 16-bit integers in *b* and rounds.

```
r0 := (a0 + b0) / 2
r1 := (a1 + b1) / 2
...
r7 := (a7 + b7) / 2
```

```
__m128i _mm_madd_epi16(__m128i a, __m128i b)
```

Multiplies the 8 signed 16-bit integers from *a* by the 8 signed 16-bit integers from *b*. Adds the signed 32-bit integer results pairwise and packs the 4 signed 32-bit integer results.

```
r0 := (a0 * b0) + (a1 * b1)
r1 := (a2 * b2) + (a3 * b3)
r2 := (a4 * b4) + (a5 * b5)
r3 := (a6 * b6) + (a7 * b7)
```

```
__m128i _mm_max_epi16(__m128i a, __m128i b)
```

Computes the pairwise maxima of the 8 signed 16-bit integers from *a* and the 8 signed 16-bit integers from *b*.

```
r0 := max(a0, b0)
r1 := max(a1, b1)
...
r7 := max(a7, b7)
```

```
__m128i _mm_max_epu8(__m128i a, __m128i b)
```

Computes the pairwise maxima of the 16 unsigned 8-bit integers from *a* and the 16 unsigned 8-bit integers from *b*.

```
r0 := max(a0, b0)
r1 := max(a1, b1)
...
r15 := max(a15, b15)
```

```
__m128i _mm_min_epi16(__m128i a, __m128i b)
```

Computes the pairwise minima of the 8 signed 16-bit integers from *a* and the 8 signed 16-bit integers from *b*.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
...
r7 := min(a7, b7)
```

```
__m128i _mm_min_epu8(__m128i a, __m128i b)
```

Computes the pairwise minima of the 16 unsigned 8-bit integers from *a* and the 16 unsigned 8-bit integers from *b*.

```
r0 := min(a0, b0)
r1 := min(a1, b1)
...
r15 := min(a15, b15)
```

```
__m128i _mm_mulhi_epi16(__m128i a, __m128i b)
```

Multiplies the 8 signed 16-bit integers from *a* by the 8 signed 16-bit integers from *b*. Packs the upper 16-bits of the 8 signed 32-bit results.

```
r0 := (a0 * b0)[31:16]
r1 := (a1 * b1)[31:16]
...
r7 := (a7 * b7)[31:16]
```

```
__m128i _mm_mulhi_epu16(__m128i a, __m128i b)
```

Multiplies the 8 unsigned 16-bit integers from *a* by the 8 unsigned 16-bit integers from *b*. Packs the upper 16-bits of the 8 unsigned 32-bit results.

```
r0 := (a0 * b0)[31:16]
r1 := (a1 * b1)[31:16]
...
r7 := (a7 * b7)[31:16]
```

```
__m128i _mm_mullo_epi16(__m128i a, __m128i b)
```

Multiplies the 8 signed or unsigned 16-bit integers from *a* by the 8 signed or unsigned 16-bit integers from *b*. Packs the lower 16-bits of the 8 signed or unsigned 32-bit results.

```
r0 := (a0 * b0)[15:0]
r1 := (a1 * b1)[15:0]
...
r7 := (a7 * b7)[15:0]
```

```
__m64 _mm_mul_su32(__m64 a, __m64 b)
```

Multiplies the lower 32-bit integer from *a* by the lower 32-bit integer from *b*, and returns the 64-bit integer result.

```
r := a0 * b0
```

```
__m128i _mm_mul_epu32(__m128i a, __m128i b)
```

Multiplies 2 unsigned 32-bit integers from *a* by 2 unsigned 32-bit integers from *b*. Packs the 2 unsigned 64-bit integer results.

```
r0 := a0 * b0
r1 := a2 * b2
```

```
__m128i _mm_sad_epu8(__m128i a, __m128i b)
```

Computes the absolute difference of the 16 unsigned 8-bit integers from *a* and the 16 unsigned 8-bit integers from *b*. Sums the upper 8 differences and lower 8 differences, and packs the resulting 2 unsigned 16-bit integers into the upper and lower 64-bit elements.

```
r0 := abs(a0 - b0) + abs(a1 - b1) + ... + abs(a7 - b7)
r1 := 0x0 ; r2 := 0x0 ; r3 := 0x0
r4 := abs(a8 - b8) + abs(a9 - b9) + ... + abs(a15 - b15)
r5 := 0x0 ; r6 := 0x0 ; r7 := 0x0
```

```
__m128i _mm_sub_epi8(__m128i a, __m128i b)
```

Subtracts the 16 signed or unsigned 8-bit integers of *b* from the 16 signed or unsigned 8-bit integers of *a*.

```
r0 := a0 - b0
r1 := a1 - b1
...
r15 := a15 - b15
```

```
__m128i_mm_sub_epi16(__m128i a, __m128i b)
```

Subtracts the 8 signed or unsigned 16-bit integers of *b* from the 8 signed or unsigned 16-bit integers of *a*.

```
r0 := a0 - b0
r1 := a1 - b1
...
r7 := a7 - b7
```

```
__m128i_mm_sub_epi32(__m128i a, __m128i b)
```

Subtracts the 4 signed or unsigned 32-bit integers of *b* from the 4 signed or unsigned 32-bit integers of *a*.

```
r0 := a0 - b0
r1 := a1 - b1
r2 := a2 - b2
r3 := a3 - b3
```

```
__m64_mm_sub_si64 (__m64 a, __m64 b)
```

Subtracts the signed or unsigned 64-bit integer *b* from the signed or unsigned 64-bit integer *a*.

```
r := a - b
```

```
__m128i_mm_sub_epi64(__m128i a, __m128i b)
```

Subtracts the 2 signed or unsigned 64-bit integers in *b* from the 2 signed or unsigned 64-bit integers in *a*.

```
r0 := a0 - b0
r1 := a1 - b1
```

```
__m128i_mm_subs_epi8(__m128i a, __m128i b)
```

Subtracts the 16 signed 8-bit integers of *b* from the 16 signed 8-bit integers of *a* using saturating arithmetic.

```
r0 := SignedSaturate(a0 - b0)
r1 := SignedSaturate(a1 - b1)
...
r15 := SignedSaturate(a15 - b15)
```

```
__m128i_mm_subs_epi16(__m128i a, __m128i b)
```

Subtracts the 8 signed 16-bit integers of *b* from the 8 signed 16-bit integers of *a* using saturating arithmetic.

```
r0 := SignedSaturate(a0 - b0)
r1 := SignedSaturate(a1 - b1)
...
r7 := SignedSaturate(a7 - b7)
```

```
__m128i_mm_subs_epu8(__m128i a, __m128i b)
```

Subtracts the 16 unsigned 8-bit integers of *b* from the 16 unsigned 8-bit integers of *a* using saturating arithmetic.

```
r0 := UnsignedSaturate(a0 - b0)
r1 := UnsignedSaturate(a1 - b1)
...
r15 := UnsignedSaturate(a15 - b15)
```

```
__m128i _mm_subs_epu16(__m128i a, __m128i b)
```

Subtracts the 8 unsigned 16-bit integers of *b* from the 8 unsigned 16-bit integers of *a* using saturating arithmetic.

```
r0 := UnsignedSaturate(a0 - b0)
r1 := UnsignedSaturate(a1 - b1)
...
r7 := UnsignedSaturate(a7 - b7)
```

Integer Logical Operations for Streaming SIMD Extensions 2

The following four logical-operation intrinsics and their respective instructions are functional as part of Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
__m128i _mm_and_si128(__m128i a, __m128i b)
```

(uses `PAND`) Computes the bitwise AND of the 128-bit value in *a* and the 128-bit value in *b*.

```
r := a & b
```

```
__m128i _mm_andnot_si128(__m128i a, __m128i b)
```

(uses `PANDN`) Computes the bitwise AND of the 128-bit value in *b* and the bitwise NOT of the 128-bit value in *a*.

```
r := (~a) & b
```

```
__m128i _mm_or_si128(__m128i a, __m128i b)
```

(uses `POR`) Computes the bitwise OR of the 128-bit value in *a* and the 128-bit value in *b*.

```
r := a | b
```

```
__m128i _mm_xor_si128(__m128i a, __m128i b)
```

(uses `PXOR`) Computes the bitwise XOR of the 128-bit value in *a* and the 128-bit value in *b*.

```
r := a ^ b
```

Integer Shift Operations for Streaming SIMD Extensions 2

The shift-operation intrinsics for Streaming SIMD Extensions 2 (SSE2) and the description for each are listed in the following table.

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

Intrinsic	Shift Direction	Shift Type	Corresponding Instruction
<code>_mm_slli_si128</code>	Left	Logical	<code>PSLLDQ</code>
<code>_mm_slli_epi16</code>	Left	Logical	<code>PSLLW</code>
<code>_mm_sll_epi16</code>	Left	Logical	<code>PSLLW</code>
<code>_mm_slli_epi32</code>	Left	Logical	<code>PSLLD</code>

Intrinsic	Shift Direction	Shift Type	Corresponding Instruction
<code>_mm_sll_epi32</code>	Left	Logical	PSLLD
<code>_mm_slli_epi64</code>	Left	Logical	PSLLQ
<code>_mm_sll_epi64</code>	Left	Logical	PSLLQ
<code>_mm_srai_epi16</code>	Right	Arithmetic	PSRAW
<code>_mm_sra_epi16</code>	Right	Arithmetic	PSRAW
<code>_mm_srai_epi32</code>	Right	Arithmetic	PSRAD
<code>_mm_sra_epi32</code>	Right	Arithmetic	PSRAD
<code>_mm_srli_si128</code>	Right	Logical	PSRLDQ
<code>_mm_srli_epi16</code>	Right	Logical	PSRLW
<code>_mm_srl_epi16</code>	Right	Logical	PSRLW
<code>_mm_srli_epi32</code>	Right	Logical	PSRLD
<code>_mm_srl_epi32</code>	Right	Logical	PSRLD
<code>_mm_srli_epi64</code>	Right	Logical	PSRLQ
<code>_mm_srl_epi64</code>	Right	Logical	PSRLQ

```
__m128i _mm_slli_si128(__m128i a, int imm)
```

Shifts the 128-bit value in `a` left by `imm` bytes while shifting in zeros. `imm` must be an immediate.
`r := a << (imm * 8)`

```
__m128i _mm_slli_epi16(__m128i a, int count)
```

Shifts the 8 signed or unsigned 16-bit integers in `a` left by `count` bits while shifting in zeros.
`r0 := a0 << count`
`r1 := a1 << count`
`...`
`r7 := a7 << count`

```
__m128i _mm_sll_epi16(__m128i a, __m128i count)
```

Shifts the 8 signed or unsigned 16-bit integers in `a` left by `count` bits while shifting in zeros.
`r0 := a0 << count`
`r1 := a1 << count`
`...`
`r7 := a7 << count`

```
__m128i _mm_slli_epi32(__m128i a, int count)
```

Shifts the 4 signed or unsigned 32-bit integers in a left by count bits while shifting in zeros.

```
r0 := a0 << count
r1 := a1 << count
r2 := a2 << count
r3 := a3 << count
```

```
__m128i _mm_sll_epi32(__m128i a, __m128i count)
```

Shifts the 4 signed or unsigned 32-bit integers in a left by count bits while shifting in zeros.

```
r0 := a0 << count
r1 := a1 << count
r2 := a2 << count
r3 := a3 << count
```

```
__m128i _mm_slli_epi64(__m128i a, int count)
```

Shifts the 2 signed or unsigned 64-bit integers in a left by count bits while shifting in zeros.

```
r0 := a0 << count
r1 := a1 << count
```

```
__m128i _mm_sll_epi64(__m128i a, __m128i count)
```

Shifts the 2 signed or unsigned 64-bit integers in a left by count bits while shifting in zeros.

```
r0 := a0 << count
r1 := a1 << count
```

```
__m128i _mm_srai_epi16(__m128i a, int count)
```

Shifts the 8 signed 16-bit integers in a right by count bits while shifting in the sign bit.

```
r0 := a0 >> count
r1 := a1 >> count
...
r7 := a7 >> count
```

```
__m128i _mm_sra_epi16(__m128i a, __m128i count)
```

Shifts the 8 signed 16-bit integers in a right by count bits while shifting in the sign bit.

```
r0 := a0 >> count
r1 := a1 >> count
...
r7 := a7 >> count
```

```
__m128i _mm_srai_epi32(__m128i a, int count)
```

Shifts the 4 signed 32-bit integers in a right by count bits while shifting in the sign bit.

```
r0 := a0 >> count
r1 := a1 >> count
r2 := a2 >> count
r3 := a3 >> count
```

```
__m128i _mm_sra_epi32(__m128i a, __m128i count)
```

Shifts the 4 signed 32-bit integers in a right by count bits while shifting in the sign bit.

```
r0 := a0 >> count
r1 := a1 >> count
r2 := a2 >> count
r3 := i3 >> count
```

```
__m128i _mm_srli_si128(__m128i a, int imm)
```

Shifts the 128-bit value in a right by imm bytes while shifting in zeros. imm must be an immediate.

```
r := srl(a, imm*8)
```

```
__m128i _mm_srli_epi16(__m128i a, int count)
```

Shifts the 8 signed or unsigned 16-bit integers in a right by count bits while shifting in zeros.

```
r0 := srl(a0, count)
r1 := srl(a1, count)
...
r7 := srl(a7, count)
```

```
__m128i _mm_srl_epi16(__m128i a, __m128i count)
```

Shifts the 8 signed or unsigned 16-bit integers in a right by count bits while shifting in zeros.

```
r0 := srl(a0, count)
r1 := srl(a1, count)
...
r7 := srl(a7, count)
```

```
__m128i _mm_srli_epi32(__m128i a, int count)
```

Shifts the 4 signed or unsigned 32-bit integers in a right by count bits while shifting in zeros.

```
r0 := srl(a0, count)
r1 := srl(a1, count)
r2 := srl(a2, count)
r3 := srl(a3, count)
```

```
__m128i _mm_srl_epi32(__m128i a, __m128i count)
```

Shifts the 4 signed or unsigned 32-bit integers in a right by count bits while shifting in zeros.

```
r0 := srl(a0, count)
r1 := srl(a1, count)
r2 := srl(a2, count)
r3 := srl(a3, count)
```

```
__m128i _mm_srli_epi64(__m128i a, int count)
```

Shifts the 2 signed or unsigned 64-bit integers in a right by count bits while shifting in zeros.

```
r0 := srl(a0, count)
r1 := srl(a1, count)
```

```
__m128i _mm_srl_epi64(__m128i a, __m128i count)
```

Shifts the 2 signed or unsigned 64-bit integers in a right by count bits while shifting in zeros.

```
r0 := srl(a0, count)
r1 := srl(a1, count)
```

Integer Comparison Operations for Streaming SIMD Extensions 2

The comparison intrinsics for Streaming SIMD Extensions 2 (SSE2) and descriptions for each are listed in the following table.

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

Intrinsic Name	Instruction	Comparison	Elements	Size of Elements
<code>_mm_cmpeq_epi8</code>	PCMPEQB	Equality	16	8
<code>_mm_cmpeq_epi16</code>	PCMPEQW	Equality	8	16
<code>_mm_cmpeq_epi32</code>	PCMPEQD	Equality	4	32
<code>_mm_cmpgt_epi8</code>	PCMPGTB	Greater Than	16	8
<code>_mm_cmpgt_epi16</code>	PCMPGTW	Greater Than	8	16
<code>_mm_cmpgt_epi32</code>	PCMPGTD	Greater Than	4	32
<code>_mm_cmplt_epi8</code>	PCMPGTB _r	Less Than	16	8
<code>_mm_cmplt_epi16</code>	PCMPGTW _r	Less Than	8	16
<code>_mm_cmplt_epi32</code>	PCMPGTD _r	Less Than	4	32

```
__m128i _mm_cmpeq_epi8(__m128i a, __m128i b)
```

Compares the 16 signed or unsigned 8-bit integers in a and the 16 signed or unsigned 8-bit integers in b for equality.

```
r0 := (a0 == b0) ? 0xff : 0x0
r1 := (a1 == b1) ? 0xff : 0x0
...
r15 := (a15 == b15) ? 0xff : 0x0
```

```
__m128i _mm_cmpeq_epi16(__m128i a, __m128i b)
```

Compares the 8 signed or unsigned 16-bit integers in a and the 8 signed or unsigned 16-bit integers in b for equality.

```
r0 := (a0 == b0) ? 0xffff : 0x0
r1 := (a1 == b1) ? 0xffff : 0x0
...
r7 := (a7 == b7) ? 0xffff : 0x0
```

```
__m128i _mm_cmpeq_epi32(__m128i a, __m128i b)
```

Compares the 4 signed or unsigned 32-bit integers in *a* and the 4 signed or unsigned 32-bit integers in *b* for equality.

```
r0 := (a0 == b0) ? 0xffffffff : 0x0
r1 := (a1 == b1) ? 0xffffffff : 0x0
r2 := (a2 == b2) ? 0xffffffff : 0x0
r3 := (a3 == b3) ? 0xffffffff : 0x0
```

```
__m128i _mm_cmpgt_epi8(__m128i a, __m128i b)
```

Compares the 16 signed 8-bit integers in *a* and the 16 signed 8-bit integers in *b* for greater than.

```
r0 := (a0 > b0) ? 0xff : 0x0
r1 := (a1 > b1) ? 0xff : 0x0
...
r15 := (a15 > b15) ? 0xff : 0x0
```

```
__m128i _mm_cmpgt_epi16(__m128i a, __m128i b)
```

Compares the 8 signed 16-bit integers in *a* and the 8 signed 16-bit integers in *b* for greater than.

```
r0 := (a0 > b0) ? 0xffff : 0x0
r1 := (a1 > b1) ? 0xffff : 0x0
...
r7 := (a7 > b7) ? 0xffff : 0x0
```

```
__m128i _mm_cmpgt_epi32(__m128i a, __m128i b)
```

Compares the 4 signed 32-bit integers in *a* and the 4 signed 32-bit integers in *b* for greater than.

```
r0 := (a0 > b0) ? 0xffff : 0x0
r1 := (a1 > b1) ? 0xffff : 0x0
r2 := (a2 > b2) ? 0xffff : 0x0
r3 := (a3 > b3) ? 0xffff : 0x0
```

```
__m128i _mm_cmplt_epi8(__m128i a, __m128i b)
```

Compares the 16 signed 8-bit integers in *a* and the 16 signed 8-bit integers in *b* for less than.

```
r0 := (a0 < b0) ? 0xff : 0x0
r1 := (a1 < b1) ? 0xff : 0x0
...
r15 := (a15 < b15) ? 0xff : 0x0
```

```
__m128i _mm_cmplt_epi16(__m128i a, __m128i b)
```

Compares the 8 signed 16-bit integers in *a* and the 8 signed 16-bit integers in *b* for less than.

```
r0 := (a0 < b0) ? 0xffff : 0x0
r1 := (a1 < b1) ? 0xffff : 0x0
...
r7 := (a7 < b7) ? 0xffff : 0x0
```

```
__m128i _mm_cmplt_epi32(__m128i a, __m128i b)
```

Compares the 4 signed 32-bit integers in *a* and the 4 signed 32-bit integers in *b* for less than.

```
r0 := (a0 < b0) ? 0xffff : 0x0
r1 := (a1 < b1) ? 0xffff : 0x0
r2 := (a2 < b2) ? 0xffff : 0x0
r3 := (a3 < b3) ? 0xffff : 0x0
```

Integer Conversions Operations for Streaming SIMD Extensions 2

The following two conversion intrinsics and their respective instructions are functional in the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
__m128i _mm_cvtsi32_si128(int a)
```

(uses `MOVD`) Moves 32-bit integer `a` to the least significant 32 bits of an `__m128i` object. Copies the sign bit of `a` into the upper 96 bits of the `__m128i` object.

```
r0 := a
r1 := 0x0 ; r2 := 0x0 ; r3 := 0x0
```

```
int _mm_cvtsi128_si32(__m128i a)
```

(uses `MOVD`) Moves the least significant 32 bits of `a` to a 32 bit integer.

```
r := a0
```

```
__m128 _mm_cvtepi32_ps(__m128i a)
```

Converts the 4 signed 32-bit integer values of `a` to SP FP values.

```
r0 := (float) a0
r1 := (float) a1
r2 := (float) a2
r3 := (float) a3
```

```
__m128i _mm_cvtps_epi32(__m128 a)
```

Converts the 4 SP FP values of `a` to signed 32-bit integer values.

```
r0 := (int) a0
r1 := (int) a1
r2 := (int) a2
r3 := (int) a3
```

```
__m128i _mm_cvttps_epi32(__m128 a)
```

Converts the 4 SP FP values of `a` to signed 32 bit integer values using truncate.

```
r0 := (int) a0
r1 := (int) a1
r2 := (int) a2
r3 := (int) a3
```

Integer Memory and Initialization Operations for Streaming SIMD Extensions 2

The integer load, set, and store intrinsics and their respective instructions provide memory and initialization operations for the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

- Load Operations
- Set Operations
- Store Operations

Integer Load Operations for Streaming SIMD Extensions 2

The following `load` operation intrinsics and their respective instructions are functional in the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
__m128i _mm_load_si128(__m128i const*p)
```

(uses `MOVDQA`) Loads 128-bit value. Address `p` must be 16-byte aligned.

```
r := *p
```

```
__m128i _mm_loadu_si128(__m128i const*p)
```

(uses `MOVDQU`) Loads 128-bit value. Address `p` not need be 16-byte aligned.

```
r := *p
```

```
__m128i _mm_loadl_epi64(__m128i const*p)
```

(uses `MOVQ`) Load the lower 64 bits of the value pointed to by `p` into the lower 64 bits of the result, zeroing the upper 64 bits of the result.

```
r0 := *p[63:0]
```

```
r1 := 0x0
```

Integer Set Operations for SSE2

The following `set` operation intrinsics and their respective instructions are functional in the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
__m128i _mm_set_epi64(__m64 q1, __m64 q0)
```

Sets the 2 64-bit integer values.

```
r0 := q0
```

```
r1 := q1
```

```
__m128i _mm_set_epi32(int i3, int i2, int i1, int i0)
```

Sets the 4 signed 32-bit integer values.

```
r0 := i0
```

```
r1 := i1
```

```
r2 := i2
```

```
r3 := i3
```

```
__m128i _mm_set_epi16(short w7, short w6, short w5, short w4, short w3,
short w2, short w1, short w0)
```

Sets the 8 signed 16-bit integer values.

```
r0 := w0
```

```
r1 := w1
```

```
...
```

```
r7 := w7
```

```
__m128i _mm_set_epi8(char b15, char b14, char b13, char b12, char b11,
char b10, char b9, char b8, char b7, char b6, char b5, char b4, char b3,
char b2, char b1, char b0)
```

Sets the 16 signed 8-bit integer values.

```
r0 := b0
r1 := b1
...
r15 := b15
```

```
__m128i _mm_set1_epi64(__m64 q)
```

Sets the 2 64-bit integer values to q.

```
r0 := q
r1 := q
```

```
__m128i _mm_set1_epi32(int i)
```

Sets the 4 signed 32-bit integer values to i.

```
r0 := i
r1 := i
r2 := i
r3 := i
```

```
__m128i _mm_set1_epi16(short w)
```

Sets the 8 signed 16-bit integer values to w.

```
r0 := w
r1 := w
...
r7 := w
```

```
__m128i _mm_set1_epi8(char b)
```

Sets the 16 signed 8-bit integer values to b.

```
r0 := b
r1 := b
...
r15 := b
```

```
__m128i _mm_setr_epi64(__m64 q0, __m64 q1)
```

Sets the 2 64-bit integer values in reverse order.

```
r0 := q0
r1 := q1
```

```
__m128i _mm_setr_epi32(int i0, int i1, int i2, int i3)
```

Sets the 4 signed 32-bit integer values in reverse order.

```
r0 := i0
r1 := i1
r2 := i2
r3 := i3
```

```
__m128i _mm_setr_epi16(short w0, short w1, short w2, short w3, short w4,
short w5, short w6, short w7)
```

Sets the 8 signed 16-bit integer values in reverse order.

```
r0 := w0
r1 := w1
...
r7 := w7
```

```
__m128i _mm_setr_epi8(char b15, char b14, char b13, char b12, char b11,
char b10, char b9, char b8, char b7, char b6, char b5, char b4, char b3,
char b2, char b1, char b0)
```

Sets the 16 signed 8-bit integer values in reverse order.

```
r0 := b0
r1 := b1
...
r15 := b15
```

```
__m128i _mm_setzero_si128()
```

Sets the 128-bit value to zero.

```
r := 0x0
```

Integer Store Operations for Streaming SIMD Extensions 2

The following store operation intrinsics and their respective instructions are functional in the Streaming SIMD Extensions 2 (SSE2).

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

```
void _mm_store_si128(__m128i *p, __m128i b)
```

(uses `MOVDQA`) Stores 128-bit value. Address `p` must be 16 byte aligned.

```
*p := a
```

```
void _mm_storeu_si128(__m128i *p, __m128i b)
```

(uses `MOVDQU`) Stores 128-bit value. Address `p` need not be 16-byte aligned.

```
*p := a
```

```
void _mm_maskmoveu_si128(__m128i d, __m128i n, char *p)
```

(uses `MASKMOVDQU`) Conditionally store byte elements of `d` to address `p`. The high bit of each byte in the selector `n` determines whether the corresponding byte in `d` will be stored. Address `p` need not be 16-byte aligned.

```
if (n0[7]) p[0] := d0
if (n1[7]) p[1] := d1
...
if (n15[7]) p[15] := d15
```

```
void _mm_storel_epi64(__m128i *p, __m128i q)
```

(uses `MOVQ`) Stores the lower 64 bits of the value pointed to by `p`.

```
*p[63:0] := a0
```

Macro Function for Shuffle

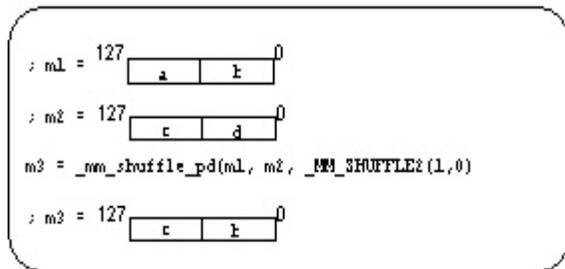
The Streaming SIMD Extensions 2 (SSE2) provide a macro function to help create constants that describe shuffle operations. The macro takes two small integers (in the range of 0 to 1) and combines them into an 2-bit immediate value used by the SHUFPD instruction. See the following example.

Shuffle Function Macro

```
_MM_SHUFFLE2(x, y)
expands to the value of
((x<<1) | y)
```

You can view the two integers as selectors for choosing which two words from the first input operand and which two words from the second are to be put into the result word.

View of Original and Result Words with Shuffle Function Macro



Cacheability Support Operations for Streaming SIMD Extensions 2

The prototypes for Streaming SIMD Extensions 2 (SSE2) intrinsics are in the `emmintrin.h` header file.

```
void _mm_stream_pd(double *p, __m128d a)
```

(uses MOVNTPD) Stores the data in `a` to the address `p` without polluting caches. The address `p` must be 16-byte aligned. If the cache line containing address `p` is already in the cache, the cache will be updated.

```
p[0] := a0
p[1] := a1
```

```
void _mm_stream_si128(__m128i *p, __m128i a)
```

Stores the data in `a` to the address `p` without polluting the caches. If the cache line containing address `p` is already in the cache, the cache will be updated. Address `p` must be 16-byte aligned.

```
*p := a
```

```
void _mm_stream_si32(int *p, int a)
```

Stores the data in `a` to the address `p` without polluting the caches. If the cache line containing address `p` is already in the cache, the cache will be updated.

```
*p := a
```

```
void _mm_clflush(void const*p)
```

Cache line containing `p` is flushed and invalidated from all caches in the coherency domain.

```
void _mm_lfence(void)
```

Guarantees that every load instruction that precedes, in program order, the load fence instruction is globally visible before any load instruction which follows the fence in program order.

```
void _mm mfence(void)
```

Guarantees that every memory access that precedes, in program order, the memory fence instruction is globally visible before any memory instruction which follows the fence in program order.

```
void _mm_pause(void)
```

The execution of the next instruction is delayed an implementation specific amount of time. The instruction does not modify the architectural state. This intrinsic provides especially significant performance gain.

PAUSE Intrinsic

The PAUSE intrinsic is used in spin-wait loops with the processors implementing dynamic execution (especially out-of-order execution). In the spin-wait loop, PAUSE improves the speed at which the code detects the release of the lock. For dynamic scheduling, the PAUSE instruction reduces the penalty of exiting from the spin-loop.

Example of loop with the PAUSE instruction:

```
spin_loop:pause
cmp eax, A
jne spin_loop
```

In this example, the program spins until memory location `A` matches the value in register `eax`. The code sequence that follows shows a test-and-test-and-set. In this example, the spin occurs only after the attempt to get a lock has failed.

```
get_lock: mov eax, 1
xchg eax, A ; Try to get lock
cmp eax, 0 ; Test if successful
jne spin_loop
critical_section code
mov A, 0 ; Release lock
jmp continue
spin_loop: pause ; Spin-loop hint
cmp 0, A ; Check lock availability
jne spin_loop
jmp get_lock
continue:
```

Note that the first branch is predicted to fall-through to the critical section in anticipation of successfully gaining access to the lock. It is highly recommended that all spin-wait loops include the PAUSE instruction. Since PAUSE is backwards compatible to all existing IA-32 processor generations, a test for processor type

(a CPUID test) is not needed. All legacy processors will execute PAUSE as a NOP, but in processors which use the PAUSE as a hint there can be significant performance benefit.

Miscellaneous Operations for Streaming SIMD Extensions 2

The miscellaneous intrinsics for Streaming SIMD Extensions 2 (SSE2) are listed in the following table followed by their descriptions.

The prototypes for SSE2 intrinsics are in the `emmintrin.h` header file.

Intrinsic	Corresponding Instruction	Operation
<code>_mm_packs_epi16</code>	PACKSSWB	Packed Saturation
<code>_mm_packs_epi32</code>	PACKSSDW	Packed Saturation
<code>_mm_packus_epi16</code>	PACKUSWB	Packed Saturation
<code>_mm_extract_epi16</code>	PEXTRW	Extraction
<code>_mm_insert_epi16</code>	PINSRW	Insertion
<code>_mm_movemask_epi8</code>	PMOVMSKB	Mask Creation
<code>_mm_shuffle_epi32</code>	PSHUFD	Shuffle
<code>_mm_shufflehi_epi16</code>	PSHUFW	Shuffle
<code>_mm_shufflelo_epi16</code>	PSHUFLW	Shuffle
<code>_mm_unpackhi_epi8</code>	PUNPCKHBW	Interleave
<code>_mm_unpackhi_epi16</code>	PUNPCKHWD	Interleave
<code>_mm_unpackhi_epi32</code>	PUNPCKHDQ	Interleave
<code>_mm_unpackhi_epi64</code>	PUNPCKHQDQ	Interleave
<code>_mm_unpacklo_epi8</code>	PUNPCKLBW	Interleave
<code>_mm_unpacklo_epi16</code>	PUNPCKLWD	Interleave
<code>_mm_unpacklo_epi32</code>	PUNPCKLDQ	Interleave
<code>_mm_unpacklo_epi64</code>	PUNPCKLQDQ	Interleave
<code>_mm_movepi64_pi64</code>	MOVDQ2Q	move
<code>_m128i_mm_movpi64_epi64</code>	MOVQ2DQ	move
<code>_mm_move_epi64</code>	MOVQ	move

```
__m128i _mm_packs_epi16(__m128i a, __m128i b)
```

Packs the 16 signed 16-bit integers from a and b into 8-bit integers and saturates.

```
r0 := SignedSaturate(a0)
r1 := SignedSaturate(a1)
...
r7 := SignedSaturate(a7)
r8 := SignedSaturate(b0)
r9 := SignedSaturate(b1)
...
r15 := SignedSaturate(b7)
```

```
__m128i _mm_packs_epi32(__m128i a, __m128i b)
```

Packs the 8 signed 32-bit integers from a and b into signed 16-bit integers and saturates.

```
r0 := SignedSaturate(a0)
r1 := SignedSaturate(a1)
r2 := SignedSaturate(a2)
r3 := SignedSaturate(a3)
r4 := SignedSaturate(b0)
r5 := SignedSaturate(b1)
r6 := SignedSaturate(b2)
r7 := SignedSaturate(b3)
```

```
__m128i _mm_packus_epi16(__m128i a, __m128i b)
```

Packs the 16 signed 16-bit integers from a and b into 8-bit unsigned integers and saturates.

```
r0 := UnsignedSaturate(a0)
r1 := UnsignedSaturate(a1)
...
r7 := UnsignedSaturate(a7)
r8 := UnsignedSaturate(b0)
r9 := UnsignedSaturate(b1)
...
r15 := UnsignedSaturate(b7)
```

```
int _mm_extract_epi16(__m128i a, int imm)
```

Extracts the selected signed or unsigned 16-bit integer from a and zero extends. The selector `imm` must be an immediate.

```
r := (imm == 0) ? a0 :
( (imm == 1) ? a1 :
...
(imm == 7) ? a7 )
```

```
__m128i _mm_insert_epi16(__m128i a, int b, int imm)
```

Inserts the least significant 16 bits of b into the selected 16-bit integer of a. The selector `imm` must be an immediate.

```
r0 := (imm == 0) ? b : a0;
r1 := (imm == 1) ? b : a1;
...
r7 := (imm == 7) ? b : a7;
```

```
int __mm_movemask_epi8(__m128i a)
```

Creates a 16-bit mask from the most significant bits of the 16 signed or unsigned 8-bit integers in `a` and zero extends the upper bits.

```
r := a15[7] << 15 |
a14[7] << 14 |
...
a1[7] << 1 |
a0[7]
```

```
__m128i __mm_shuffle_epi32(__m128i a, int imm)
```

Shuffles the 4 signed or unsigned 32-bit integers in `a` as specified by `imm`. The shuffle value, `imm`, must be an immediate. See Macro Function for Shuffle for a description of shuffle semantics.

```
__m128i __mm_shufflehi_epi16(__m128i a, int imm)
```

Shuffles the upper 4 signed or unsigned 16-bit integers in `a` as specified by `imm`. The shuffle value, `imm`, must be an immediate. See Macro Function for Shuffle for a description of shuffle semantics.

```
__m128i __mm_shufflelo_epi16(__m128i a, int imm)
```

Shuffles the lower 4 signed or unsigned 16-bit integers in `a` as specified by `imm`. The shuffle value, `imm`, must be an immediate. See Macro Function for Shuffle for a description of shuffle semantics.

```
__m128i __mm_unpackhi_epi8(__m128i a, __m128i b)
```

Interleaves the upper 8 signed or unsigned 8-bit integers in `a` with the upper 8 signed or unsigned 8-bit integers in `b`.

```
r0 := a8 ; r1 := b8
r2 := a9 ; r3 := b9
...
r14 := a15 ; r15 := b15
```

```
__m128i __mm_unpackhi_epi16(__m128i a, __m128i b)
```

Interleaves the upper 4 signed or unsigned 16-bit integers in `a` with the upper 4 signed or unsigned 16-bit integers in `b`.

```
r0 := a4 ; r1 := b4
r2 := a5 ; r3 := b5
r4 := a6 ; r5 := b6
r6 := a7 ; r7 := b7
```

```
__m128i __mm_unpackhi_epi32(__m128i a, __m128i b)
```

Interleaves the upper 2 signed or unsigned 32-bit integers in `a` with the upper 2 signed or unsigned 32-bit integers in `b`.

```
r0 := a2 ; r1 := b2
r2 := a3 ; r3 := b3
```

```
__m128i _mm_unpackhi_epi64(__m128i a, __m128i b)
```

Interleaves the upper signed or unsigned 64-bit integer in a with the upper signed or unsigned 64-bit integer in b.

```
r0 := a1 ; r1 := b1
```

```
__m128i _mm_unpacklo_epi8(__m128i a, __m128i b)
```

Interleaves the lower 8 signed or unsigned 8-bit integers in a with the lower 8 signed or unsigned 8-bit integers in b.

```
r0 := a0 ; r1 := b0
```

```
r2 := a1 ; r3 := b1
```

```
...
```

```
r14 := a7 ; r15 := b7
```

```
__m128i _mm_unpacklo_epi16(__m128i a, __m128i b)
```

Interleaves the lower 4 signed or unsigned 16-bit integers in a with the lower 4 signed or unsigned 16-bit integers in b.

```
r0 := a0 ; r1 := b0
```

```
r2 := a1 ; r3 := b1
```

```
r4 := a2 ; r5 := b2
```

```
r6 := a3 ; r7 := b3
```

```
__m128i _mm_unpacklo_epi32(__m128i a, __m128i b)
```

Interleaves the lower 2 signed or unsigned 32-bit integers in a with the lower 2 signed or unsigned 32-bit integers in b.

```
r0 := a0 ; r1 := b0
```

```
r2 := a1 ; r3 := b1
```

```
__m128i _mm_unpacklo_epi64(__m128i a, __m128i b)
```

Interleaves the lower signed or unsigned 64-bit integer in a with the lower signed or unsigned 64-bit integer in b.

```
r0 := a0 ; r1 := b0
```

```
__m64 _mm_movepi64_pi64(__m128i a)
```

Returns the lower 64 bits of a as an __m64 type.

```
r0 := a0 ;
```

```
__128i _mm_movpi64_pi64(__m64 a)
```

Moves the 64 bits of a to the lower 64 bits of the result, zeroing the upper bits.

```
r0 := a0 ; r1 := 0X0 ;
```

```
__128i _mm_move_epi64(__128i a)
```

Moves the lower 64 bits of the lower 64 bits of the result, zeroing the upper bits.

```
r0 := a0 ; r1 := 0X0 ;
```

Additional Miscellaneous Intrinsics

The prototypes for Streaming SIMD Extensions 2 (SSE2) intrinsics are in the `emmintrin.h` header file.

```
__m128d _mm_unpackhi_pd(__m128d a, __m128d b)
```

(uses UNPCKHPD) Interleaves the upper DP FP values of a and b.

```
r0 := a1  
r1 := b1
```

```
__m128d _mm_unpacklo_pd(__m128d a, __m128d b)
```

(uses UNPCKLPD) Interleaves the lower DP FP values of a and b.

```
r0 := a0  
r1 := b0
```

```
int _mm_movemask_pd(__m128d a)
```

(uses MOVMSKPD) Creates a two-bit mask from the sign bits of the two DP FP values of a.

```
r := sign(a1) << 1 | sign(a0)
```

```
__m128d _mm_shuffle_pd(__m128d a, __m128d b, int i)
```

(uses SHUFPS) Selects two specific DP FP values from a and b, based on the mask i. The mask must be an immediate. See Macro Function for Shuffle for a description of the shuffle semantics.

Intrinsics for Casting Support

This version of the Intel C++ Compiler supports casting between various SP, DP, and INT vector types. These intrinsics do not convert values; they just change the type.

```
extern __m128  _mm_castpd_ps(__m128d in);
```

```
extern __m128i _mm_castpd_si128(__m128d in);
```

```
extern __m128d _mm_castps_pd(__m128 in);
```

```
extern __m128i _mm_castps_si128(__m128 in);
```

```
extern __m128  _mm_castsi128_ps(__m128i in);
```

```
extern __m128d _mm_castsi128_pd(__m128i in);
```

Streaming SIMD Extensions 3

The Intel® C++ intrinsics listed in this section are designed for the Intel® Pentium® 4 processor with Streaming SIMD Extensions 3 (SSE3). They will not function correctly on other IA-32 processors. New SSE3 intrinsics include:

- Floating-point Vector Intrinsics
- Integer Vector Intrinsics
- Miscellaneous Intrinsics
- Macro Functions

The prototypes for these intrinsics are in the `pmmintrin.h` header file.



Note

You can also use the single `ia32intrin.h` header file for any IA-32 intrinsics.

Floating-point Vector Intrinsics for Streaming SIMD Extensions 3

The floating-point intrinsics listed here are designed for the Intel® Pentium® 4 processor with Streaming SIMD Extensions 3 (SSE3).

The prototypes for these intrinsics are in the `pmmintrin.h` header file.

Single-precision Floating-point Vector Intrinsics

```
extern __m128 _mm_addsub_ps(__m128 a, __m128 b);
```

Subtracts even vector elements while adding odd vector elements.

```
r0 := a0 - b0;
r1 := a1 + b1;
r2 := a2 - b2;
r3 := a3 + b3;
```

```
extern __m128 _mm_hadd_ps(__m128 a, __m128 b);
```

Adds adjacent vector elements.

```
r0 := a0 + a1;
r1 := a2 + a3;
r2 := b0 + b1;
r3 := b2 + b3;
```

```
extern __m128 _mm_hsub_ps(__m128 a, __m128 b);
```

Subtracts adjacent vector elements.

```
r0 := a0 - a1;
r1 := a2 - a3;
r2 := b0 - b1;
r3 := b2 - b3;
```

```
extern __m128 _mm_movehdup_ps(__m128 a);
```

Duplicates odd vector elements into even vector elements.

```
r0 := a1;
r1 := a1;
r2 := a3;
r3 := a3;
```

```
extern __m128 _mm_moveldup_ps(__m128 a);
```

Duplicates even vector elements into odd vector elements.

```
r0 := a0;
r1 := a0;
r2 := a2;
r3 := a2;
```

Double-precision Floating-point Vector Intrinsics

```
extern __m128d _mm_addsub_pd(__m128d a, __m128d b);
```

Adds upper vector element while subtracting lower vector element.

```
r0 := a0 - b0;
r1 := a1 + b1;
```

```
extern __m128d _mm_hadd_pd(__m128d a, __m128d b);
```

Adds adjacent vector elements.

```
r0 := a0 + a1;
r1 := b0 + b1;
```

```
extern __m128d _mm_hsub_pd(__m128d a, __m128d b);
```

Subtracts adjacent vector elements.

```
r0 := a0 - a1;
r1 := b0 - b1;
```

```
extern __m128d _mm_loaddup_pd(double const * dp);
```

Duplicates a double value into upper and lower vector elements.

```
r0 := *dp;
r1 := *dp;
```

```
extern __m128d _mm_movedup_pd(__m128d a);
```

Duplicates lower vector element into upper vector element.

```
r0 := a0;
r1 := a0;
```

Integer Vector Intrinsics for Streaming SIMD Extensions 3

The integer vector intrinsic listed here is designed for the Intel® Pentium® 4 processor with Streaming SIMD Extensions 3 (SSE3).

The prototypes for these intrinsics are in the `pmmmintrin.h` header file.

```
extern __m128i _mm_lddqu_si128(__m128i const *p);
```

Loads an unaligned 128-bit value. This differs from `movdqu` in that it can provide higher performance in some cases. However, it also may provide lower performance than `movdqu` if the memory value being read was just previously written.

```
r := *p;
```

Macro Functions for Streaming SIMD Extensions 3

The macro function intrinsics listed here are designed for the Intel® Pentium® 4 processor with Streaming SIMD Extensions 3 (SSE3).

The prototypes for these intrinsics are in the `pmmmintrin.h` header file.

```
_MM_SET_DENORMALS_ZERO_MODE(x)
```

Macro arguments: one of `__MM_DENORMALS_ZERO_ON`, `__MM_DENORMALS_ZERO_OFF`

This causes "denormals are zero" mode to be turned on or off by setting the appropriate bit of the control register.

```
_MM_GET_DENORMALS_ZERO_MODE()
```

No arguments. This returns the current value of the denormals are zero mode bit of the control register.

Miscellaneous Intrinsics for Streaming SIMD Extensions 3

The miscellaneous intrinsics listed here are designed for the Intel® Pentium® 4 processor with Streaming SIMD Extensions 3 (SSE3).

The prototypes for these intrinsics are in the `pmmmintrin.h` header file.

```
extern void _mm_monitor(void const *p, unsigned extensions, unsigned hints);
```

Generates the `MONITOR` instruction. This sets up an address range for the monitor hardware using `p` to provide the logical address, and will be passed to the monitor instruction in register `eax`. The `extensions` parameter contains optional extensions to the monitor hardware which will be passed in `ecx`. The `hints` parameter will contain hints to the monitor hardware, which will be passed in `edx`. A non-zero value for `extensions` will cause a general protection fault.

```
extern void _mm_mwait(unsigned extensions, unsigned hints);
```

Generates the `MWAIT` instruction. This instruction is a hint that allows the processor to stop execution and enter an implementation-dependent optimized state until occurrence of a class of events. In future processor designs `extensions` and `hints` parameters may be used to convey additional information to the processor. All

non-zero values of extensions and hints are reserved. A non-zero value for extensions will cause a general protection fault.

Intrinsics for Itanium® Instructions

This section lists and describes the native intrinsics for Itanium® instructions. These intrinsics cannot be used on the IA-32 architecture. The intrinsics for Itanium instructions give programmers access to Itanium instructions that cannot be generated using the standard constructs of the C and C++ languages.

The prototypes for these intrinsics are in the `ia64intrin.h` header file.

Native Intrinsics for Itanium® Instructions

The prototypes for these intrinsics are in the `ia64intrin.h` header file.

Integer Operations

Intrinsic	Corresponding Instruction
<code>__int64 _m64_dep_mr(__int64 r, __int64 s, const int pos, const int len)</code>	<code>dep</code> (Deposit)
<code>__int64 _m64_dep_mi(const int v, __int64 s, const int p, const int len)</code>	<code>dep</code> (Deposit)
<code>__int64 _m64_dep_zr(__int64 s, const int pos, const int len)</code>	<code>dep.z</code> (Deposit)
<code>__int64 _m64_dep_zi(const int v, const int pos, const int len)</code>	<code>dep.z</code> (Deposit)
<code>__int64 _m64_extr(__int64 r, const int pos, const int len)</code>	<code>extr</code> (Extract)
<code>__int64 _m64_extru(__int64 r, const int pos, const int len)</code>	<code>extr.u</code> (Extract)
<code>__int64 _m64_xmal(__int64 a, __int64 b, __int64 c)</code>	<code>xma.l</code> (Fixed-point multiply add using the low 64 bits of the 128-bit result. The result is signed.)
<code>__int64 _m64_xmalu(__int64 a, __int64 b, __int64 c)</code>	<code>xma.lu</code> (Fixed-point multiply add using the low 64 bits of the 128-bit result. The result is unsigned.)
<code>__int64 _m64_xmah(__int64 a, __int64 b, __int64 c)</code>	<code>xma.h</code> (Fixed-point multiply add using the high 64 bits of the 128-bit result. The result is signed.)
<code>__int64 _m64_xmahu(__int64 a, __int64 b, __int64 c)</code>	<code>xma.hu</code> (Fixed-point multiply add using the high 64 bits of the 128-bit result. The result is unsigned.)
<code>__int64 _m64_popcnt(__int64 a)</code>	<code>popcnt</code> (Population count)

Intrinsic	Corresponding Instruction
<code>__int64 _m64_shladd(__int64 a, const int count, __int64 b)</code>	<code>shladd</code> (Shift left and add)
<code>__int64 _m64_shrp(__int64 a, __int64 b, const int count)</code>	<code>shrp</code> (Shift right pair)

FSR Operations

Intrinsic	Description
<code>void _fsetc(int amask, int omask)</code>	Sets the control bits of <code>FPSR.sf0</code> . Maps to the <code>fsetc.sf0 r, r</code> instruction. There is no corresponding instruction to read the control bits. Use <code>_mm_getfpsr()</code> .
<code>void _fclrf(void)</code>	Clears the floating point status flags (the 6-bit flags of <code>FPSR.sf0</code>). Maps to the <code>fclrf.sf0</code> instruction.

```
__int64 _m64_dep_mr(__int64 r, __int64 s, const int pos, const int len)
```

The right-justified 64-bit value `r` is deposited into the value in `s` at an arbitrary bit position and the result is returned. The deposited bit field begins at bit position `pos` and extends to the left (toward the most significant bit) the number of bits specified by `len`.

```
__int64 _m64_dep_mi(const int v, __int64 s, const int p, const int len)
```

The sign-extended value `v` (either all 1s or all 0s) is deposited into the value in `s` at an arbitrary bit position and the result is returned. The deposited bit field begins at bit position `p` and extends to the left (toward the most significant bit) the number of bits specified by `len`.

```
__int64 _m64_dep_zr(__int64 s, const int pos, const int len)
```

The right-justified 64-bit value `s` is deposited into a 64-bit field of all zeros at an arbitrary bit position and the result is returned. The deposited bit field begins at bit position `pos` and extends to the left (toward the most significant bit) the number of bits specified by `len`.

```
__int64 _m64_dep_zi(const int v, const int pos, const int len)
```

The sign-extended value `v` (either all 1s or all 0s) is deposited into a 64-bit field of all zeros at an arbitrary bit position and the result is returned. The deposited bit field begins at bit position `pos` and extends to the left (toward the most significant bit) the number of bits specified by `len`.

```
__int64 _m64_extr(__int64 r, const int pos, const int len)
```

A field is extracted from the 64-bit value `r` and is returned right-justified and sign extended. The extracted field begins at position `pos` and extends `len` bits to the left. The sign is taken from the most significant bit of the extracted field.

```
__int64 _m64_extru(__int64 r, const int pos, const int len)
```

A field is extracted from the 64-bit value `r` and is returned right-justified and zero extended. The extracted field begins at position `pos` and extends `len` bits to the left.

```
__int64 _m64_xmal(__int64 a, __int64 b, __int64 c)
```

The 64-bit values `a` and `b` are treated as signed integers and multiplied to produce a full 128-bit signed result. The 64-bit value `c` is zero-extended and added to the product. The least significant 64 bits of the sum are then returned.

```
__int64 _m64_xmalu(__int64 a, __int64 b, __int64 c)
```

The 64-bit values `a` and `b` are treated as signed integers and multiplied to produce a full 128-bit unsigned result. The 64-bit value `c` is zero-extended and added to the product. The least significant 64 bits of the sum are then returned.

```
__int64 _m64_xmah(__int64 a, __int64 b, __int64 c)
```

The 64-bit values `a` and `b` are treated as signed integers and multiplied to produce a full 128-bit signed result. The 64-bit value `c` is zero-extended and added to the product. The most significant 64 bits of the sum are then returned.

```
__int64 _m64_xmahu(__int64 a, __int64 b, __int64 c)
```

The 64-bit values `a` and `b` are treated as unsigned integers and multiplied to produce a full 128-bit unsigned result. The 64-bit value `c` is zero-extended and added to the product. The most significant 64 bits of the sum are then returned.

```
__int64 _m64_popcnt(__int64 a)
```

The number of bits in the 64-bit integer `a` that have the value 1 are counted, and the resulting sum is returned.

```
__int64 _m64_shladd(__int64 a, const int count, __int64 b)
```

`a` is shifted to the left by `count` bits and then added to `b`. The result is returned.

```
__int64 _m64_shrp(__int64 a, __int64 b, const int count)
```

`a` and `b` are concatenated to form a 128-bit value and shifted to the right `count` bits. The least significant 64 bits of the result are returned.

Lock and Atomic Operation Related Intrinsics

The prototypes for these intrinsics are in the `ia64intrin.h` header file.

Intrinsic	Description
<code>unsigned __int64 _InterlockedExchange8(volatile unsigned char *Target, unsigned __int64 value)</code>	Map to the <code>xchg1</code> instruction. Atomically write the least significant byte of its 2nd argument to address specified by its 1st argument.
<code>unsigned __int64 _InterlockedCompareExchange8_rel(volatile unsigned char *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</code>	Compare and exchange atomically the least significant byte at the address specified by its 1st argument. Maps to the <code>cmpxchg1_rel</code> instruction with appropriate setup.
<code>unsigned __int64 _InterlockedCompareExchange8_acq(volatile unsigned char *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</code>	Same as the previous intrinsic, but using <code>acquire</code> semantic.
<code>unsigned __int64 _InterlockedExchange16(volatile unsigned short *Target, unsigned __int64 value)</code>	Map to the <code>xchg2</code> instruction. Atomically write the least significant word of its 2nd argument to address specified by its 1st argument.
<code>unsigned __int64 _InterlockedCompareExchange16_rel(volatile unsigned short *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</code>	Compare and exchange atomically the least significant word at the address specified by its 1st argument. Maps to the <code>cmpxchg2_rel</code> instruction with appropriate setup.
<code>unsigned __int64 _InterlockedCompareExchange16_acq(volatile unsigned short *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</code>	Same as the previous intrinsic, but using <code>acquire</code> semantic.
<code>int _InterlockedIncrement(volatile int *addend)</code>	Atomically increment by one the value specified by its argument. Maps to the <code>fetchadd4</code> instruction.
<code>int _InterlockedDecrement(volatile int *addend)</code>	Atomically decrement by one the value specified by its argument. Maps to the <code>fetchadd4</code> instruction.
<code>int _InterlockedExchange(volatile int *Target, long value)</code>	Do an exchange operation atomically. Maps to the <code>xchg4</code> instruction.

Intrinsic	Description
<pre>int __InterlockedCompareExchange(volatile int *Destination, int Exchange, int Comparand)</pre>	<p>Do a compare and exchange operation atomically. Maps to the <code>cmpxchg4</code> instruction with appropriate setup.</p>
<pre>int __InterlockedExchangeAdd(volatile int *addend, int increment)</pre>	<p>Use compare and exchange to do an atomic add of the increment value to the <code>addend</code>. Maps to a loop with the <code>cmpxchg4</code> instruction to guarantee atomicity.</p>
<pre>int __InterlockedAdd(volatile int *addend, int increment)</pre>	<p>Same as the previous intrinsic, but returns new value, not the original one.</p>
<pre>void * __InterlockedCompareExchangePointer(void * volatile *Destination, void *Exchange, void *Comparand)</pre>	<p>Map the <code>exch8</code> instruction; Atomically compare and exchange the pointer value specified by its first argument (all arguments are pointers)</p>
<pre>unsigned __int64 __InterlockedExchangeU(volatile unsigned int *Target, unsigned __int64 value)</pre>	<p>Atomically exchange the 32-bit quantity specified by the 1st argument. Maps to the <code>xchg4</code> instruction.</p>
<pre>unsigned __int64 __InterlockedCompareExchange_rel(volatile unsigned int *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</pre>	<p>Maps to the <code>cmpxchg4.rel</code> instruction with appropriate setup. Atomically compare and exchange the value specified by the first argument (a 64-bit pointer).</p>
<pre>unsigned __int64 __InterlockedCompareExchange_acq(volatile unsigned int *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</pre>	<p>Same as the previous intrinsic, but map the <code>cmpxchg4.acq</code> instruction.</p>
<pre>void __ReleaseSpinLock(volatile int *x)</pre>	<p>Release spin lock.</p>
<pre>__int64 __InterlockedIncrement64(volatile __int64 *addend)</pre>	<p>Increment by one the value specified by its argument. Maps to the <code>fetchadd</code> instruction.</p>
<pre>__int64 __InterlockedDecrement64(volatile __int64 *addend)</pre>	<p>Decrement by one the value specified by its argument. Maps to the <code>fetchadd</code> instruction.</p>
<pre>__int64 __InterlockedExchange64(volatile __int64 *Target, __int64 value)</pre>	<p>Do an exchange operation atomically. Maps to the <code>xchg</code> instruction.</p>

Intrinsic	Description
<code>unsigned __int64 _InterlockedExchangeU64(volatile unsigned __int64 *Target, unsigned __int64 value)</code>	Same as <code>InterlockedExchange64</code> (for unsigned quantities).
<code>unsigned __int64 InterlockedCompareExchange64 rel(volatile unsigned __int64 *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</code>	Maps to the <code>cmpxchg.rel</code> instruction with appropriate setup. Atomically compare and exchange the value specified by the first argument (a 64-bit pointer).
<code>unsigned __int64 InterlockedCompareExchange64 acq(volatile unsigned __int64 *Destination, unsigned __int64 Exchange, unsigned __int64 Comparand)</code>	Maps to the <code>cmpxchg.acq</code> instruction with appropriate setup. Atomically compare and exchange the value specified by the first argument (a 64-bit pointer).
<code>__int64 _InterlockedCompareExchange64(volatile __int64 *Destination, __int64 Exchange, __int64 Comparand)</code>	Same as the previous intrinsic for signed quantities.
<code>__int64 InterlockedExchangeAdd64(volatile __int64 *addend, __int64 increment)</code>	Use compare and exchange to do an atomic add of the increment value to the addend. Maps to a loop with the <code>cmpxchg</code> instruction to guarantee atomicity
<code>__int64 InterlockedAdd64(volatile __int64 *addend, __int64 increment);</code>	Same as the previous intrinsic, but returns the new value, not the original value. See Note.

 **Note**

`_InterlockedSub64` is provided as a macro definition based on `_InterlockedAdd64`.

```
#define _InterlockedSub64(target, incr) _InterlockedAdd64((target), (-  
(incr))).
```

Uses `cmpxchg` to do an atomic sub of the `incr` value to the `target`. Maps to a loop with the `cmpxchg` instruction to guarantee atomicity.

Load and Store

You can use the load and store intrinsic to force the strict memory access ordering of specific data objects. This intended use is for the case when the user suppresses the strict memory access ordering by using the `-serialize-volatile-` option.

Intrinsic	Prototype	Description
<code>__st1_rel</code>	<code>void __st1_rel(void *dst, const char value);</code>	Generates an <code>st1.rel</code> instruction.
<code>__st2_rel</code>	<code>void __st2_rel(void *dst, const short value);</code>	Generates an <code>st2.rel</code> instruction.
<code>__st4_rel</code>	<code>void __st4_rel(void *dst, const int value);</code>	Generates an <code>st4.rel</code> instruction.
<code>__st8_rel</code>	<code>void __st8_rel(void *dst, const __int64 value);</code>	Generates an <code>st8.rel</code> instruction.
<code>__ld1_acq</code>	<code>unsigned char __ld1_acq(void *src);</code>	Generates an <code>ld1.acq</code> instruction.
<code>__ld2_acq</code>	<code>unsigned short __ld2_acq(void *src);</code>	Generates an <code>ld2.acq</code> instruction.
<code>__ld4_acq</code>	<code>unsigned int __ld4_acq(void *src);</code>	Generates an <code>ld4.acq</code> instruction.
<code>__ld8_acq</code>	<code>unsigned __int64 __ld8_acq(void *src);</code>	Generates an <code>ld8.acq</code> instruction.

Operating System Related Intrinsics

The prototypes for these intrinsics are in the `ia64intrin.h` header file.

Intrinsic	Description
<code>unsigned __int64 __getReg(const int whichReg)</code>	Gets the value from a hardware register based on the index passed in. Produces a corresponding <code>mov = r</code> instruction. Provides access to the following registers: See Register Names for <code>getReg()</code> and <code>setReg()</code> .
<code>void __setReg(const int whichReg, unsigned __int64 value)</code>	Sets the value for a hardware register based on the index passed in. Produces a corresponding <code>mov = r</code> instruction. See Register Names for <code>getReg()</code> and <code>setReg()</code> .
<code>unsigned __int64 __getIndReg(const int whichIndReg, __int64 index)</code>	Return the value of an indexed register. The index is the 2nd argument; the register file is the first argument.

Intrinsic	Description
void __setIndReg(const int whichIndReg, __int64 index, unsigned __int64 value)	Copy a value in an indexed register. The index is the 2nd argument; the register file is the first argument.
void *__ptr64 _rdteb(void)	Gets TEB address. The TEB address is kept in r13 and maps to the move r=tp instruction
void __isrlz(void)	Executes the serialize instruction. Maps to the srlz.i instruction.
void __dsrlz(void)	Serializes the data. Maps to the srlz.d instruction.
unsigned __int64 __fetchadd4_acq(unsigned int *addend, const int increment)	Map the fetchadd4.acq instruction.
unsigned __int64 __fetchadd4_rel(unsigned int *addend, const int increment)	Map the fetchadd4.rel instruction.
unsigned __int64 __fetchadd8_acq(unsigned __int64 *addend, const int increment)	Map the fetchadd8.acq instruction.
unsigned __int64 __fetchadd8_rel(unsigned __int64 *addend, const int increment)	Map the fetchadd8.rel instruction.
void __fwb(void)	Flushes the write buffers. Maps to the fwb instruction.
void __ldfs(const int whichFloatReg, void *src)	Map the ldfs instruction. Load a single precision value to the specified register.
void __ldfd(const int whichFloatReg, void *src)	Map the ldfd instruction. Load a double precision value to the specified register.
void __ldfe(const int whichFloatReg, void *src)	Map the ldfe instruction. Load an extended precision value to the specified register.
void __ldf8(const int whichFloatReg, void *src)	Map the ldf8 instruction.
void __ldf_fill(const int whichFloatReg, void *src)	Map the ldf.fill instruction.
void __stfs(void *dst, const int whichFloatReg)	Map the stfs instruction.
void __stfd(void *dst, const int whichFloatReg)	Map the stfd instruction.
void __stfe(void *dst, const int whichFloatReg)	Map the stfe instruction.

Intrinsic	Description
<code>void __stf8(void *dst, const int whichFloatReg)</code>	Map the <code>stf8</code> instruction.
<code>void __stf_spill(void *dst, const int whichFloatReg)</code>	Map the <code>stf.spill</code> instruction.
<code>void __mf(void)</code>	Executes a memory fence instruction. Maps to the <code>mf</code> instruction.
<code>void __mfa(void)</code>	Executes a memory fence, acceptance form instruction. Maps to the <code>mf.a</code> instruction.
<code>void __synci(void)</code>	Enables memory synchronization. Maps to the <code>sync.i</code> instruction.
<code>void __thash(__int64)</code>	Generates a translation hash entry address. Maps to the <code>thash r = r</code> instruction.
<code>void __ttag(__int64)</code>	Generates a translation hash entry tag. Maps to the <code>ttag r=r</code> instruction.
<code>void __itcd(__int64 pa)</code>	Insert an entry into the data translation cache (Map <code>itc.d</code> instruction).
<code>void __itci(__int64 pa)</code>	Insert an entry into the instruction translation cache (Map <code>itc.i</code>).
<code>void __itr_d(__int64 whichTransReg, __int64 pa)</code>	Map the <code>itr.d</code> instruction.
<code>void __itri(__int64 whichTransReg, __int64 pa)</code>	Map the <code>itr.i</code> instruction.
<code>void __ptce(__int64 va)</code>	Map the <code>ptc.e</code> instruction.
<code>void __ptcl(__int64 va, __int64 pagesz)</code>	Purges the local translation cache. Maps to the <code>ptc.l r, r</code> instruction.
<code>void __ptcg(__int64 va, __int64 pagesz)</code>	Purges the global translation cache. Maps to the <code>ptc.g r, r</code> instruction.
<code>void __ptcga(__int64 va, __int64 pagesz)</code>	Purges the global translation cache and ALAT. Maps to the <code>ptc.ga r, r</code> instruction.
<code>void __ptri(__int64 va, __int64 pagesz)</code>	Purges the translation register. Maps to the <code>ptr.i r, r</code> instruction.
<code>void __ptrd(__int64 va, __int64 pagesz)</code>	Purges the translation register. Maps to the <code>ptr.d r, r</code> instruction.
<code>__int64 __tpa(__int64 va)</code>	Map the <code>tpa</code> instruction.
<code>void __invalat(void)</code>	Invalidates ALAT. Maps to the <code>invala</code> instruction.
<code>void __invala (void)</code>	Same as <code>void __invalat(void)</code>
<code>void __invala_gr(const int whichGeneralReg)</code>	<code>whichGeneralReg = 0-127</code>

Intrinsic	Description
void __invala_fr(const int whichFloatReg)	whichFloatReg = 0-127
void __break(const int)	Generates a break instruction with an immediate.
void __nop(const int)	Generate a nop instruction.
void __debugbreak(void)	Generates a Debug Break Instruction fault.
void __fc(__int64)	Flushes a cache line associated with the address given by the argument. Maps to the fc instruction.
void __sum(int mask)	Sets the user mask bits of PSR. Maps to the sum imm24 instruction.
void __rum(int mask)	Resets the user mask.
__int64 __ReturnAddress(void)	Get the caller's address.
void __lfetch(int lfhint, void *y)	Generate the lfetch.lfhint instruction. The value of the first argument specifies the hint type.
void __lfetch_fault(int lfhint, void *y)	Generate the lfetch.fault.lfhint instruction. The value of the first argument specifies the hint type.
void __lfetch_excl(int lfhint, void *y)	Generate the lfetch.excl.lfhint instruction. The value {0 1 2 3} of the first argument specifies the hint type.
void __lfetch_fault_excl(int lfhint, void *y)	Generate the lfetch.fault.excl.lfhint instruction. The value of the first argument specifies the hint type.
unsigned int __cacheSize(unsigned int cacheLevel)	__cacheSize(n) returns the size in bytes of the cache at level n. 1 represents the first-level cache. 0 is returned for a non-existent cache level. For example, an application may query the cache size and use it to select block sizes in algorithms that operate on matrices.
void __memory_barrier(void)	Creates a barrier across which the compiler will not schedule any data access instruction. The compiler may allocate local data in registers across a memory barrier, but not global data.
void __ssm(int mask)	Sets the system mask. Maps to the ssm imm24 instruction.
void __rsm(int mask)	Resets the system mask bits of PSR. Maps to the rsm imm24 instruction.

Conversion Ininsics

The prototypes for these intrinsics are in the `ia64intrin.h` header file.

Intrinsic	Description
<code>__int64 _m_to_int64(__m64 a)</code>	Convert <code>a</code> of type <code>__m64</code> to type <code>__int64</code> . Translates to <code>nop</code> since both types reside in the same register on Itanium-based systems.
<code>__m64 _m_from_int64(__int64 a)</code>	Convert <code>a</code> of type <code>__int64</code> to type <code>__m64</code> . Translates to <code>nop</code> since both types reside in the same register on Itanium-based systems.
<code>__int64 __round_double_to_int64(double d)</code>	Convert its double precision argument to a signed integer.
<code>unsigned __int64 __getf_exp(double d)</code>	Map the <code>getf.exp</code> instruction and return the 16-bit exponent and the sign of its operand.

Register Names for `getReg()` and `setReg()`

The prototypes for `getReg()` and `setReg()` intrinsics are in the `ia64regs.h` header file.

Name	whichReg
<code>_IA64_REG_IP</code>	1016
<code>_IA64_REG_PSR</code>	1019
<code>_IA64_REG_PSR_L</code>	1019

General Integer Registers

Name	whichReg
<code>_IA64_REG_GP</code>	1025
<code>_IA64_REG_SP</code>	1036
<code>_IA64_REG_TP</code>	1037

Application Registers

Name	whichReg
_IA64_REG_AR_KR0	3072
_IA64_REG_AR_KR1	3073
_IA64_REG_AR_KR2	3074
_IA64_REG_AR_KR3	3075
_IA64_REG_AR_KR4	3076
_IA64_REG_AR_KR5	3077
_IA64_REG_AR_KR6	3078
_IA64_REG_AR_KR7	3079
_IA64_REG_AR_RSC	3088
_IA64_REG_AR_BSP	3089
_IA64_REG_AR_BSPSTORE	3090
_IA64_REG_AR_RNAT	3091
_IA64_REG_AR_FCR	3093
_IA64_REG_AR_EFLAG	3096
_IA64_REG_AR_CSD	3097
_IA64_REG_AR_SSD	3098
_IA64_REG_AR_CFLAG	3099
_IA64_REG_AR_FSR	3100
_IA64_REG_AR_FIR	3101
_IA64_REG_AR_FDR	3102
_IA64_REG_AR_CCV	3104
_IA64_REG_AR_UNAT	3108
_IA64_REG_AR_FPSR	3112
_IA64_REG_AR_ITC	3116
_IA64_REG_AR_PFS	3136
_IA64_REG_AR_LC	3137
_IA64_REG_AR_EC	3138

Control Registers

Name	whichReg
_IA64_REG_CR_DCR	4096
_IA64_REG_CR_ITM	4097
_IA64_REG_CR_IVA	4098
_IA64_REG_CR_PTA	4104
_IA64_REG_CR_IPSR	4112
_IA64_REG_CR_ISR	4113
_IA64_REG_CR_IIP	4115
_IA64_REG_CR_IFA	4116
_IA64_REG_CR_ITIR	4117
_IA64_REG_CR_IIPA	4118
_IA64_REG_CR_IFS	4119
_IA64_REG_CR_IIM	4120
_IA64_REG_CR_IHA	4121
_IA64_REG_CR_LID	4160
_IA64_REG_CR_IVR	4161 *
_IA64_REG_CR_TPR	4162
_IA64_REG_CR_EOI	4163
_IA64_REG_CR_IRR0	4164 *
_IA64_REG_CR_IRR1	4165 *
_IA64_REG_CR_IRR2	4166 *
_IA64_REG_CR_IRR3	4167 *
_IA64_REG_CR_ITV	4168
_IA64_REG_CR_PMV	4169
_IA64_REG_CR_CMCV	4170
_IA64_REG_CR_LRR0	4176
_IA64_REG_CR_LRR1	4177

* getReg only

Indirect Registers for getIndReg() and setIndReg()

Name	whichReg
_IA64_REG_INDR_CPUID	9000 *
_IA64_REG_INDR_DBR	9001
_IA64_REG_INDR_IBR	9002
_IA64_REG_INDR_PKR	9003
_IA64_REG_INDR_PMC	9004
_IA64_REG_INDR_PMD	9005
_IA64_REG_INDR_RR	9006
_IA64_REG_INDR_RESERVED	9007

* getIndReg only

Multimedia Additions

The prototypes for these intrinsics are in the `ia64intrin.h` header file.

Intrinsic	Corresponding Instruction
<code>__int64 _m64_czx1l(__m64 a)</code>	<code>cxz1.l</code> (Compute Zero Index)
<code>__int64 _m64_czx1r(__m64 a)</code>	<code>cxz1.r</code> (Compute Zero Index)
<code>__int64 _m64_czx2l(__m64 a)</code>	<code>cxz2.l</code> (Compute Zero Index)
<code>__int64 _m64_czx2r(__m64 a)</code>	<code>cxz2.r</code> (Compute Zero Index)
<code>__m64 _m64_mix1l(__m64 a, __m64 b)</code>	<code>mix1.l</code> (Mix)
<code>__m64 _m64_mix1r(__m64 a, __m64 b)</code>	<code>mix1.r</code> (Mix)
<code>__m64 _m64_mix2l(__m64 a, __m64 b)</code>	<code>mix2.l</code> (Mix)
<code>__m64 _m64_mix2r(__m64 a, __m64 b)</code>	<code>mix2.r</code> (Mix)
<code>__m64 _m64_mix4l(__m64 a, __m64 b)</code>	<code>mix4.l</code> (Mix)
<code>__m64 _m64_mix4r(__m64 a, __m64 b)</code>	<code>mix4.r</code> (Mix)
<code>__m64 _m64_mux1(__m64 a, const int n)</code>	<code>mux1</code> (Mux)
<code>__m64 _m64_mux2(__m64 a, const int n)</code>	<code>mux2</code> (Mux)
<code>__m64 _m64_paddluus(__m64 a, __m64 b)</code>	<code>padd1.uus</code> (Parallel add)
<code>__m64 _m64_padd2uus(__m64 a, __m64 b)</code>	<code>padd2.uus</code> (Parallel add)
<code>__m64 _m64_pavg1_nraz(__m64 a, __m64 b)</code>	<code>pavg1</code> (Parallel average)

Intrinsic	Corresponding Instruction
<code>__m64 __m64_pavg2_nraz(__m64 a, __m64 b)</code>	<code>pavg2</code> (Parallel average)
<code>__m64 __m64_pavgsub1(__m64 a, __m64 b)</code>	<code>pavgsub1</code> (Parallel average subtract)
<code>__m64 __m64_pavgsub2(__m64 a, __m64 b)</code>	<code>pavgsub2</code> (Parallel average subtract)
<code>__m64 __m64_pmpy2r(__m64 a, __m64 b)</code>	<code>pmpy2.r</code> (Parallel multiply)
<code>__m64 __m64_pmpy2l(__m64 a, __m64 b)</code>	<code>pmpy2.l</code> (Parallel multiply)
<code>__m64 __m64_pmpyshr2(__m64 a, __m64 b, const int count)</code>	<code>pmpyshr2</code> (Parallel multiply and shift right)
<code>__m64 __m64_pmpyshr2u(__m64 a, __m64 b, const int count)</code>	<code>pmpyshr2.u</code> (Parallel multiply and shift right)
<code>__m64 __m64_pshladd2(__m64 a, const int count, __m64 b)</code>	<code>pshladd2</code> (Parallel shift left and add)
<code>__m64 __m64_pshradd2(__m64 a, const int count, __m64 b)</code>	<code>pshradd2</code> (Parallel shift right and add)
<code>__m64 __m64_psubluus(__m64 a, __m64 b)</code>	<code>psub1.uus</code> (Parallel subtract)
<code>__m64 __m64_psub2uus(__m64 a, __m64 b)</code>	<code>psub2.uus</code> (Parallel subtract)

```
__int64 __m64_czx1l(__m64 a)
```

The 64-bit value `a` is scanned for a zero element from the most significant element to the least significant element, and the index of the first zero element is returned. The element width is 8 bits, so the range of the result is from 0 - 7. If no zero element is found, the default result is 8.

```
__int64 __m64_czx1r(__m64 a)
```

The 64-bit value `a` is scanned for a zero element from the least significant element to the most significant element, and the index of the first zero element is returned. The element width is 8 bits, so the range of the result is from 0 - 7. If no zero element is found, the default result is 8.

```
__int64 __m64_czx2l(__m64 a)
```

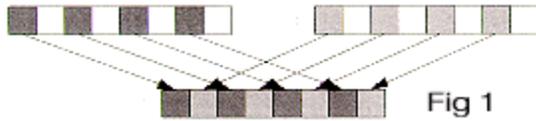
The 64-bit value `a` is scanned for a zero element from the most significant element to the least significant element, and the index of the first zero element is returned. The element width is 16 bits, so the range of the result is from 0 - 3. If no zero element is found, the default result is 4.

```
__int64 __m64_czx2r(__m64 a)
```

The 64-bit value `a` is scanned for a zero element from the least significant element to the most significant element, and the index of the first zero element is returned. The element width is 16 bits, so the range of the result is from 0 - 3. If no zero element is found, the default result is 4.

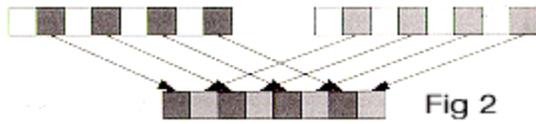
```
__m64 _m64_mix1l(__m64 a, __m64 b)
```

Interleave 64-bit quantities a and b in 1-byte groups, starting from the left, as shown in Figure 1, and return the result.



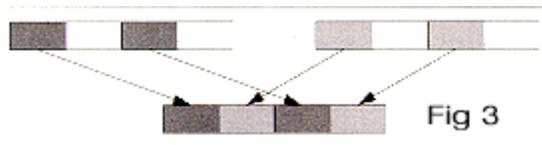
```
__m64 _m64_mix1r(__m64 a, __m64 b)
```

Interleave 64-bit quantities a and b in 1-byte groups, starting from the right, as shown in Figure 2, and return the result.



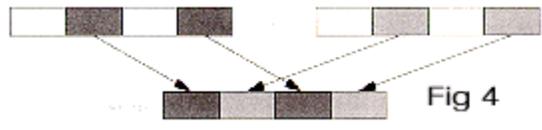
```
__m64 _m64_mix2l(__m64 a, __m64 b)
```

Interleave 64-bit quantities a and b in 2-byte groups, starting from the left, as shown in Figure 3, and return the result.



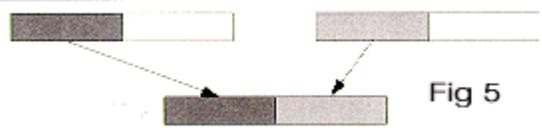
```
__m64 _m64_mix2r(__m64 a, __m64 b)
```

Interleave 64-bit quantities a and b in 2-byte groups, starting from the right, as shown in Figure 4, and return the result.



```
__m64 _m64_mix4l(__m64 a, __m64 b)
```

Interleave 64-bit quantities a and b in 4-byte groups, starting from the left, as shown in Figure 5, and return the result.



```
__m64 _m64_mix4r(__m64 a, __m64 b)
```

Interleave 64-bit quantities a and b in 4-byte groups, starting from the right, as shown in Figure 6, and return the result.

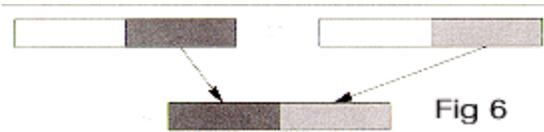


Fig 6

```
__m64 _m64_mux1(__m64 a, const int n)
```

Based on the value of n, a permutation is performed on a as shown in Figure 7, and the result is returned. Table 1 shows the possible values of n.

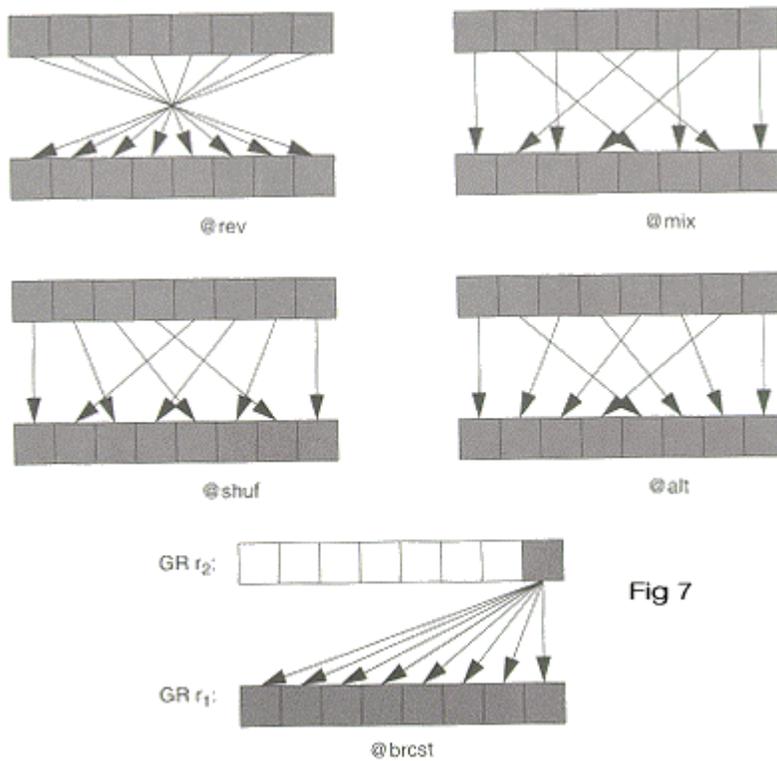


Fig 7

Values of n for m64_mux1 Operation

	n
@brcst	0
@mix	8
@shuf	9
@alt	0xA
@rev	0xB

```
__m64 __m64_mux2(__m64 a, const int n)
```

Based on the value of n, a permutation is performed on a as shown in Figure 8, and the result is returned.

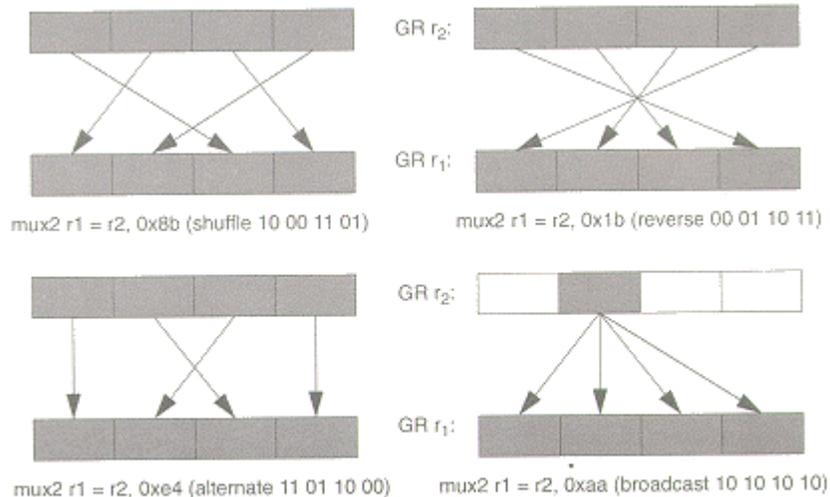


Fig 8

```
__m64 __m64_pavgsub1(__m64 a, __m64 b)
```

The unsigned data elements (bytes) of b are subtracted from the unsigned data elements (bytes) of a and the results of the subtraction are then each independently shifted to the right by one position. The high-order bits of each element are filled with the borrow bits of the subtraction.

```
__m64 __m64_pavgsub2(__m64 a, __m64 b)
```

The unsigned data elements (double bytes) of b are subtracted from the unsigned data elements (double bytes) of a and the results of the subtraction are then each independently shifted to the right by one position. The high-order bits of each element are filled with the borrow bits of the subtraction.

```
__m64 __m64_pmpy2l(__m64 a, __m64 b)
```

Two signed 16-bit data elements of a, starting with the most significant data element, are multiplied by the corresponding two signed 16-bit data elements of b, and the two 32-bit results are returned as shown in Figure 9.

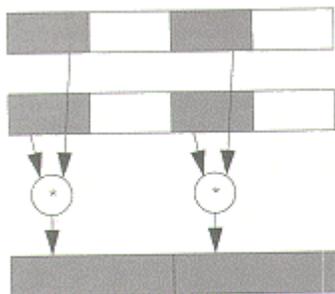


Fig 9

```
__m64 __m64_pmpy2r(__m64 a, __m64 b)
```

Two signed 16-bit data elements of *a*, starting with the least significant data element, are multiplied by the corresponding two signed 16-bit data elements of *b*, and the two 32-bit results are returned as shown in Figure 10.

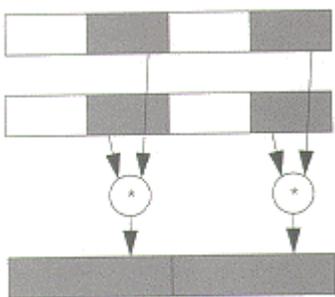


Fig 10

```
__m64 __m64_pmpyshr2(__m64 a, __m64 b, const int count)
```

The four signed 16-bit data elements of *a* are multiplied by the corresponding signed 16-bit data elements of *b*, yielding four 32-bit products. Each product is then shifted to the right *count* bits and the least significant 16 bits of each shifted product form 4 16-bit results, which are returned as one 64-bit word.

```
__m64 __m64_pmpyshr2u(__m64 a, __m64 b, const int count)
```

The four unsigned 16-bit data elements of *a* are multiplied by the corresponding unsigned 16-bit data elements of *b*, yielding four 32-bit products. Each product is then shifted to the right *count* bits and the least significant 16 bits of each shifted product form 4 16-bit results, which are returned as one 64-bit word.

```
__m64 __m64_pshladd2(__m64 a, const int count, __m64 b)
```

a is shifted to the left by *count* bits and then is added to *b*. The upper 32 bits of the result are forced to 0, and then bits [31:30] of *b* are copied to bits [62:61] of the result. The result is returned.

```
__m64 __m64_pshradd2(__m64 a, const int count, __m64 b)
```

The four signed 16-bit data elements of *a* are each independently shifted to the right by *count* bits (the high order bits of each element are filled with the initial value of the sign bits of the data elements in *a*); they are then added to the four signed 16-bit data elements of *b*. The result is returned.

```
__m64 _m64_padd1uus(__m64 a, __m64 b)
```

a is added to b as eight separate byte-wide elements. The elements of a are treated as unsigned, while the elements of b are treated as signed. The results are treated as unsigned and are returned as one 64-bit word.

```
__m64 _m64_padd2uus(__m64 a, __m64 b)
```

a is added to b as four separate 16-bit wide elements. The elements of a are treated as unsigned, while the elements of b are treated as signed. The results are treated as unsigned and are returned as one 64-bit word.

```
__m64 _m64_psub1uus(__m64 a, __m64 b)
```

a is subtracted from b as eight separate byte-wide elements. The elements of a are treated as unsigned, while the elements of b are treated as signed. The results are treated as unsigned and are returned as one 64-bit word.

```
__m64 _m64_psub2uus(__m64 a, __m64 b)
```

a is subtracted from b as four separate 16-bit wide elements. The elements of a are treated as unsigned, while the elements of b are treated as signed. The results are treated as unsigned and are returned as one 64-bit word.

```
__m64 _m64_pavg1_nraz(__m64 a, __m64 b)
```

The unsigned byte-wide data elements of a are added to the unsigned byte-wide data elements of b and the results of each add are then independently shifted to the right by one position. The high-order bits of each element are filled with the carry bits of the sums.

```
__m64 _m64_pavg2_nraz(__m64 a, __m64 b)
```

The unsigned 16-bit wide data elements of a are added to the unsigned 16-bit wide data elements of b and the results of each add are then independently shifted to the right by one position. The high-order bits of each element are filled with the carry bits of the sums.

Synchronization Primitives

The synchronization primitive intrinsics provide a variety of operations. Besides performing these operations, each intrinsic has two key properties:

- the function performed is guaranteed to be atomic
- associated with each intrinsic are certain memory barrier properties that restrict the movement of memory references to visible data across the intrinsic operation by either the compiler or the processor

For the following intrinsics, <type> is either a 32-bit or 64-bit integer.

Atomic Fetch-and-op Operations

```
<type> __sync_fetch_and_add(<type> *ptr, <type> val)
<type> __sync_fetch_and_and(<type> *ptr, <type> val)
<type> __sync_fetch_and_nand(<type> *ptr, <type> val)
<type> __sync_fetch_and_or(<type> *ptr, <type> val)
<type> __sync_fetch_and_sub(<type> *ptr, <type> val)
<type> __sync_fetch_and_xor(<type> *ptr, <type> val)
```

Atomic Op-and-fetch Operations

```

<type> __sync_add_and_fetch(<type> *ptr,<type> val)
<type> __sync_sub_and_fetch(<type> *ptr,<type> val)
<type> __sync_or_and_fetch(<type> *ptr,<type> val)
<type> __sync_and_and_fetch(<type> *ptr,<type> val)
<type> __sync_nand_and_fetch(<type> *ptr,<type> val)
<type> __sync_xor_and_fetch(<type> *ptr,<type> val)

```

Atomic Compare-and-swap Operations

```

<type> __sync_val_compare_and_swap(<type> *ptr, <type> old_val, <type>
new_val)
int __sync_bool_compare_and_swap(<type> *ptr, <type> old_val, <type>
new_val)

```

Atomic Synchronize Operation

```
void __sync_synchronize (void);
```

Atomic Lock-test-and-set Operation

```
<type> __sync_lock_test_and_set(<type> *ptr,<type> val)
```

Atomic Lock-release Operation

```
void __sync_lock_release(<type> *ptr)
```

Miscellaneous Intrinsics

```
void* __get_return_address(unsigned int level);
```

This intrinsic yields the return address of the current function. The `level` argument must be a constant value. A value of 0 yields the return address of the current function. Any other value yields a zero return address. On Linux systems, this intrinsic is synonymous with `__builtin_return_address`. The name and the argument are provided for compatibility with gcc*.

```
void __set_return_address(void* addr);
```

This intrinsic overwrites the default return address of the current function with the address indicated by its argument. On return from the current invocation, program execution continues at the address provided.

```
void* __get_frame_address(unsigned int level);
```

This intrinsic returns the frame address of the current function. The `level` argument must be a constant value. A value of 0 yields the frame address of the current function. Any other value yields a zero return value. On Linux systems, this intrinsic is synonymous with `__builtin_frame_address`. The name and the argument are provided for compatibility with gcc.

Data Alignment, Memory Allocation Intrinsics, and Inline Assembly

This section describes features that support usage of the intrinsics. The following topics are described:

- Alignment Support
- Allocating and Freeing Aligned Memory Blocks

Alignment Support

To improve intrinsics performance, you need to align data. For example, when you are using the Streaming SIMD Extensions, you should align data to 16 bytes in memory operations to improve performance. Specifically, you must align `__m128` objects as addresses passed to the `_mm_load` and `_mm_store` intrinsics. If you want to declare arrays of floats and treat them as `__m128` objects by casting, you need to ensure that the float arrays are properly aligned.

Use `__declspec(align)` to direct the compiler to align data more strictly than it otherwise does on both IA-32 and Itanium®-based systems. For example, a data object of type `int` is allocated at a byte address which is a multiple of 4 by default (the size of an `int`). However, by using `__declspec(align)`, you can direct the compiler to instead use an address which is a multiple of 8, 16, or 32 with the following restrictions on IA-32:

- 32-byte addresses must be statically allocated
- 16-byte addresses can be locally or statically allocated

You can use this data alignment support as an advantage in optimizing cache line usage. By clustering small objects that are commonly used together into a `struct`, and forcing the `struct` to be allocated at the beginning of a cache line, you can effectively guarantee that each object is loaded into the cache as soon as any one is accessed, resulting in a significant performance benefit.

The syntax of this extended-attribute is as follows:

```
align(n)
```

where *n* is an integral power of 2, less than or equal to 32. The value specified is the requested alignment.



Caution

In this release, `__declspec(align(8))` does not function correctly. Use `__declspec(align(16))` instead.



Note

If a value is specified that is less than the alignment of the affected data type, it has no effect. In other words, data is aligned to the maximum of its own alignment or the alignment specified with `__declspec(align)`.

You can request alignments for individual variables, whether of static or automatic storage duration. (Global and static variables have static storage duration; local variables have automatic storage duration by default.) You cannot adjust the alignment of a parameter, nor a field of a `struct` or `class`. You can, however, increase the alignment of a `struct` (or union or `class`), in which case every object of that type is affected.

As an example, suppose that a function uses local variables `i` and `j` as subscripts into a 2-dimensional array. They might be declared as follows:

```
int i, j;
```

These variables are commonly used together. But they can fall in different cache lines, which could be detrimental to performance. You can instead declare them as follows:

```
__declspec(align(8)) struct { int i, j; } sub;
```

The compiler now ensures that they are allocated in the same cache line. In C++, you can omit the `struct` variable name (written as `sub` in the previous example). In C, however, it is required, and you must write references to `i` and `j` as `sub.i` and `sub.j`.

If you use many functions with such subscript pairs, it is more convenient to declare and use a `struct` type for them, as in the following example:

```
typedef struct __declspec(align(8)) { int i, j; } Sub;
```

By placing the `__declspec(align)` after the keyword `struct`, you are requesting the appropriate alignment for all objects of that type. However, that allocation of parameters is unaffected by `__declspec(align)`. (If necessary, you can assign the value of a parameter to a local variable with the appropriate alignment.)

You can also force alignment of global variables, such as arrays:

```
__declspec(align(16)) float array[1000];
```

Allocating and Freeing Aligned Memory Blocks

Use the `_mm_malloc` and `_mm_free` intrinsics to allocate and free aligned blocks of memory. These intrinsics are based on `malloc` and `free`, which are in the `libirc.a` library. You need to include `malloc.h`. The syntax for these intrinsics is as follows:

```
void* _mm_malloc (int size, int align)
```

```
void _mm_free (void *p)
```

The `_mm_malloc` routine takes an extra parameter, which is the alignment constraint. This constraint must be a power of two. The pointer that is returned from `_mm_malloc` is guaranteed to be aligned on the specified boundary.



Note

Memory that is allocated using `_mm_malloc` must be freed using `_mm_free`. Calling `free` on memory allocated with `_mm_malloc` or calling `_mm_free` on memory allocated with `malloc` will cause unpredictable behavior.

Inline Assembly

By default, the compiler inlines a number of standard C, C++, and math library functions. This usually results in faster execution of your program.

Sometimes inline expansion of library functions can cause unexpected results. The inlined library functions do not set the `errno` variable. So, in code that relies upon the setting of the `errno` variable, you should use the `-nolib_inline` option, which turns off inline expansion of library functions. Also, if one of your functions has the same name as one of the compiler's supplied library functions, the compiler assumes that it is one of the latter and replaces the call with the inlined version. Consequently, 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 program's function is the one used.



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.c` without expanding the library functions, but with inline expansion from interprocedural optimizations (IPO):

```
prompt>icpc -ip -nolib_inline sum.cpp
```

For details on IPO, see Interprocedural Optimizations.

MASM* Style Inline Assembly

The Intel® C++ Compiler supports MASM style inline assembly with the `-use_msasm` option. See your MASM documentation for the proper syntax.

GNU*-like Style Inline Assembly (IA-32 only)

The Intel® C++ Compiler supports GNU-like style inline assembly. The syntax is as follows:

```
asm-keyword [ volatile-keyword ] ( asm-template [ asm-interface ] ) ;
```



Caution

Under the `-use_msasm` compilation flag, Gnu asm aliases will only work if you use the `__asm__` keyword, they will not work correctly if you use the alternate `__asm` or `asm` keywords.

Syntax Element	Description
<code>asm-keyword</code>	<code>asm</code> statements begin with the keyword <code>asm</code> . Alternatively, either <code>__asm</code> or <code>__asm__</code> may be used for compatibility. See Caution statement.
<code>volatile-keyword</code>	If the optional keyword <code>volatile</code> is given, the <code>asm</code> is volatile. Two <code>volatile asm</code> statements will never be moved past each other, and a reference to a <code>volatile</code> variable will not be moved relative to a <code>volatile asm</code> . Alternate keywords <code>__volatile</code> and <code>__volatile__</code> may be used for compatibility.

Syntax Element	Description
asm-template	The <code>asm-template</code> is a C language ASCII string which specifies how to output the assembly code for an instruction. Most of the template is a fixed string; everything but the substitution-directives, if any, is passed through to the assembler. The syntax for a substitution directive is a <code>%</code> followed by one or two characters. The supported substitution directives are specified in a subsequent section.
asm-interface	The <code>asm-interface</code> consists of three parts: 1. an optional <code>output-list</code> 2. an optional <code>input-list</code> 3. an optional <code>clobber-list</code> These are separated by colon (<code>:</code>) characters. If the <code>output-list</code> is missing, but an <code>input-list</code> is given, the input list may be preceded by two colons (<code>::</code>) to take the place of the missing <code>output-list</code> . If the <code>asm-interface</code> is omitted altogether, the <code>asm</code> statement is considered <code>volatile</code> regardless of whether a <code>volatile-keyword</code> was specified.
output-list	An <code>output-list</code> consists of one or more <code>output-specs</code> separated by commas. For the purposes of substitution in the <code>asm-template</code> , each <code>output-spec</code> is numbered. The first operand in the <code>output-list</code> is numbered 0, the second is 1, and so on. Numbering is continuous through the <code>output-list</code> and into the <code>input-list</code> . The total number of operands is limited to 10 (i.e. 0-9).
input-list	Similar to an <code>output-list</code> , an <code>input-list</code> consists of one or more <code>input-specs</code> separated by commas. For the purposes of substitution in the <code>asm-template</code> , each <code>input-spec</code> is numbered, with the numbers continuing from those in the <code>output-list</code> .
clobber-list	A <code>clobber-list</code> tells the compiler that the <code>asm</code> uses or changes a specific machine register that is either coded directly into the <code>asm</code> or is changed implicitly by the assembly instruction. The <code>clobber-list</code> is a comma-separated list of <code>clobber-specs</code> .
input-spec	The <code>input-specs</code> tell the compiler about expressions whose values may be needed by the inserted assembly instruction. In order to describe fully the input requirements of the <code>asm</code> , you can list <code>input-specs</code> that are not actually referenced in the <code>asm-template</code> .
clobber-spec	Each <code>clobber-spec</code> specifies the name of a single machine register that is clobbered. The register name may optionally be preceded by a <code>%</code> . The following are the valid register names: <code>eax</code> , <code>ebx</code> , <code>ecx</code> , <code>edx</code> , <code>esi</code> , <code>edi</code> , <code>ebp</code> , <code>esp</code> , <code>ax</code> , <code>bx</code> , <code>cx</code> , <code>dx</code> , <code>si</code> , <code>di</code> , <code>bp</code> , <code>sp</code> , <code>al</code> , <code>bl</code> , <code>cl</code> , <code>dl</code> , <code>ah</code> , <code>bh</code> , <code>ch</code> , <code>dh</code> , <code>st</code> , <code>st(1) - st(7)</code> , <code>mm0 - mm7</code> , <code>xmm0 - xmm7</code> , and <code>cc</code> . It is also legal to specify "memory" in a <code>clobber-spec</code> . This prevents the compiler from keeping data cached in registers across the <code>asm</code> statement.

Intrinsics Cross-processor Implementation

This section provides a series of tables that compare intrinsics performance across architectures. Before implementing intrinsics across architectures, please note the following.

- Intrinsics may generate code that does not run on all IA processors. Therefore the programmer is responsible for using `CPUID` to detect the processor and generating the appropriate code.
- Implement intrinsics by processor family, not by specific processor. The guiding principle for which family -- IA-32 or Itanium® processors -- the intrinsic is implemented on is performance, not compatibility. Where there is added performance on both families, the intrinsic will be identical.

Intrinsics For Implementation Across All IA

The following intrinsics provide significant performance gain over a non-intrinsic-based code equivalent.

<code>int abs(int)</code>
<code>long labs(long)</code>
<code>unsigned long __lrotl(unsigned long value, int shift)</code>
<code>unsigned long __lrotr(unsigned long value, int shift)</code>
<code>unsigned int __rotl(unsigned int value, int shift)</code>
<code>unsigned int __rotr(unsigned int value, int shift)</code>
<code>__int64 __i64_rotl(__int64 value, int shift)</code>
<code>__int64 __i64_rotr(__int64 value, int shift)</code>
<code>double fabs(double)</code>
<code>double log(double)</code>
<code>float logf(float)</code>
<code>double log10(double)</code>
<code>float log10f(float)</code>
<code>double exp(double)</code>
<code>float expf(float)</code>
<code>double pow(double, double)</code>
<code>float powf(float, float)</code>
<code>double sin(double)</code>
<code>float sinf(float)</code>
<code>double cos(double)</code>
<code>float cosf(float)</code>
<code>double tan(double)</code>
<code>float tanf(float)</code>

<code>double acos(double)</code>
<code>float acosf(float)</code>
<code>double acosh(double)</code>
<code>float acoshf(float)</code>
<code>double asin(double)</code>
<code>float asinf(float)</code>
<code>double asinh(double)</code>
<code>float asinhf(float)</code>
<code>double atan(double)</code>
<code>float atanf(float)</code>
<code>double atanh(double)</code>
<code>float atanhf(float)</code>
<code>float cabs(double)*</code>
<code>double ceil(double)</code>
<code>float ceilf(float)</code>
<code>double cosh(double)</code>
<code>float coshf(float)</code>
<code>float fabsf(float)</code>
<code>double floor(double)</code>
<code>float floorf(float)</code>
<code>double fmod(double)</code>
<code>float fmodf(float)</code>
<code>double hypot(double, double)</code>
<code>float hypotf(float)</code>
<code>double rint(double)</code>
<code>float rintf(float)</code>
<code>double sinh(double)</code>
<code>float sinhf(float)</code>
<code>float sqrtf(float)</code>
<code>double tanh(double)</code>
<code>float tanhf(float)</code>
<code>char *_strset(char *, _int32)</code>
<code>void *memcmp(const void *cs, const void *ct, size_t n)</code>
<code>void *memcpy(void *s, const void *ct, size_t n)</code>

<code>void *memset(void * s, int c, size_t n)</code>
<code>char *Strcat(char * s, const char * ct)</code>
<code>int *strcmp(const char *, const char *)</code>
<code>char *strcpy(char * s, const char * ct)</code>
<code>size_t strlen(const char * cs)</code>
<code>int strncmp(char *, char *, int)</code>
<code>int strncpy(char *, char *, int)</code>
<code>void *__alloca(int)</code>
<code>int _setjmp(jmp_buf)</code>
<code>_exception_code(void)</code>
<code>_exception_info(void)</code>
<code>_abnormal_termination(void)</code>
<code>void _enable()</code>
<code>void _disable()</code>
<code>int _bswap(int)</code>
<code>int _in_byte(int)</code>
<code>int _in_dword(int)</code>
<code>int _in_word(int)</code>
<code>int _inp(int)</code>
<code>int _inpd(int)</code>
<code>int _inpw(int)</code>
<code>int _out_byte(int, int)</code>
<code>int _out_dword(int, int)</code>
<code>int _out_word(int, int)</code>
<code>int _outp(int, int)</code>
<code>int _outpd(int, int)</code>
<code>int _outpw(int, int)</code>

MMX(TM) Technology Intrinsic Implementation

Key to the table entries

- A = Expected to give significant performance gain over non-intrinsic-based code equivalent.
- B = Non-intrinsic-based source code would be better; the intrinsic's implementation may map directly to native instructions, but they offer no significant performance gain.
- C = Requires contorted implementation for particular microarchitecture. Will result in very poor performance if used.

Intrinsic Name	Alternate Name	Across All IA	MMX(TM) Technology Streaming SIMD Extensions Streaming SIMD Extensions 2	Itanium® Architecture
<code>_m_empty</code>	<code>_mm_empty</code>	N/A	A	B
<code>_m_from_int</code>	<code>_mm_cvtsi32_si64</code>	N/A	A	A
<code>_m_to_int</code>	<code>_mm_cvtsi64_si32</code>	N/A	A	A
<code>_m_packsswb</code>	<code>_mm_packs_pi16</code>	N/A	A	A
<code>_m_packssdw</code>	<code>_mm_packs_pi32</code>	N/A	A	A
<code>_m_packuswb</code>	<code>_mm_packs_pu16</code>	N/A	A	A
<code>_m_punpckhbw</code>	<code>_mm_unpackhi_pi8</code>	N/A	A	A
<code>_m_punpckhwd</code>	<code>_mm_unpackhi_pi16</code>	N/A	A	A
<code>_m_punpckhdq</code>	<code>_mm_unpackhi_pi32</code>	N/A	A	A
<code>_m_punpcklbw</code>	<code>_mm_unpacklo_pi8</code>	N/A	A	A
<code>_m_punpcklwd</code>	<code>_mm_unpacklo_pi16</code>	N/A	A	A
<code>_m_punpckldq</code>	<code>_mm_unpacklo_pi32</code>	N/A	A	A
<code>_m_paddb</code>	<code>_mm_add_pi8</code>	N/A	A	A
<code>_m_paddw</code>	<code>_mm_add_pi16</code>	N/A	A	A
<code>_m_paddd</code>	<code>_mm_add_pi32</code>	N/A	A	A
<code>_m_paddsb</code>	<code>_mm_adds_pi8</code>	N/A	A	A
<code>_m_paddsw</code>	<code>_mm_adds_pi16</code>	N/A	A	A
<code>_m_paddusb</code>	<code>_mm_adds_pu8</code>	N/A	A	A

_m_paddusw	_mm_adds_pu16	N/A	A	A
_m_psubb	_mm_sub_pi8	N/A	A	A
_m_psubw	_mm_sub_pi16	N/A	A	A
_m_psubd	_mm_sub_pi32	N/A	A	A
_m_psubsb	_mm_subs_pi8	N/A	A	A
_m_psubsw	_mm_subs_pi16	N/A	A	A
_m_psubusb	_mm_subs_pu8	N/A	A	A
_m_psubusw	_mm_subs_pu16	N/A	A	A
_m_pmaddwd	_mm_madd_pi16	N/A	A	C
_m_pmulhw	_mm_mulhi_pi16	N/A	A	A
_m_pmullw	_mm_mullo_pi16	N/A	A	A
_m_psllw	_mm_sll_pi16	N/A	A	A
_m_psllwi	_mm_slli_pi16	N/A	A	A
_m_pslld	_mm_sll_pi32	N/A	A	A
_m_pslldi	_mm_slli_pi32	N/A	A	A
_m_psllq	_mm_sll_si64	N/A	A	A
_m_psllqi	_mm_slli_si64	N/A	A	A
_m_psraw	_mm_sra_pi16	N/A	A	A
_m_psrawi	_mm_srai_pi16	N/A	A	A
_m_psrads	_mm_sra_pi32	N/A	A	A
_m_psradi	_mm_srai_pi32	N/A	A	A
_m_psrslw	_mm_srl_pi16	N/A	A	A
_m_psrslwi	_mm_srli_pi16	N/A	A	A
_m_psrld	_mm_srl_pi32	N/A	A	A
_m_psrldi	_mm_srli_pi32	N/A	A	A
_m_psrslq	_mm_srl_si64	N/A	A	A
_m_psrslqi	_mm_srli_si64	N/A	A	A
_m_pand	_mm_and_si64	N/A	A	A
_m_pandn	_mm_andnot_si64	N/A	A	A
_m_por	_mm_or_si64	N/A	A	A
_m_pxor	_mm_xor_si64	N/A	A	A

<code>_m_pcmpeqb</code>	<code>_mm_cmpeq_pi8</code>	N/A	A	A
<code>_m_pcmpeqw</code>	<code>_mm_cmpeq_pi16</code>	N/A	A	A
<code>_m_pcmpeqd</code>	<code>_mm_cmpeq_pi32</code>	N/A	A	A
<code>_m_pcmpgtb</code>	<code>_mm_cmpgt_pi8</code>	N/A	A	A
<code>_m_pcmpgtw</code>	<code>_mm_cmpgt_pi16</code>	N/A	A	A
<code>_m_pcmpgtd</code>	<code>_mm_cmpgt_pi32</code>	N/A	A	A
<code>mm_setzero_si64</code>		N/A	A	A
<code>_mm_set_pi32</code>		N/A	A	A
<code>_mm_set_pi16</code>		N/A	A	C
<code>_mm_set_pi8</code>		N/A	A	C
<code>_mm_set1_pi32</code>		N/A	A	A
<code>_mm_set1_pi16</code>		N/A	A	A
<code>_mm_set1_pi8</code>		N/A	A	A
<code>_mm_setr_pi32</code>		N/A	A	A
<code>_mm_setr_pi16</code>		N/A	A	C
<code>_mm_setr_pi8</code>		N/A	A	C

`_mm_empty` is implemented in Itanium instructions as a NOP for source compatibility only.

Streaming SIMD Extensions Intrinsic Implementation

Regular Streaming SIMD Extensions intrinsics work on 4 32-bit single precision values. On Itanium®-based systems basic operations like add or compare will require two SIMD instructions. Both can be executed in the same cycle so the throughput is one basic Streaming SIMD Extensions operation per cycle or 4 32-bit single precision operations per cycle.

Key to the table entries

- A = Expected to give significant performance gain over non-intrinsic-based code equivalent.
- B = Non-intrinsic-based source code would be better; the intrinsic's implementation may map directly to native instructions but they offer no significant performance gain.
- C = Requires contorted implementation for particular microarchitecture. Will result in very poor performance if used.

Intrinsic Name	Alternate Name	Across All IA	MMX(TM Technology	Streaming SIMD Extensions Streaming SIMD Extensions 2	Itanium® Architecture
<code>_mm_add_ss</code>		N/A	N/A	B	B
<code>_mm_add_ps</code>		N/A	N/A	A	A
<code>_mm_sub_ss</code>		N/A	N/A	B	B
<code>_mm_sub_ps</code>		N/A	N/A	A	A
<code>_mm_mul_ss</code>		N/A	N/A	B	B
<code>_mm_mul_ps</code>		N/A	N/A	A	A
<code>_mm_div_ss</code>		N/A	N/A	B	B
<code>_mm_div_ps</code>		N/A	N/A	A	A
<code>_mm_sqrt_ss</code>		N/A	N/A	B	B
<code>_mm_sqrt_ps</code>		N/A	N/A	A	A
<code>_mm_rcp_ss</code>		N/A	N/A	B	B
<code>_mm_rcp_ps</code>		N/A	N/A	A	A
<code>_mm_rsqrt_ss</code>		N/A	N/A	B	B
<code>_mm_rsqrt_ps</code>		N/A	N/A	A	A
<code>_mm_min_ss</code>		N/A	N/A	B	B
<code>_mm_min_ps</code>		N/A	N/A	A	A
<code>_mm_max_ss</code>		N/A	N/A	B	B

_mm_max_ps		N/A	N/A	A	A
_mm_and_ps		N/A	N/A	A	A
_mm_andnot_ps		N/A	N/A	A	A
_mm_or_ps		N/A	N/A	A	A
_mm_xor_ps		N/A	N/A	A	A
_mm_cmpeq_ss		N/A	N/A	B	B
_mm_cmpeq_ps		N/A	N/A	A	A
_mm_cmplt_ss		N/A	N/A	B	B
_mm_cmplt_ps		N/A	N/A	A	A
_mm_cmple_ss		N/A	N/A	B	B
_mm_cmple_ps		N/A	N/A	A	A
_mm_cmpgt_ss		N/A	N/A	B	B
_mm_cmpgt_ps		N/A	N/A	A	A
_mm_cmpge_ss		N/A	N/A	B	B
_mm_cmpge_ps		N/A	N/A	A	A
_mm_cmpneq_ss		N/A	N/A	B	B
_mm_cmpneq_ps		N/A	N/A	A	A
_mm_cmpnlt_ss		N/A	N/A	B	B
_mm_cmpnlt_ps		N/A	N/A	A	A
_mm_cmpnle_ss		N/A	N/A	B	B
_mm_cmpnle_ps		N/A	N/A	A	A
_mm_cmpngt_ss		N/A	N/A	B	B
_mm_cmpngt_ps		N/A	N/A	A	A
_mm_cmpnge_ss		N/A	N/A	B	B
_mm_cmpnge_ps		N/A	N/A	A	A
_mm_cmpord_ss		N/A	N/A	B	B
_mm_cmpord_ps		N/A	N/A	A	A
_mm_cmpunord_ss		N/A	N/A	B	B
_mm_cmpunord_ps		N/A	N/A	A	A
_mm_comieq_ss		N/A	N/A	B	B
_mm_comilt_ss		N/A	N/A	B	B

_mm_comile_ss		N/A	N/A	B	B
_mm_comigt_ss		N/A	N/A	B	B
_mm_comige_ss		N/A	N/A	B	B
_mm_comineq_ss		N/A	N/A	B	B
_mm_ucomieq_ss		N/A	N/A	B	B
_mm_ucomilt_ss		N/A	N/A	B	B
_mm_ucomile_ss		N/A	N/A	B	B
_mm_ucomigt_ss		N/A	N/A	B	B
_mm_ucomige_ss		N/A	N/A	B	B
_mm_ucomineq_ss		N/A	N/A	B	B
_mm_cvt_ss2si	_mm_cvtss_si32	N/A	N/A	A	B
_mm_cvt_ps2pi	_mm_cvtps_pi32	N/A	N/A	A	A
_mm_cvtt_ss2si	_mm_cvttss_si32	N/A	N/A	A	B
_mm_cvtt_ps2pi	_mm_cvttps_pi32	N/A	N/A	A	A
_mm_cvt_si2ss	_mm_cvtsi32_ss	N/A	N/A	A	B
_mm_cvt_pi2ps	_mm_cvtpi32_ps	N/A	N/A	A	C
_mm_cvtpi16_ps		N/A	N/A	A	C
_mm_cvtpul6_ps		N/A	N/A	A	C
_mm_cvtpi8_ps		N/A	N/A	A	C
_mm_cvtpu8_ps		N/A	N/A	A	C
_mm_cvtpi32x2_ps		N/A	N/A	A	C
_mm_cvtps_pi16		N/A	N/A	A	C
_mm_cvtps_pi8		N/A	N/A	A	C
_mm_move_ss		N/A	N/A	A	A
_mm_shuffle_ps		N/A	N/A	A	A
_mm_unpackhi_ps		N/A	N/A	A	A
_mm_unpacklo_ps		N/A	N/A	A	A
_mm_movehl_ps		N/A	N/A	A	A
_mm_movelh_ps		N/A	N/A	A	A
_mm_movemask_ps		N/A	N/A	A	C
_mm_getcsr		N/A	N/A	A	A

_mm_setcsr		N/A	N/A	A	A
_mm_loadh_pi		N/A	N/A	A	A
_mm_loadl_pi		N/A	N/A	A	A
_mm_load_ss		N/A	N/A	A	B
_mm_load_ps1	_mm_loadl_ps	N/A	N/A	A	A
_mm_load_ps		N/A	N/A	A	A
_mm_loadu_ps		N/A	N/A	A	A
_mm_loadr_ps		N/A	N/A	A	A
_mm_storeh_pi		N/A	N/A	A	A
_mm_storel_pi		N/A	N/A	A	A
_mm_store_ss		N/A	N/A	A	A
_mm_store_ps		N/A	N/A	A	A
_mm_store_ps1	_mm_storel_ps	N/A	N/A	A	A
_mm_storeu_ps		N/A	N/A	A	A
_mm_storer_ps		N/A	N/A	A	A
_mm_set_ss		N/A	N/A	A	A
_mm_set_ps1	_mm_setl_ps	N/A	N/A	A	A
_mm_set_ps		N/A	N/A	A	A
_mm_setr_ps		N/A	N/A	A	A
_mm_setzero_ps		N/A	N/A	A	A
_mm_prefetch		N/A	N/A	A	A
_mm_stream_pi		N/A	N/A	A	A
_mm_stream_ps		N/A	N/A	A	A
_mm_sfence		N/A	N/A	A	A
_m_pextrw	_mm_extract_pi16	N/A	N/A	A	A
_m_pinsrw	_mm_insert_pi16	N/A	N/A	A	A
_m_pmaxsw	_mm_max_pi16	N/A	N/A	A	A
_m_pmaxub	_mm_max_pu8	N/A	N/A	A	A
_m_pminsw	_mm_min_pi16	N/A	N/A	A	A
_m_pminub	_mm_min_pu8	N/A	N/A	A	A
_m_pmovmskb	_mm_movemask_pi8	N/A	N/A	A	C

_m_pmulhuw	_mm_mulhi_pu16	N/A	N/A	A	A
_m_pshufw	_mm_shuffle_pi16	N/A	N/A	A	A
_m_maskmovq	_mm_maskmove_si64	N/A	N/A	A	C
_m_pavgb	_mm_avg_pu8	N/A	N/A	A	A
_m_pavgw	_mm_avg_pu16	N/A	N/A	A	A
_m_psadbw	_mm_sad_pu8	N/A	N/A	A	A

Streaming SIMD Extensions 2 Intrinsic Implementation

Streaming SIMD Extensions 2 operate on 128-bit quantities with 64-bit double precision floating-point values. The Intel® Itanium® processor does not support parallel double precision computation, so Streaming SIMD Extensions 2 are not implemented on Itanium-based systems.

Key to the table entries

- A = Expected to give significant performance gain over non-intrinsic-based code equivalent.
- B = Non-intrinsic-based source code would be better; the intrinsic's implementation may map directly to native instructions, but they offer no significant performance gain.
- C = Requires contorted implementation for particular microarchitecture. Will result in very poor performance if used.

Intrinsic	Across All IA	MMX(TM) Technology	Streaming SIMD Extensions	Streaming SIMD Extensions 2	Itanium® Architecture
_mm_add_sd	N/A	N/A	N/A	A	N/A
_mm_add_pd	N/A	N/A	N/A	A	N/A
_mm_sub_sd	N/A	N/A	N/A	A	N/A
_mm_sub_pd	N/A	N/A	N/A	A	N/A
_mm_mul_sd	N/A	N/A	N/A	A	N/A
_mm_mul_pd	N/A	N/A	N/A	A	N/A
_mm_sqrt_sd	N/A	N/A	N/A	A	N/A
_mm_sqrt_pd	N/A	N/A	N/A	A	N/A
_mm_div_sd	N/A	N/A	N/A	A	N/A
_mm_div_pd	N/A	N/A	N/A	A	N/A
_mm_min_sd	N/A	N/A	N/A	A	N/A
_mm_min_pd	N/A	N/A	N/A	A	N/A
_mm_max_sd	N/A	N/A	N/A	A	N/A

_mm_max_pd	N/A	N/A	N/A	A	N/A
_mm_and_pd	N/A	N/A	N/A	A	N/A
_mm_andnot_pd	N/A	N/A	N/A	A	N/A
_mm_or_pd	N/A	N/A	N/A	A	N/A
_mm_xor_pd	N/A	N/A	N/A	A	N/A
_mm_cmpeq_sd	N/A	N/A	N/A	A	N/A
_mm_cmpeq_pd	N/A	N/A	N/A	A	N/A
_mm_cmplt_sd	N/A	N/A	N/A	A	N/A
_mm_cmplt_pd	N/A	N/A	N/A	A	N/A
_mm_cmple_sd	N/A	N/A	N/A	A	N/A
_mm_cmple_pd	N/A	N/A	N/A	A	N/A
_mm_cmpgt_sd	N/A	N/A	N/A	A	N/A
_mm_cmpgt_pd	N/A	N/A	N/A	A	N/A
_mm_cmpge_sd	N/A	N/A	N/A	A	N/A
_mm_cmpge_pd	N/A	N/A	N/A	A	N/A
_mm_cmpneq_sd	N/A	N/A	N/A	A	N/A
_mm_cmpneq_pd	N/A	N/A	N/A	A	N/A
_mm_cmpnlt_sd	N/A	N/A	N/A	A	N/A
_mm_cmpnlt_pd	N/A	N/A	N/A	A	N/A
_mm_cmpnle_sd	N/A	N/A	N/A	A	N/A
_mm_cmpnle_pd	N/A	N/A	N/A	A	N/A
_mm_cmpngt_sd	N/A	N/A	N/A	A	N/A
_mm_cmpngt_pd	N/A	N/A	N/A	A	N/A
_mm_cmpnge_sd	N/A	N/A	N/A	A	N/A
_mm_cmpnge_pd	N/A	N/A	N/A	A	N/A
_mm_cmpord_pd	N/A	N/A	N/A	A	N/A
_mm_cmpord_sd	N/A	N/A	N/A	A	N/A
_mm_cmpunord_pd	N/A	N/A	N/A	A	N/A
_mm_cmpunord_sd	N/A	N/A	N/A	A	N/A
_mm_comieq_sd	N/A	N/A	N/A	A	N/A
_mm_comilt_sd	N/A	N/A	N/A	A	N/A

_mm_comile_sd	N/A	N/A	N/A	A	N/A
_mm_comigt_sd	N/A	N/A	N/A	A	N/A
_mm_comige_sd	N/A	N/A	N/A	A	N/A
_mm_comineq_sd	N/A	N/A	N/A	A	N/A
_mm_ucomieq_sd	N/A	N/A	N/A	A	N/A
_mm_ucomilt_sd	N/A	N/A	N/A	A	N/A
_mm_ucomile_sd	N/A	N/A	N/A	A	N/A
_mm_ucomigt_sd	N/A	N/A	N/A	A	N/A
_mm_ucomige_sd	N/A	N/A	N/A	A	N/A
_mm_ucomineq_sd	N/A	N/A	N/A	A	N/A
_mm_cvtepi32_pd	N/A	N/A	N/A	A	N/A
_mm_cvtpd_epi32	N/A	N/A	N/A	A	N/A
_mm_cvttpd_epi32	N/A	N/A	N/A	A	N/A
_mm_cvtepi32_ps	N/A	N/A	N/A	A	N/A
_mm_cvtps_epi32	N/A	N/A	N/A	A	N/A
_mm_cvttps_epi32	N/A	N/A	N/A	A	N/A
_mm_cvtpd_ps	N/A	N/A	N/A	A	N/A
_mm_cvtps_pd	N/A	N/A	N/A	A	N/A
_mm_cvtsd_ss	N/A	N/A	N/A	A	N/A
_mm_cvtsd_sd	N/A	N/A	N/A	A	N/A
_mm_cvtsd_si32	N/A	N/A	N/A	A	N/A
_mm_cvttss_si32	N/A	N/A	N/A	A	N/A
_mm_cvtsi32_sd	N/A	N/A	N/A	A	N/A
_mm_cvtpd_pi32	N/A	N/A	N/A	A	N/A
_mm_cvttpd_pi32	N/A	N/A	N/A	A	N/A
_mm_cvtpi32_pd	N/A	N/A	N/A	A	N/A
_mm_unpackhi_pd	N/A	N/A	N/A	A	N/A
_mm_unpacklo_pd	N/A	N/A	N/A	A	N/A
_mm_unpacklo_pd	N/A	N/A	N/A	A	N/A
_mm_shuffle_pd	N/A	N/A	N/A	A	N/A
_mm_load_pd	N/A	N/A	N/A	A	N/A

_mm_loadl_pd	N/A	N/A	N/A	A	N/A
_mm_loadr_pd	N/A	N/A	N/A	A	N/A
_mm_loadu_pd	N/A	N/A	N/A	A	N/A
_mm_load_sd	N/A	N/A	N/A	A	N/A
_mm_loadh_pd	N/A	N/A	N/A	A	N/A
_mm_loadl_pd	N/A	N/A	N/A	A	N/A
_mm_set_sd	N/A	N/A	N/A	A	N/A
_mm_setl_pd	N/A	N/A	N/A	A	N/A
_mm_set_pd	N/A	N/A	N/A	A	N/A
_mm_setr_pd	N/A	N/A	N/A	A	N/A
_mm_setzero_pd	N/A	N/A	N/A	A	N/A
_mm_move_sd	N/A	N/A	N/A	A	N/A
_mm_store_sd	N/A	N/A	N/A	A	N/A
_mm_storel_pd	N/A	N/A	N/A	A	N/A
_mm_store_pd	N/A	N/A	N/A	A	N/A
_mm_storeu_pd	N/A	N/A	N/A	A	N/A
_mm_storer_pd	N/A	N/A	N/A	A	N/A
_mm_storeh_pd	N/A	N/A	N/A	A	N/A
_mm_storel_pd	N/A	N/A	N/A	A	N/A
_mm_add_epi8	N/A	N/A	N/A	A	N/A
_mm_add_epi16	N/A	N/A	N/A	A	N/A
_mm_add_epi32	N/A	N/A	N/A	A	N/A
_mm_add_si64	N/A	N/A	N/A	A	N/A
_mm_add_epi64	N/A	N/A	N/A	A	N/A
_mm_adds_epi8	N/A	N/A	N/A	A	N/A
_mm_adds_epi16	N/A	N/A	N/A	A	N/A
_mm_adds_epu8	N/A	N/A	N/A	A	N/A
_mm_adds_epu16	N/A	N/A	N/A	A	N/A
_mm_avg_epu8	N/A	N/A	N/A	A	N/A
_mm_avg_epu16	N/A	N/A	N/A	A	N/A
_mm_madd_epi16	N/A	N/A	N/A	A	N/A

_mm_max_epi16	N/A	N/A	N/A	A	N/A
_mm_max_epu8	N/A	N/A	N/A	A	N/A
_mm_min_epi16	N/A	N/A	N/A	A	N/A
_mm_min_epu8	N/A	N/A	N/A	A	N/A
_mm_mulhi_epi16	N/A	N/A	N/A	A	N/A
_mm_mulhi_epu16	N/A	N/A	N/A	A	N/A
_mm_mullo_epi16	N/A	N/A	N/A	A	N/A
_mm_mul_su32	N/A	N/A	N/A	A	N/A
_mm_mul_epu32	N/A	N/A	N/A	A	N/A
_mm_sad_epu8	N/A	N/A	N/A	A	N/A
_mm_sub_epi8	N/A	N/A	N/A	A	N/A
_mm_sub_epi16	N/A	N/A	N/A	A	N/A
_mm_sub_epi32	N/A	N/A	N/A	A	N/A
_mm_sub_si64	N/A	N/A	N/A	A	N/A
_mm_sub_epi64	N/A	N/A	N/A	A	N/A
_mm_subs_epi8	N/A	N/A	N/A	A	N/A
_mm_subs_epi16	N/A	N/A	N/A	A	N/A
_mm_subs_epu8	N/A	N/A	N/A	A	N/A
_mm_subs_epu16	N/A	N/A	N/A	A	N/A
_mm_and_si128	N/A	N/A	N/A	A	N/A
_mm_andnot_si128	N/A	N/A	N/A	A	N/A
_mm_or_si128	N/A	N/A	N/A	A	N/A
_mm_xor_si128	N/A	N/A	N/A	A	N/A
_mm_slli_si128	N/A	N/A	N/A	A	N/A
_mm_slli_epi16	N/A	N/A	N/A	A	N/A
_mm_sll_epi16	N/A	N/A	N/A	A	N/A
_mm_slli_epi32	N/A	N/A	N/A	A	N/A
_mm_sll_epi32	N/A	N/A	N/A	A	N/A
_mm_slli_epi64	N/A	N/A	N/A	A	N/A
_mm_sll_epi64	N/A	N/A	N/A	A	N/A
_mm_srai_epi16	N/A	N/A	N/A	A	N/A

_mm_sra_epi16	N/A	N/A	N/A	A	N/A
_mm_srai_epi32	N/A	N/A	N/A	A	N/A
_mm_sra_epi32	N/A	N/A	N/A	A	N/A
_mm_srli_si128	N/A	N/A	N/A	A	N/A
_mm_srli_epi16	N/A	N/A	N/A	A	N/A
_mm_srl_epi16	N/A	N/A	N/A	A	N/A
_mm_srli_epi32	N/A	N/A	N/A	A	N/A
_mm_srl_epi32	N/A	N/A	N/A	A	N/A
_mm_srli_epi64	N/A	N/A	N/A	A	N/A
_mm_srl_epi64	N/A	N/A	N/A	A	N/A
_mm_cmpeq_epi8	N/A	N/A	N/A	A	N/A
_mm_cmpeq_epi16	N/A	N/A	N/A	A	N/A
_mm_cmpeq_epi32	N/A	N/A	N/A	A	N/A
_mm_cmpgt_epi8	N/A	N/A	N/A	A	N/A
_mm_cmpgt_epi16	N/A	N/A	N/A	A	N/A
_mm_cmpgt_epi32	N/A	N/A	N/A	A	N/A
_mm_cmplt_epi8	N/A	N/A	N/A	A	N/A
_mm_cmplt_epi16	N/A	N/A	N/A	A	N/A
_mm_cmplt_epi32	N/A	N/A	N/A	A	N/A
_mm_cvtsi32_si128	N/A	N/A	N/A	A	N/A
_mm_cvtsi128_si32	N/A	N/A	N/A	A	N/A
_mm_packs_epi16	N/A	N/A	N/A	A	N/A
_mm_packs_epi32	N/A	N/A	N/A	A	N/A
_mm_packus_epi16	N/A	N/A	N/A	A	N/A
_mm_extract_epi16	N/A	N/A	N/A	A	N/A
_mm_insert_epi16	N/A	N/A	N/A	A	N/A
_mm_movemask_epi8	N/A	N/A	N/A	A	N/A
_mm_shuffle_epi32	N/A	N/A	N/A	A	N/A
_mm_shufflehi_epi16	N/A	N/A	N/A	A	N/A
_mm_shufflelo_epi16	N/A	N/A	N/A	A	N/A
_mm_unpackhi_epi8	N/A	N/A	N/A	A	N/A

_mm_unpackhi_epi16	N/A	N/A	N/A	A	N/A
_mm_unpackhi_epi32	N/A	N/A	N/A	A	N/A
_mm_unpackhi_epi64	N/A	N/A	N/A	A	N/A
_mm_unpacklo_epi8	N/A	N/A	N/A	A	N/A
_mm_unpacklo_epi16	N/A	N/A	N/A	A	N/A
_mm_unpacklo_epi32	N/A	N/A	N/A	A	N/A
_mm_unpacklo_epi64	N/A	N/A	N/A	A	N/A
_mm_move_epi64	N/A	N/A	N/A	A	N/A
_mm_movpi64_epi64	N/A	N/A	N/A	A	N/A
_mm_movepi64_pi64	N/A	N/A	N/A	A	N/A
_mm_load_si128	N/A	N/A	N/A	A	N/A
_mm_loadu_si128	N/A	N/A	N/A	A	N/A
_mm_loadl_epi64	N/A	N/A	N/A	A	N/A
_mm_set_epi64	N/A	N/A	N/A	A	N/A
_mm_set_epi32	N/A	N/A	N/A	A	N/A
_mm_set_epi16	N/A	N/A	N/A	A	N/A
_mm_set_epi8	N/A	N/A	N/A	A	N/A
_mm_setl_epi64	N/A	N/A	N/A	A	N/A
_mm_setl_epi32	N/A	N/A	N/A	A	N/A
_mm_setl_epi16	N/A	N/A	N/A	A	N/A
_mm_setl_epi8	N/A	N/A	N/A	A	N/A
_mm_setr_epi64	N/A	N/A	N/A	A	N/A
_mm_setr_epi32	N/A	N/A	N/A	A	N/A
_mm_setr_epi16	N/A	N/A	N/A	A	N/A
_mm_setr_epi8	N/A	N/A	N/A	A	N/A
_mm_setzero_si128	N/A	N/A	N/A	A	N/A
_mm_store_si128	N/A	N/A	N/A	A	N/A
_mm_storeu_si128	N/A	N/A	N/A	A	N/A
_mm_storel_epi64	N/A	N/A	N/A	A	N/A
_mm_maskmoveu_si128	N/A	N/A	N/A	A	N/A
_mm_stream_pd	N/A	N/A	N/A	A	N/A

<code>_mm_stream_si128</code>	N/A	N/A	N/A	A	N/A
<code>_mm_clflush</code>	N/A	N/A	N/A	A	N/A
<code>_mm_lfence</code>	N/A	N/A	N/A	A	N/A
<code>_mm_mfence</code>	N/A	N/A	N/A	A	N/A
<code>_mm_stream_si32</code>	N/A	N/A	N/A	A	N/A
<code>_mm_pause</code>	N/A	N/A	N/A	A	N/A

Intel® C++ Class Libraries

The Intel® C++ Class Libraries enable Single-Instruction, Multiple-Data (SIMD) operations. The principle of SIMD operations is to exploit microprocessor architecture through parallel processing. The effect of parallel processing is increased data throughput using fewer clock cycles. The objective is to improve application performance of complex and computation-intensive audio, video, and graphical data bit streams.

Hardware and Software Requirements

You must have the Intel® C++ Compiler version 4.0 or higher installed on your system to use the class libraries. The Intel® C++ Class Libraries are functions abstracted from the instruction extensions available on Intel processors as specified in the table that follows.

Processor Requirements for Use of Class Libraries

Header File	Extension Set	Available on These Processors
<code>ivec.h</code>	MMX(TM) technology	Pentium® with MMX technology, Pentium II, Pentium III, Pentium 4, Intel® Xeon(TM), and Itanium® processors
<code>fvec.h</code>	Streaming SIMD Extensions	Pentium III, Pentium 4, Intel Xeon, and Itanium processors
<code>dvec.h</code>	Streaming SIMD Extensions 2	Pentium 4 and Intel Xeon processors

About the Classes

The Intel® C++ Class Libraries for SIMD Operations include:

- Integer vector (`Ivec`) classes
- Floating-point vector (`Fvec`) classes

You can find the definitions for these operations in three header files: `ivec.h`, `fvec.h`, and `dvec.h`. The classes themselves are not partitioned like this. The classes are named according to the underlying type of operation. The header files are partitioned according to architecture:

- `ivec.h` is specific to architectures with MMX(TM) technology
- `fvec.h` is specific to architectures with Streaming SIMD Extensions
- `dvec.h` is specific to architectures with Streaming SIMD Extensions 2

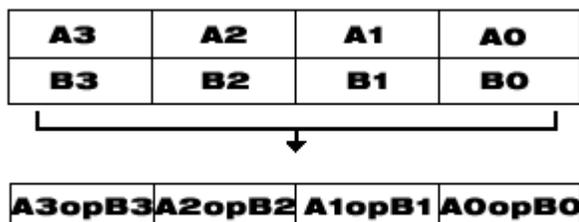
Streaming SIMD Extensions 2 intrinsics cannot be used on Itanium®-based systems. The `mmclass.h` header file includes the classes that are usable on the Itanium architecture.

This documentation is intended for programmers writing code for the Intel architecture, particularly code that would benefit from the use of SIMD instructions. You should be familiar with C++ and the use of C++ classes.

Details About the Libraries

The Intel® C++ Class Libraries for SIMD Operations provide a convenient interface to access the underlying instructions for processors as specified in Processor Requirements for Use of Class Libraries. These processor-instruction extensions enable parallel processing using the single instruction-multiple data (SIMD) technique as illustrated in the following figure.

SIMD Data Flow



Performing four operations with a single instruction improves efficiency by a factor of four for that particular instruction.

These new processor instructions can be implemented using assembly inlining, intrinsics, or the C++ SIMD classes. Compare the coding required to add four 32-bit floating-point values, using each of the available interfaces:

Comparison Between Inlining, Intrinsics and Class Libraries

Assembly Inlining	Intrinsics	SIMD Class Libraries
<pre>... __m128 a,b,c; __asm{ movaps xmm0,b movaps xmm1,c addps xmm0,xmm1 movaps a, xmm0 } ...</pre>	<pre>#include <mmintrin.h> ... __m128 a,b,c; a = _mm_add_ps(b,c); ...</pre>	<pre>#include <fvec.h> ... F32vec4 a,b,c; a = b +c; ...</pre>

This table shows an addition of two single-precision floating-point values using assembly inlining, intrinsics, and the libraries. You can see how much easier it is to code with the Intel C++ SIMD Class Libraries. Besides using fewer keystrokes and fewer lines of code, the notation is like the standard notation in C++, making it much easier to implement over other methods.

C++ Classes and SIMD Operations

The use of C++ classes for SIMD operations is based on the concept of operating on arrays, or vectors of data, in parallel. Consider the addition of two vectors, A and B, where each vector contains four elements. Using the integer vector (Ivec) class, the elements A[i] and B[i] from each array are summed as shown in the following example.

Typical Method of Adding Elements Using a Loop

```
short a[4], b[4], c[4];
for (i=0; i<4; i++) /* needs four iterations */
c[i] = a[i] + b[i]; /* returns c[0], c[1], c[2], c[3] *
```

The following example shows the same results using one operation with `Ivec` Classes.

SIMD Method of Adding Elements Using Ivec Classes

```
sIs16vec4 ivecA, ivecB, ivec C; /*needs one iteration */
ivecC = ivecA + ivecB; /*returns ivecC0, ivecC1, ivecC2, ivecC3 */
```

Available Classes

The Intel C++ SIMD classes provide parallelism, which is not easily implemented using typical mechanisms of C++. The following table shows how the Intel C++ SIMD classes use the classes and libraries.

SIMD Vector Classes

Instruction Set	Class	Signedness	Data Type	Size	Elements	Header File
MMX(TM) technology (available for IA-32- and Itanium®-based systems)	<code>I64vec1</code>	unspecified	<code>__m64</code>	64	1	<code>ivec.h</code>
	<code>I32vec2</code>	unspecified	<code>int</code>	32	2	<code>ivec.h</code>
	<code>Is32vec2</code>	signed	<code>int</code>	32	2	<code>ivec.h</code>
	<code>Iu32vec2</code>	unsigned	<code>int</code>	32	2	<code>ivec.h</code>
	<code>I16vec4</code>	unspecified	<code>short</code>	16	4	<code>ivec.h</code>
	<code>Is16vec4</code>	signed	<code>short</code>	16	4	<code>ivec.h</code>
	<code>Iu16vec4</code>	unsigned	<code>short</code>	16	4	<code>ivec.h</code>
	<code>I8vec8</code>	unspecified	<code>char</code>	8	8	<code>ivec.h</code>
	<code>Is8vec8</code>	signed	<code>char</code>	8	8	<code>ivec.h</code>
	<code>Iu8vec8</code>	unsigned	<code>char</code>	8	8	<code>ivec.h</code>
Streaming SIMD Extensions (available for IA-32 and Itanium-based systems)	<code>F32vec4</code>	signed	<code>float</code>	32	4	<code>fvec.h</code>
	<code>F32vec1</code>	signed	<code>float</code>	32	1	<code>fvec.h</code>
Streaming SIMD Extensions 2 (available for IA-32-based systems only)	<code>F64vec2</code>	signed	<code>double</code>	64	2	<code>dvec.h</code>

Instruction Set	Class	Signedness	Data Type	Size	Elements	Header File
	I128vec1	unspecified	__m128i	128	1	dvec.h
	I64vec2	unspecified	long int	64	4	dvec.h
	Is64vec2	signed	long int	64	4	dvec.h
	Iu64vec2	unsigned	long int	32	4	dvec.h
	I32vec4	unspecified	int	32	4	dvec.h
	Is32vec4	signed	int	32	4	dvec.h
	Iu32vec4	unsigned	int	32	4	dvec.h
	I16vec8	unspecified	int	16	8	dvec.h
	Is16vec8	signed	int	16	8	dvec.h
	Iu16vec8	unsigned	int	16	8	dvec.h
	I8vec16	unspecified	char	8	16	dvec.h
	Is8vec16	signed	char	8	16	dvec.h
	Iu8vec16	unsigned	char	8	16	dvec.h

Most classes contain similar functionality for all data types and are represented by all available intrinsics. However, some capabilities do not translate from one data type to another without suffering from poor performance, and are therefore excluded from individual classes.



Note

Intrinsics that take immediate values and cannot be expressed easily in classes are not implemented. (For example, `_mm_shuffle_ps`, `_mm_shuffle_pi16`, `_mm_extract_pi16`, `_mm_insert_pi16`).

Access to Classes Using Header Files

The required class header files are installed in the include directory with the Intel® C++ Compiler. To enable the classes, use the `#include` directive in your program file as shown in the table that follows.

Include Directives for Enabling Classes

Instruction Set Extension	Include Directive
MMX Technology	<code>#include <ivec.h></code>
Streaming SIMD Extensions	<code>#include <fvec.h></code>
Streaming SIMD Extensions 2	<code>#include <dvec.h></code>

Each succeeding file from the top down includes the preceding class. You only need to include `fvec.h` if you want to use both the `Ivec` and `Fvec` classes. Similarly, to use all the classes including those for the Streaming SIMD Extensions 2, you need only to include the `dvec.h` file.

Usage Precautions

When using the C++ classes, you should follow some general guidelines. More detailed usage rules for each class are listed in Integer Vector Classes, and Floating-point Vector Classes.

Clear MMX Registers

If you use both the `Ivec` and `Fvec` classes at the same time, your program could mix MMX instructions, called by `Ivec` classes, with Intel x87 architecture floating-point instructions, called by `Fvec` classes. Floating-point instructions exist in the following `Fvec` functions:

- `fvec` constructors
- debug functions (`cout` and element access)
- `rsqrt_nr`



Note

MMX registers are aliased on the floating-point registers, so you should clear the MMX state with the EMMS instruction intrinsic before issuing an x87 floating-point instruction, as in the following example.

<code>ivecA = ivecA & ivecB;</code>	Ivec logical operation that uses MMX instructions
<code>empty ();</code>	clear state
<code>cout << f32vec4a;</code>	F32vec4 operation that uses x87 floating-point instructions



Caution

Failure to clear the MMX registers can result in incorrect execution or poor performance due to an incorrect register state.

Follow EMMS Instruction Guidelines

Intel strongly recommends that you follow the guidelines for using the EMMS instruction. Refer to this topic before coding with the `Ivec` classes.

Capabilities

The fundamental capabilities of each C++ SIMD class include:

- computation
- horizontal data motion
- branch compression/elimination
- caching hints

Understanding each of these capabilities and how they interact is crucial to achieving desired results.

Computation

The SIMD C++ classes contain vertical operator support for most arithmetic operations, including shifting and saturation.

Computation operations include: +, -, *, /, reciprocal (`rcp` and `rcp_nr`), square root (`sqrt`), reciprocal square root (`rsqrt` and `rsqrt_nr`).

Operations `rcp` and `rsqrt` are new approximating instructions with very short latencies that produce results with at least 12 bits of accuracy. Operations `rcp_nr` and `rsqrt_nr` use software refining techniques to enhance the accuracy of the approximations, with a minimal impact on performance. (The "nr" stands for Newton-Raphson, a mathematical technique for improving performance using an approximate result.)

Horizontal Data Support

The C++ SIMD classes provide horizontal support for some arithmetic operations. The term "horizontal" indicates computation across the elements of one vector, as opposed to the vertical, element-by-element operations on two different vectors.

The `add_horizontal`, `unpack_low` and `pack_sat` functions are examples of horizontal data support. This support enables certain algorithms that cannot exploit the full potential of SIMD instructions.

Shuffle intrinsics are another example of horizontal data flow. Shuffle intrinsics are not expressed in the C++ classes due to their immediate arguments. However, the C++ class implementation enables you to mix shuffle intrinsics with the other C++ functions. For example:

```
F32vec4 fveca, fvecb, fvecd;
fveca += fvecb;
fvecd = _mm_shuffle_ps(fveca, fvecb, 0);
```

Typically every instruction with horizontal data flow contains some inefficiency in the implementation. If possible, implement your algorithms without using the horizontal capabilities.

Branch Compression/Elimination

Branching in SIMD architectures can be complicated and expensive, possibly resulting in poor predictability and code expansion. The SIMD C++ classes provide functions to eliminate branches, using logical operations, max and min functions, conditional selects, and compares. Consider the following example:

```
short a[4], b[4], c[4];
for (i=0; i<4; i++)
c[i] = a[i] > b[i] ? a[i] : b[i];
```

This operation is independent of the value of `i`. For each `i`, the result could be either A or B depending on the actual values. A simple way of removing the branch altogether is to use the `select_gt` function, as follows:

```
Is16vec4 a, b, c
c = select_gt(a, b, a, b)
```

Caching Hints

Streaming SIMD Extensions provide prefetching and streaming hints. Prefetching data can minimize the effects of memory latency. Streaming hints allow you to indicate that certain data should not be cached. This results in higher performance for data that should be cached.

Integer Vector Classes

The `Ivec` classes provide an interface to SIMD processing using integer vectors of various sizes. The class hierarchy is represented in the following figure.

Ivec Class Hierarchy



OM00834

The `M64` and `M128` classes define the `__m64` and `__m128i` data types from which the rest of the `Ivec` classes are derived. The first generation of child classes are derived based solely on bit sizes of 128, 64, 32, 16, and 8 respectively for the `I128vec1`, `I64vec1`, `I64vec2`, `I32vec2`, `I32vec4`, `I16vec4`, `I16vec8`, `I8vec16`, and `I8vec8` classes. The latter seven of these classes require specification of signedness and saturation.



Caution

Do not intermix the `M64` and `M128` data types. You will get unexpected behavior if you do.

The signedness is indicated by the `s` and `u` in the class names:

```

Is64vec2
Iu64vec2
Is32vec4
Iu32vec4
Is16vec8
Iu16vec8
Is8vec16
Iu8vec16
Is32vec2
Iu32vec2
Is16vec4
Iu16vec4
Is8vec8
Iu8vec8
  
```

Terms, Conventions, and Syntax

The following are special terms and syntax used in this chapter to describe functionality of the classes with respect to their associated operations.

Ivec Class Syntax Conventions

The name of each class denotes the data type, signedness, bit size, number of elements using the following generic format:

```
<type><signedness><bits>vec<elements>
```

```
{ F | I } { s | u } { 64 | 32 | 16 | 8 } vec { 8 | 4 | 2 | 1 }
```

where

type	indicates floating point (F) or integer (I)
signedness	indicates signed (s) or unsigned (u). For the Ivec class, leaving this field blank indicates an intermediate class. There are no unsigned Fvec classes, therefore for the Fvec classes, this field is blank.
bits	specifies the number of bits per element
elements	specifies the number of elements

Special Terms and Conventions

The following terms are used to define the functionality and characteristics of the classes and operations defined in this manual.

- **Nearest Common Ancestor** -- This is the intermediate or parent class of two classes of the same size. For example, the nearest common ancestor of `Iu8vec8` and `Is8vec8` is `I8vec8`. Also, the nearest common ancestor between `Iu8vec8` and `I16vec4` is `M64`.
- **Casting** -- Changes the data type from one class to another. When an operation uses different data types as operands, the return value of the operation must be assigned to a single data type. Therefore, one or more of the data types must be converted to a required data type. This conversion is known as a typecast. Sometimes, typecasting is automatic, other times you must use special syntax to explicitly typecast it yourself.
- **Operator Overloading** -- This is the ability to use various operators on the same user-defined data type of a given class. Once you declare a variable, you can add, subtract, multiply, and perform a range of operations. Each family of classes accepts a specified range of operators, and must comply by rules and restrictions regarding typecasting and operator overloading as defined in the header files. The following table shows the notation used in this documentation to address typecasting, operator overloading, and other rules.

Class Syntax Notation Conventions

Class Name	Description
<code>I[s u][N]vec[N]</code>	Any value except <code>I128vec1</code> nor <code>I64vec1</code>
<code>I64vec1</code>	<code>__m64</code> data type
<code>I[s u]64vec2</code>	two 64-bit values of any signedness
<code>I[s u]32vec4</code>	four 32-bit values of any signedness
<code>I[s u]8vec16</code>	eight 16-bit values of any signedness
<code>I[s u]16vec8</code>	sixteen 8-bit values of any signedness
<code>I[s u]32vec2</code>	two 32-bit values of any signedness
<code>I[s u]16vec4</code>	four 16-bit values of any signedness
<code>I[s u]8vec8</code>	eight 8-bit values of any signedness

Rules for Operators

To use operators with the `Ivec` classes you must use one of the following three syntax conventions:

```
[ Ivec_Class ] R = [ Ivec_Class ] A [ operator ][ Ivec_Class ] B
```

Example 1: `I64vec1 R = I64vec1 A & I64vec1 B;`

```
[ Ivec_Class ] R =[ operator ] ([ Ivec_Class ] A,[ Ivec_Class ] B)
```

Example 2: `I64vec1 R = andnot(I64vec1 A, I64vec1 B);`

```
[ Ivec_Class ] R [ operator ]= [ Ivec_Class ] A
```

Example 3: `I64vec1 R &= I64vec1 A;`

[operator] an operator (for example, `&`, `|`, or `^`)

[Ivec_Class] an `Ivec` class

R, A, B variables declared using the pertinent `Ivec` classes

The table that follows shows automatic and explicit sign and size typecasting. "Explicit" means that it is illegal to mix different types without an explicit typecasting. "Automatic" means that you can mix types freely and the compiler will do the typecasting for you.

Summary of Rules Major Operators

Operators	Sign Typecasting	Size Typecasting	Other Typecasting Requirements
Assignment	N/A	N/A	N/A
Logical	Automatic	Automatic (to left)	Explicit typecasting is required for different types used in non-logical expressions on the right side of the assignment.
Addition and Subtraction	Automatic	Explicit	N/A
Multiplication	Automatic	Explicit	N/A
Shift	Automatic	Explicit	Casting Required to ensure arithmetic shift.
Compare	Automatic	Explicit	Explicit casting is required for signed classes for the less-than or greater-than operations.
Conditional Select	Automatic	Explicit	Explicit casting is required for signed classes for less-than or greater-than operations.

Data Declaration and Initialization

The following table shows literal examples of constructor declarations and data type initialization for all class sizes. All values are initialized with the most significant element on the left and the least significant to the right.

Declaration and Initialization Data Types for Ivec Classes

Operation	Class	Syntax
Declaration	M128	<code>I128vec1 A; Iu8vec16 A;</code>
Declaration	M64	<code>I64vec1 A; Iu8vec16 A;</code>
<code>__m128</code> Initialization	M128	<code>I128vec1 A(__m128 m); Iu16vec8(__m128 m);</code>
<code>__m64</code> Initialization	M64	<code>I64vec1 A(__m64 m); Iu8vec8 A(__m64 m);</code>
<code>__int64</code> Initialization	M64	<code>I64vec1 A = __int64 m; Iu8vec8 A = __int64 m;</code>
<code>int i</code> Initialization	M64	<code>I64vec1 A = int i; Iu8vec8 A = int i;</code>

Operation	Class	Syntax
int initialization	I32vec2	I32vec2 A(int A1, int A0); Is32vec2 A(signed int A1, signed int A0); Iu32vec2 A(unsigned int A1, unsigned int A0);
int Initialization	I32vec4	I32vec4 A(short A3, short A2, short A1, short A0); Is32vec4 A(signed short A3, ..., signed short A0); Iu32vec4 A(unsigned short A3, ..., unsigned short A0);
short int Initialization	I16vec4	I16vec4 A(short A3, short A2, short A1, short A0); Is16vec4 A(signed short A3, ..., signed short A0); Iu16vec4 A(unsigned short A3, ..., unsigned short A0);
short int Initialization	I16vec8	I16vec8 A(short A7, short A6, ..., short A1, short A0); Is16vec8 A(signed A7, ..., signed short A0); Iu16vec8 A(unsigned short A7, ..., unsigned short A0);
char Initialization	I8vec8	I8vec8 A(char A7, char A6, ..., char A1, char A0); Is8vec8 A(signed char A7, ..., signed char A0); Iu8vec8 A(unsigned char A7, ..., unsigned char A0);
char Initialization	I8vec16	I8vec16 A(char A15, ..., char A0); Is8vec16 A(signed char A15, ..., signed char A0); Iu8vec16 A(unsigned char A15, ..., unsigned char A0);

Assignment Operator

Any `Ivec` object can be assigned to any other `Ivec` object; conversion on assignment from one `Ivec` object to another is automatic.

Assignment Operator Examples

```
Is16vec4 A;

Is8vec8 B;

I64vec1 C;

A = B; /* assign Is8vec8 to Is16vec4 */

B = C; /* assign I64vec1 to Is8vec8 */

B = A & C; /* assign M64 result of '&' to Is8vec8 */
```

Logical Operators

The logical operators use the symbols and intrinsics listed in the following table.

Bitwise Operation	Operator Symbols		Syntax Usage		Corresponding Intrinsic
	Standard	w/assign	Standard	w/assign	
AND	&	&=	R = A & B	R &= A	<code>_mm_and_si64</code> <code>_mm_and_si128</code>
OR		=	R = A B	R = A	<code>_mm_and_si64</code> <code>_mm_and_si128</code>
XOR	^	^=	R = A ^ B	R ^= A	<code>_mm_and_si64</code> <code>_mm_and_si128</code>
ANDNOT	<code>andnot</code>	N/A	R = A <code>andnot</code> B	N/A	<code>_mm_and_si64</code> <code>_mm_and_si128</code>

Logical Operators and Miscellaneous Exceptions.

A and B converted to M64. Result assigned to `Iu8vec8`.

```
I64vec1 A;

Is8vec8 B;

Iu8vec8 C;

C = A & B;
```

Same size and signedness operators return the nearest common ancestor.

```
I32vec2 R = Is32vec2 A ^ Iu32vec2 B;
```

A&B returns M64, which is cast to Iu8vec8.

```
C = Iu8vec8(A&B)+ C;
```

When A and B are of the same class, they return the same type. When A and B are of different classes, the return value is the return type of the nearest common ancestor.

The logical operator returns values for combinations of classes, listed in the following tables, apply when A and B are of different classes.

Ivec Logical Operator Overloading

Return ®	AND	OR	XOR	NAND	A Operand	B Operand
I64vec1 R	&		^	andnot	I[s u]64vec2 A	I[s u]64vec2 B
I64vec2 R	&		^	andnot	I[s u]64vec2 A	I[s u]64vec2 B
I32vec2 R	&		^	andnot	I[s u]32vec2 A	I[s u]32vec2 B
I32vec4 R	&		^	andnot	I[s u]32vec4 A	I[s u]32vec4 B
I16vec4 R	&		^	andnot	I[s u]16vec4 A	I[s u]16vec4 B
I16vec8 R	&		^	andnot	I[s u]16vec8 A	I[s u]16vec8 B
I8vec8 R	&		^	andnot	I[s u]8vec8 A	I[s u]8vec8 B
I8vec16 R	&		^	andnot	I[s u]8vec16 A	I[s u]8vec16 B

For logical operators with assignment, the return value of R is always the same data type as the pre-declared value of R as listed in the table that follows.

Ivec Logical Operator Overloading with Assignment

Return Type	Left Side ®	AND	OR	XOR	Right Side (Any Ivec Type)
I128vec1	I128vec1 R	&=	=	^=	I[s u][N]vec[N] A;
I64vec1	I64vec1 R	&=	=	^=	I[s u][N]vec[N] A;
I64vec2	I64vec2 R	&=	=	^=	I[s u][N]vec[N] A;
I[x]32vec4	I[x]32vec4 R	&=	=	^=	I[s u][N]vec[N] A;
I[x]32vec2	I[x]32vec2 R	&=	=	^=	I[s u][N]vec[N] A;
I[x]16vec8	I[x]16vec8 R	&=	=	^=	I[s u][N]vec[N] A;
I[x]16vec4	I[x]16vec4 R	&=	=	^=	I[s u][N]vec[N] A;
I[x]8vec16	I[x]8vec16 R	&=	=	^=	I[s u][N]vec[N] A;
I[x]8vec8	I[x]8vec8 R	&=	=	^=	I[s u][N]vec[N] A;

Addition and Subtraction Operators

The addition and subtraction operators return the class of the nearest common ancestor when the right-side operands are of different signs. The following code provides examples of usage and miscellaneous exceptions.

Syntax Usage for Addition and Subtraction Operators

Return nearest common ancestor type, I16vec4.

```
Is16vec4 A;
```

```
Iu16vec4 B;
```

```
I16vec4 C;
```

```
C = A + B;
```

Returns type left-hand operand type.

```
Is16vec4 A;
```

```
Iu16vec4 B;
```

```
A += B;
```

```
B -= A;
```

Explicitly convert B to Is16vec4.

```
Is16vec4 A,C;
```

```

Iu32vec24 B;

C = A + C;

C = A + (Is16vec4)B;

```

Addition and Subtraction Operators with Corresponding Intrinsics

Operation	Symbols	Syntax	Corresponding Intrinsics
Addition	+ +=	R = A + B R += A	_mm_add_epi64 _mm_add_epi32 _mm_add_epi16 _mm_add_epi8 _mm_add_pi32 _mm_add_pi16 _mm_add_pi8
Subtraction	- -=	R = A - B R -= A	_mm_sub_epi64 _mm_sub_epi32 _mm_sub_epi16 _mm_sub_epi8 _mm_sub_pi32 _mm_sub_pi16 _mm_sub_pi8

The following table lists addition and subtraction return values for combinations of classes when the right side operands are of different signedness. The two operands must be the same size, otherwise you must explicitly indicate the typecasting.

Addition and Subtraction Operator Overloading

Return Value	Available Operators		Right Side Operands	
	Add	Sub	A	B
I64vec2 R	+	-	I[s u]64vec2 A	I[s u]64vec2 B
I32vec4 R	+	-	I[s u]32vec4 A	I[s u]32vec4 B
I32vec2 R	+	-	I[s u]32vec2 A	I[s u]32vec2 B
I16vec8 R	+	-	I[s u]16vec8 A	I[s u]16vec8 B
I16vec4 R	+	-	I[s u]16vec4 A	I[s u]16vec4 B
I8vec8 R	+	-	I[s u]8vec8 A	I[s u]8vec8 B
I8vec16 R	+	-	I[s u]8vec2 A	I[s u]8vec16 B

The following table shows the return data type values for operands of the addition and subtraction operators with assignment. The left side operand determines the size and signedness of the return value. The right side operand must be the same size as the left operand; otherwise, you must use an explicit typecast.

Addition and Subtraction with Assignment

Return Value ®	Left Side ®	Add	Sub	Right Side (A)
I[x]32vec4	I[x]32vec2 R	+=	-=	I[s u]32vec4 A;
I[x]32vec2 R	I[x]32vec2 R	+=	-=	I[s u]32vec2 A;
I[x]16vec8	I[x]16vec8	+=	-=	I[s u]16vec8 A;
I[x]16vec4	I[x]16vec4	+=	-=	I[s u]16vec4 A;
I[x]8vec16	I[x]8vec16	+=	-=	I[s u]8vec16 A;
I[x]8vec8	I[x]8vec8	+=	-=	I[s u]8vec8 A;

Multiplication Operators

The multiplication operators can only accept and return data types from the I[s|u]16vec4 or I[s|u]16vec8 classes, as shown in the following example.

Syntax Usage for Multiplication Operators

Explicitly convert B to Is16vec4.

```
Is16vec4 A,C;
```

```
Iu32vec2 B;
```

```
C = A * C;
```

```
C = A * (Is16vec4)B;
```

Return nearest common ancestor type, I16vec4

```
Is16vec4 A;
```

```
Iu16vec4 B;
```

```
I16vec4 C;
```

```
C = A + B;
```

The mul_high and mul_add functions take Is16vec4 data only.

```
Is16vec4 A,B,C,D;
```

```
C = mul_high(A,B);
```

```
D = mul_add(A,B);
```

Multiplication Operators with Corresponding Intrinsics

Symbols		Syntax Usage	Intrinsic
*	*=	R = A * B R *= A	_mm_mullo_pi16 _mm_mullo_epi16
mul_high	N/A	R = mul_high(A, B)	_mm_mulhi_pi16 _mm_mulhi_epi16
mul_add	N/A	R = mul_high(A, B)	_mm_madd_pi16 _mm_madd_epi16

The multiplication return operators always return the nearest common ancestor as listed in the table that follows. The two operands must be 16 bits in size, otherwise you must explicitly indicate typecasting.

Multiplication Operator Overloading

R	Mul	A	B
I16vec4 R	*	I[s u]16vec4 A	I[s u]16vec4 B
I16vec8 R	*	I[s u]16vec8 A	I[s u]16vec8 B
Is16vec4 R	mul_add	Is16vec4 A	Is16vec4 B
Is16vec8	mul_add	Is16vec8 A	Is16vec8 B
Is32vec2 R	mul_high	Is16vec4 A	Is16vec4 B
Is32vec4 R	mul_high	s16vec8 A	Is16vec8 B

The following table shows the return values and data type assignments for operands of the multiplication operators with assignment. All operands must be 16 bytes in size. If the operands are not the right size, you must use an explicit typecast.

Multiplication with Assignment

Return Value ®	Left Side ®	Mul	Right Side (A)
I[x]16vec8	I[x]16vec8	*=	I[s u]16vec8 A;
I[x]16vec4	I[x]16vec4	*=	I[s u]16vec4 A;

Shift Operators

The right shift argument can be any integer or `Ivec` value, and is implicitly converted to a M64 data type. The first or left operand of a `<<` can be of any type except `I[s|u]8vec[8|16]`.

Example Syntax Usage for Shift Operators

Automatic size and sign conversion.

```
Is16vec4 A, C;
```

```
Iu32vec2 B;
```

```
C = A;
```

`A&B` returns `I16vec4`, which must be cast to `Iu16vec4` to ensure logical shift, not arithmetic shift.

```
Is16vec4 A, C;
```

```
Iu16vec4 B, R;
```

```
R = (Iu16vec4)(A & B) C;
```

`A&B` returns `I16vec4`, which must be cast to `Is16vec4` to ensure arithmetic shift, not logical shift.

```
R = (Is16vec4)(A & B) C;
```

Shift Operators with Corresponding Intrinsics

Operation	Symbols	Syntax Usage	Intrinsic
Shift Left	<code><<</code> <code>&=</code>	<code>R = A << B</code> <code>R &= A</code>	<code>_mm_sll_si64</code> <code>_mm_slli_si64</code> <code>_mm_sll_pi32</code> <code>_mm_slli_pi32</code> <code>_mm_sll_pi16</code> <code>_mm_slli_pi16</code>
Shift Right	<code>>></code>	<code>R = A >> B</code> <code>R >>= A</code>	<code>_mm_srl_si64</code> <code>_mm_srli_si64</code> <code>_mm_srl_pi32</code> <code>_mm_srli_pi32</code> <code>_mm_srl_pi16</code> <code>_mm_srli_pi16</code> <code>_mm_sra_pi32</code> <code>_mm_srai_pi32</code> <code>_mm_sra_pi16</code> <code>_mm_srai_pi16</code>

Right shift operations with signed data types use arithmetic shifts. All unsigned and intermediate classes correspond to logical shifts. The following table shows how the return type is determined by the first argument type.

Shift Operator Overloading

Operation	R	Right Shift		Left Shift		A	B
Logical	I64vec1	>>	>>=	<<	<<=	I64vec1 A;	I64vec1 B;
Logical	I32vec2	>>	>>=	<<	<<=	I32vec2 A	I32vec2 B;
Arithmetic	Is32vec2	>>	>>=	<<	<<=	Is32vec2 A	I[s u][N]vec[N] B;
Logical	Iu32vec2	>>	>>=	<<	<<=	Iu32vec2 A	I[s u][N]vec[N] B;
Logical	I16vec4	>>	>>=	<<	<<=	I16vec4 A	I16vec4 B
Arithmetic	Is16vec4	>>	>>=	<<	<<=	Is16vec4 A	I[s u][N]vec[N] B;
Logical	Iu16vec4	>>	>>=	<<	<<=	Iu16vec4 A	I[s u][N]vec[N] B;

Comparison Operators

The equality and inequality comparison operands can have mixed signedness, but they must be of the same size. The comparison operators for less-than and greater-than must be of the same sign and size.

Example of Syntax Usage for Comparison Operator

The nearest common ancestor is returned for compare for equal/not-equal operations.

```
Iu8vec8 A;
```

```
Is8vec8 B;
```

```
I8vec8 C;
```

```
C = cmpneq(A,B);
```

Type cast needed for different-sized elements for equal/not-equal comparisons.

```
Iu8vec8 A, C;
```

```
Is16vec4 B;
```

```
C = cmpeq(A, (Iu8vec8)B);
```

Type cast needed for sign or size differences for less-than and greater-than comparisons.

```
Iu16vec4 A;
```

```
Is16vec4 B, C;
```

```
C = cmpge((Is16vec4)A,B);
```

```
C = cmpgt(B,C);
```

Inequality Comparison Symbols and Corresponding Intrinsics

Compare For:	Operators	Syntax	Intrinsic	
Equality	cmpeq	R = cmpeq(A, B)	_mm_cmpeq_pi32 _mm_cmpeq_pi16 _mm_cmpeq_pi8	
Inequality	cmpneq	R = cmpneq(A, B)	_mm_cmpeq_pi32 _mm_cmpeq_pi16 _mm_cmpeq_pi8	_mm_andnot_si64
Greater Than	cmpgt	R = cmpgt(A, B)	_mm_cmpgt_pi32 _mm_cmpgt_pi16 _mm_cmpgt_pi8	
Greater Than or Equal To	cmpge	R = cmpge(A, B)	_mm_cmpgt_pi32 _mm_cmpgt_pi16 _mm_cmpgt_pi8	_mm_andnot_si64
Less Than	cmplt	R = cmplt(A, B)	_mm_cmpgt_pi32 _mm_cmpgt_pi16 _mm_cmpgt_pi8	
Less Than or Equal To	cmple	R = cmple(A, B)	_mm_cmpgt_pi32 _mm_cmpgt_pi16 _mm_cmpgt_pi8	_mm_andnot_si64

Comparison operators have the restriction that the operands must be the size and sign as listed in the Compare Operator Overloading table.

Compare Operator Overloading

R	Comparison	A	B
I32vec2 R	cmpeq cmpne	I[s u]32vec2 B	I[s u]32vec2 B
I16vec4 R		I[s u]16vec4 B	I[s u]16vec4 B
I8vec8 R		I[s u]8vec8 B	I[s u]8vec8 B
I32vec2 R	cmpgt cmpge cmplt cmple	Is32vec2 B	Is32vec2 B
I16vec4 R		Is16vec4 B	Is16vec4 B
I8vec8 R		Is8vec8 B	Is8vec8 B

Conditional Select Operators

For conditional select operands, the third and fourth operands determine the type returned. Third and fourth operands with same size, but different signedness, return the nearest common ancestor data type.

Conditional Select Syntax Usage

Return the nearest common ancestor data type if third and fourth operands are of the same size, but different signs.

```
I16vec4 R = select_neq(Is16vec4, Is16vec4, Is16vec4, Iu16vec4);
```

Conditional Select for Equality

```
R0 := (A0 == B0) ? C0 : D0;
```

```
R1 := (A1 == B1) ? C1 : D1;
```

```
R2 := (A2 == B2) ? C2 : D2;
```

```
R3 := (A3 == B3) ? C3 : D3;
```

Conditional Select for Inequality

```
R0 := (A0 != B0) ? C0 : D0;
```

```
R1 := (A1 != B1) ? C1 : D1;
```

```
R2 := (A2 != B2) ? C2 : D2;
```

```
R3 := (A3 != B3) ? C3 : D3;
```

Conditional Select Symbols and Corresponding Intrinsics

Conditional Select For:	Operators	Syntax	Corresponding Intrinsic	Additional Intrinsic (Applies to All)
Equality	<code>select_eq</code>	<code>R = select_eq(A, B, C, D)</code>	<code>_mm_cmpeq_pi32</code> <code>_mm_cmpeq_pi16</code> <code>_mm_cmpeq_pi8</code>	<code>_mm_and_si64</code> <code>_mm_or_si64</code> <code>_mm_andnot_si64</code>
Inequality	<code>select_neq</code>	<code>R = select_neq(A, B, C, D)</code>	<code>_mm_cmpeq_pi32</code> <code>_mm_cmpeq_pi16</code> <code>_mm_cmpeq_pi8</code>	
Greater Than	<code>select_gt</code>	<code>R = select_gt(A, B, C, D)</code>	<code>_mm_cmpgt_pi32</code> <code>_mm_cmpgt_pi16</code> <code>_mm_cmpgt_pi8</code>	

Conditional Select For:	Operators	Syntax	Corresponding Intrinsic	Additional Intrinsic (Applies to All)
Greater Than or Equal To	<code>select_ge</code>	R = <code>select_gt(A, B, C, D)</code>	<code>_mm_cmpge_pi32</code> <code>_mm_cmpge_pi16</code> <code>_mm_cmpge_pi8</code>	
Less Than	<code>select_lt</code>	R = <code>select_lt(A, B, C, D)</code>	<code>_mm_cmplt_pi32</code> <code>_mm_cmplt_pi16</code> <code>_mm_cmplt_pi8</code>	
Less Than or Equal To	<code>select_le</code>	R = <code>select_le(A, B, C, D)</code>	<code>_mm_cmple_pi32</code> <code>_mm_cmple_pi16</code> <code>_mm_cmple_pi8</code>	

All conditional select operands must be of the same size. The return data type is the nearest common ancestor of operands C and D. For conditional select operations using greater-than or less-than operations, the first and second operands must be signed as listed in the table that follows.

Conditional Select Operator Overloading

R	Comparison	A and B	C	D
I32vec2 R	<code>select_eq</code> <code>select_ne</code>	I[s u]32vec2	I[s u]32vec2	I[s u]32vec2
I16vec4 R		I[s u]16vec4	I[s u]16vec4	I[s u]16vec4
I8vec8 R		I[s u]8vec8	I[s u]8vec8	I[s u]8vec8
I32vec2 R	<code>select_gt</code> <code>select_ge</code> <code>select_lt</code> <code>select_le</code>	Is32vec2	Is32vec2	Is32vec2
I16vec4 R		Is16vec4	Is16vec4	Is16vec4
I8vec8 R		Is8vec8	Is8vec8	Is8vec8

The following table shows the mapping of return values from R0 to R7 for any number of elements. The same return value mappings also apply when there are fewer than four return values.

Conditional Select Operator Return Value Mapping

Return Value	A and B Operands						C and D operands		
	A0	Available Operators				B0			
R0:=	A0	==	!=	>	>=	<	<=	B0	? C0 : D0;
R1:=	A0	==	!=	>	>=	<	<=	B0	? C1 : D1;
R2:=	A0	==	!=	>	>=	<	<=	B0	? C2 : D2;
R3:=	A0	==	!=	>	>=	<	<=	B0	? C3 : D3;
R4:=	A0	==	!=	>	>=	<	<=	B0	? C4 : D4;
R5:=	A0	==	!=	>	>=	<	<=	B0	? C5 : D5;
R6:=	A0	==	!=	>	>=	<	<=	B0	? C6 : D6;
R7:=	A0	==	!=	>	>=	<	<=	B0	? C7 : D7;

Debug

The debug operations do not map to any compiler intrinsics for MMX(TM) instructions. They are provided for debugging programs only. Use of these operations may result in loss of performance, so you should not use them outside of debugging.

Output

The four 32-bit values of A are placed in the output buffer and printed in the following format (default in decimal):

```
cout << Is32vec4 A;

cout << Iu32vec4 A;

cout << hex << Iu32vec4 A; /* print in hex format */

"[3]:A3 [2]:A2 [1]:A1 [0]:A0"
```

Corresponding Intrinsics: none

The two 32-bit values of A are placed in the output buffer and printed in the following format (default in decimal):

```
cout << Is32vec2 A;

cout << Iu32vec2 A;

cout << hex << Iu32vec2 A; /* print in hex format */

"[1]:A1 [0]:A0"
```

Corresponding Intrinsic: none

The eight 16-bit values of A are placed in the output buffer and printed in the following format (default in decimal):

```
cout << Is16vec8 A;
cout << Iu16vec8 A;
cout << hex << Iu16vec8 A; /* print in hex format */
"[7]:A7 [6]:A6 [5]:A5 [4]:A4 [3]:A3 [2]:A2 [1]:A1 [0]:A0"
```

Corresponding Intrinsic: none

The four 16-bit values of A are placed in the output buffer and printed in the following format (default in decimal):

```
cout << Is16vec4 A;
cout << Iu16vec4 A;
cout << hex << Iu16vec4 A; /* print in hex format */
"[3]:A3 [2]:A2 [1]:A1 [0]:A0"
```

Corresponding Intrinsic: none

The sixteen 8-bit values of A are placed in the output buffer and printed in the following format (default is decimal):

```
cout << Is8vec16 A; cout << Iu8vec16 A; cout << hex << Iu8vec8 A;
/* print in hex format instead of decimal*/
"[15]:A15 [14]:A14 [13]:A13 [12]:A12 [11]:A11 [10]:A10 [9]:A9 [8]:A8
[7]:A7 [6]:A6 [5]:A5 [4]:A4 [3]:A3 [2]:A2 [1]:A1 [0]:A0"
```

Corresponding Intrinsic: none

The eight 8-bit values of A are placed in the output buffer and printed in the following format (default is decimal):

```
cout << Is8vec8 A; cout << Iu8vec8 A; cout << hex << Iu8vec8 A;
/* print in hex format instead of decimal*/
"[7]:A7 [6]:A6 [5]:A5 [4]:A4 [3]:A3 [2]:A2 [1]:A1 [0]:A0"
```

Corresponding Intrinsic: none

Element Access Operators

```

int R = Is64vec2 A[i];

unsigned int R = Iu64vec2 A[i];

int R = Is32vec4 A[i];

unsigned int R = Iu32vec4 A[i];

int R = Is32vec2 A[i];

unsigned int R = Iu32vec2 A[i];

short R = Is16vec8 A[i];

unsigned short R = Iu16vec8 A[i];

short R = Is16vec4 A[i];

unsigned short R = Iu16vec4 A[i];

signed char R = Is8vec16 A[i];

unsigned char R = Iu8vec16 A[i];

signed char R = Is8vec8 A[i];

unsigned char R = Iu8vec8 A[i];

```

Access and read element *i* of *A*. If *DEBUG* is enabled and the user tries to access an element outside of *A*, a diagnostic message is printed and the program aborts.

Corresponding Intrinsics: none

Element Assignment Operators

```

Is64vec2 A[i] = int R;

Is32vec4 A[i] = int R;

Iu32vec4 A[i] = unsigned int R;

Is32vec2 A[i] = int R;

Iu32vec2 A[i] = unsigned int R;

Is16vec8 A[i] = short R;

Iu16vec8 A[i] = unsigned short R;

Is16vec4 A[i] = short R;

```

```
Iu16vec4 A[i] = unsigned short R;
```

```
Is8vec16 A[i] = signed char R;
```

```
Iu8vec16 A[i] = unsigned char R;
```

```
Is8vec8 A[i] = signed char R;
```

```
Iu8vec8 A[i] = unsigned char R;
```

Assign R to element i of A. If DEBUG is enabled and the user tries to assign a value to an element outside of A, a diagnostic message is printed and the program aborts.

Corresponding Intrinsics: none

Unpack Operators

Interleave the 64-bit value from the high half of A with the 64-bit value from the high half of B.

```
I364vec2 unpack_high(I64vec2 A, I64vec2 B);
```

```
Is64vec2 unpack_high(Is64vec2 A, Is64vec2 B);
```

```
Iu64vec2 unpack_high(Iu64vec2 A, Iu64vec2 B);
```

```
R0 = A1;
```

```
R1 = B1;
```

Corresponding intrinsic: `_mm_unpackhi_epi64`

Interleave the two 32-bit values from the high half of A with the two 32-bit values from the high half of B .

```
I32vec4 unpack_high(I32vec4 A, I32vec4 B);
```

```
Is32vec4 unpack_high(Is32vec4 A, Is32vec4 B);
```

```
Iu32vec4 unpack_high(Iu32vec4 A, Iu32vec4 B);
```

```
R0 = A1;
```

```
R1 = B1;
```

```
R2 = A2;
```

```
R3 = B2;
```

Corresponding intrinsic: `_mm_unpackhi_epi32`

Interleave the 32-bit value from the high half of A with the 32-bit value from the high half of B.

```
I32vec2 unpack_high(I32vec2 A, I32vec2 B);
```

```
Is32vec2 unpack_high(Is32vec2 A, Is32vec2 B);
```

```
Iu32vec2 unpack_high(Iu32vec2 A, Iu32vec2 B);
```

```
R0 = A1;
R1 = B1;
```

Corresponding intrinsic: `_mm_unpackhi_pi32`

Interleave the four 16-bit values from the high half of A with the two 16-bit values from the high half of B.

```
I16vec8  unpack_high(I16vec8 A, I16vec8 B);
Is16vec8 unpack_high(Is16vec8 A, Is16vec8 B);
Iu16vec8 unpack_high(Iu16vec8 A, Iu16vec8 B);
```

```
R0 = A2;
R1 = B2;
R2 = A3;
R3 = B3;
```

Corresponding intrinsic: `_mm_unpackhi_epi16`

Interleave the two 16-bit values from the high half of A with the two 16-bit values from the high half of B.

```
I16vec4  unpack_high(I16vec4 A, I16vec4 B);
Is16vec4 unpack_high(Is16vec4 A, Is16vec4 B);
Iu16vec4 unpack_high(Iu16vec4 A, Iu16vec4 B);
```

```
R0 = A2;R1 = B2;
R2 = A3;R3 = B3;
```

Corresponding intrinsic: `_mm_unpackhi_pi16`

Interleave the four 8-bit values from the high half of A with the four 8-bit values from the high half of B.

```
I8vec8  unpack_high(I8vec8 A, I8vec8 B);
Is8vec8 unpack_high(Is8vec8 A, I8vec8 B);
Iu8vec8 unpack_high(Iu8vec8 A, I8vec8 B);
```

```
R0 = A4;
R1 = B4;
R2 = A5;
R3 = B5;
R4 = A6;
R5 = B6;
R6 = A7;
R7 = B7;
```

Corresponding intrinsic: `_mm_unpackhi_pi8`

Interleave the sixteen 8-bit values from the high half of A with the four 8-bit values from the high half of B.

```
I8vec16 unpack_high(I8vec16 A, I8vec16 B);
```

```
Is8vec16 unpack_high(Is8vec16 A, I8vec16 B);
```

```
Iu8vec16 unpack_high(Iu8vec16 A, I8vec16 B);
```

```
R0 = A8;
R1 = B8;
R2 = A9;
R3 = B9;
R4 = A10;
R5 = B10;
R6 = A11;
R7 = B11;
R8 = A12;
R9 = B12;
R2 = A13;
R3 = B13;
R4 = A14;
R5 = B14;
R6 = A15;
R7 = B15;
```

Corresponding intrinsic: `_mm_unpackhi_epi16`

Interleave the 32-bit value from the low half of A with the 32-bit value from the low half of B

```
R0 = A0;
R1 = B0;
```

Corresponding intrinsic: `_mm_unpacklo_epi32`

Interleave the 64-bit value from the low half of A with the 64-bit values from the low half of B

```
I64vec2 unpack_low(I64vec2 A, I64vec2 B);
```

```
Is64vec2 unpack_low(Is64vec2 A, Is64vec2 B);
```

```
Iu64vec2 unpack_low(Iu64vec2 A, Iu64vec2 B);
```

```
R0 = A0;
R1 = B0;
R2 = A1;
R3 = B1;
```

Corresponding intrinsic: `_mm_unpacklo_epi32`

Interleave the two 32-bit values from the low half of A with the two 32-bit values from the low half of B

```
I32vec4 unpack_low(I32vec4 A, I32vec4 B);
```

```
Is32vec4 unpack_low(Is32vec4 A, Is32vec4 B);
```

```
Iu32vec4 unpack_low(Iu32vec4 A, Iu32vec4 B);
```

```
R0 = A0;
R1 = B0;
R2 = A1;
R3 = B1;
```

Corresponding intrinsic: `_mm_unpacklo_epi32`

Interleave the 32-bit value from the low half of A with the 32-bit value from the low half of B.

```
I32vec2  unpack_low(I32vec2 A, I32vec2 B);
Is32vec2 unpack_low(Is32vec2 A, Is32vec2 B);
Iu32vec2 unpack_low(Iu32vec2 A, Iu32vec2 B);

R0 = A0;
R1 = B0;
```

Corresponding intrinsic: `_mm_unpacklo_pi32`

Interleave the two 16-bit values from the low half of A with the two 16-bit values from the low half of B.

```
I16vec8  unpack_low(I16vec8 A, I16vec8 B);
Is16vec8 unpack_low(Is16vec8 A, Is16vec8 B);
Iu16vec8 unpack_low(Iu16vec8 A, Iu16vec8 B);

R0 = A0;
R1 = B0;
R2 = A1;
R3 = B1;
R4 = A2;
R5 = B2;
R6 = A3;
R7 = B3;
```

Corresponding intrinsic: `_mm_unpacklo_epi16`

Interleave the two 16-bit values from the low half of A with the two 16-bit values from the low half of B.

```
I16vec4  unpack_low(I16vec4 A, I16vec4 B);
Is16vec4 unpack_low(Is16vec4 A, Is16vec4 B);
Iu16vec4 unpack_low(Iu16vec4 A, Iu16vec4 B);

R0 = A0;
R1 = B0;
R2 = A1;
R3 = B1;
```

Corresponding intrinsic: `_mm_unpacklo_pi16`

Interleave the four 8-bit values from the high low of A with the four 8-bit values from the low half of B.

```
I8vec16  unpack_low(I8vec16 A, I8vec16 B);
Is8vec16 unpack_low(Is8vec16 A, Is8vec16 B);
```

```
Iu8vec16 unpack_low(Iu8vec16 A, Iu8vec16 B);
```

```
R0 = A0;
R1 = B0;
R2 = A1;
R3 = B1;
R4 = A2;
R5 = B2;
R6 = A3;
R7 = B3;
R8 = A4;
R9 = B4;
R10 = A5;
R11 = B5;
R12 = A6;
R13 = B6;
R14 = A7;
R15 = B7;
```

Corresponding intrinsic: `_mm_unpacklo_epi8`

Interleave the four 8-bit values from the high low of A with the four 8-bit values from the low half of B.

```
I8vec8 unpack_low(I8vec8 A, I8vec8 B);
```

```
Is8vec8 unpack_low(Is8vec8 A, Is8vec8 B);
```

```
Iu8vec8 unpack_low(Iu8vec8 A, Iu8vec8 B);
```

```
R0 = A0;
R1 = B0;
R2 = A1;
R3 = B1;
R4 = A2;
R5 = B2;
R6 = A3;
R7 = B3;
```

Corresponding intrinsic: `_mm_unpacklo_pi8`

Pack Operators

Pack the eight 32-bit values found in A and B into eight 16-bit values with signed saturation.

```
Is16vec8 pack_sat(Is32vec2 A, Is32vec2 B);
Corresponding intrinsic: _mm_packs_epi32
```

Pack the four 32-bit values found in A and B into eight 16-bit values with signed saturation.

```
Is16vec4 pack_sat(Is32vec2 A, Is32vec2 B);
Corresponding intrinsic: _mm_packs_pi32
```

Pack the sixteen 16-bit values found in A and B into sixteen 8-bit values with signed saturation.

```
Is8vec16 pack_sat(Is16vec4 A, Is16vec4 B);
Corresponding intrinsic: _mm_packs_epi16
```

Pack the eight 16-bit values found in A and B into eight 8-bit values with signed saturation.

```
Is8vec8 pack_sat(Is16vec4 A, Is16vec4 B);
Corresponding intrinsic: _mm_packs_pi16
```

Pack the sixteen 16-bit values found in A and B into sixteen 8-bit values with unsigned saturation .

```
Iu8vec16 packu_sat(Is16vec4 A, Is16vec4 B);
Corresponding intrinsic: _mm_packus_epi16
```

Pack the eight 16-bit values found in A and B into eight 8-bit values with unsigned saturation.

```
Iu8vec8 packu_sat(Is16vec4 A, Is16vec4 B);
Corresponding intrinsic: _mm_packs_pu16
```

Clear MMX(TM) Instructions State Operator

Empty the MMX(TM) registers and clear the MMX state. Read the guidelines for using the EMMS instruction intrinsic.

```
void empty(void);
Corresponding intrinsic: _mm_empty
```

Integer Intrinsics for Streaming SIMD Extensions



Note

You must include `fvec.h` header file for the following functionality.

Compute the element-wise maximum of the respective signed integer words in A and B.

```
Is16vec4 simd_max(Is16vec4 A, Is16vec4 B);
Corresponding intrinsic: _mm_max_pi16
```

Compute the element-wise minimum of the respective signed integer words in A and B.

```
Is16vec4 simd_min(Is16vec4 A, Is16vec4 B);
Corresponding intrinsic: _mm_min_pi16
```

Compute the element-wise maximum of the respective unsigned bytes in A and B.

```
Iu8vec8 simd_max(Iu8vec8 A, Iu8vec8 B);
Corresponding intrinsic: _mm_max_pu8
```

Compute the element-wise minimum of the respective unsigned bytes in A and B.

```
Iu8vec8 simd_min(Iu8vec8 A, Iu8vec8 B);
Corresponding intrinsic: _mm_min_pu8
```

Create an 8-bit mask from the most significant bits of the bytes in A.

```
int move_mask(I8vec8 A);
Corresponding intrinsic: _mm_movemask_pi8
```

Conditionally store byte elements of A to address p. The high bit of each byte in the selector B determines whether the corresponding byte in A will be stored.

```
void mask_move(I8vec8 A, I8vec8 B, signed char *p);
Corresponding intrinsic: _mm_maskmove_si64
```

Store the data in A to the address p without polluting the caches. A can be any Ivec type.

```
void store_nta(__m64 *p, M64 A);
Corresponding intrinsic: _mm_stream_pi
```

Compute the element-wise average of the respective unsigned 8-bit integers in A and B.

```
Iu8vec8 simd_avg(Iu8vec8 A, Iu8vec8 B);
Corresponding intrinsic: _mm_avg_pu8
```

Compute the element-wise average of the respective unsigned 16-bit integers in A and B.

```
Iu16vec4 simd_avg(Iu16vec4 A, Iu16vec4 B);
Corresponding intrinsic: _mm_avg_pu16
```

Conversions Between Fvec and Ivec

Convert the lower double-precision floating-point value of A to a 32-bit integer with truncation.

```
int F64vec2ToInt(F64vec42 A);
r := (int)A0;
```

Convert the four floating-point values of A to two the two least significant double-precision floating-point values.

```
F64vec2 F32vec4ToF64vec2(F32vec4 A);
r0 := (double)A0;
r1 := (double)A1;
```

Convert the two double-precision floating-point values of A to two single-precision floating-point values.

```
F32vec4 F64vec2ToF32vec4(F64vec2 A);
r0 := (float)A0;
r1 := (float)A1;
```

Convert the signed int in B to a double-precision floating-point value and pass the upper double-precision value from A through to the result.

```
F64vec2 InttoF64vec2(F64vec2 A, int B);
r0 := (double)B;
r1 := A1;
```

Convert the lower floating-point value of A to a 32-bit integer with truncation.

```
int F32vec4ToInt(F32vec4 A);
r := (int)A0;
```

Convert the two lower floating-point values of A to two 32-bit integer with truncation, returning the integers in packed form.

```
Is32vec2 F32vec4ToIs32vec2 (F32vec4 A);
r0 := (int)A0;
r1 := (int)A1;
```

Convert the 32-bit integer value B to a floating-point value; the upper three floating-point values are passed through from A.

```
F32vec4 IntToF32vec4(F32vec4 A, int B);
r0 := (float)B;
r1 := A1;
r2 := A2;
r3 := A3;
```

Convert the two 32-bit integer values in packed form in B to two floating-point values; the upper two floating-point values are passed through from A.

```
F32vec4 Is32vec2ToF32vec4(F32vec4 A, Is32vec2 B);
r0 := (float)B0;
r1 := (float)B1;
r2 := A2;
r3 := A3;
```

Floating-point Vector Classes

Floating-point Vector Classes

The floating-point vector classes, `F64vec2`, `F32vec4`, and `F32vec1`, provide an interface to SIMD operations. The class specifications are as follows:

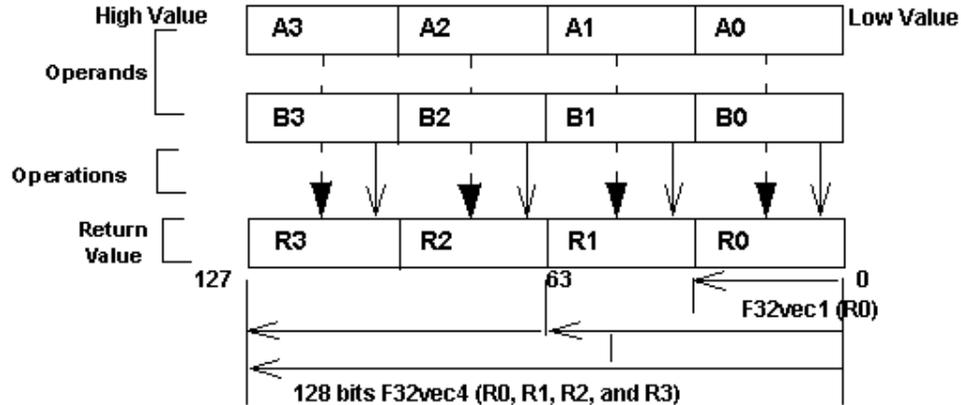
```
F64vec2 A(double x, double y);
```

```
F32vec4 A(float z, float y, float x, float w);
```

```
F32vec1 B(float w);
```

The packed floating-point input values are represented with the right-most value lowest as shown in the following table.

Single-Precision Floating-point Elements



F32vec4 returns **four** packed **single-precision floating point** values (**R0, R1, R2, and R3**).
F32vec2 returns **one** **single-precision floating point** value (**R0**).

Fvec Notation Conventions

This reference uses the following conventions for syntax and return values.

Fvec Classes Syntax Notation

Fvec classes use the syntax conventions shown the following examples:

```
[Fvec_Class] R = [Fvec_Class] A [operator][Ivec_Class] B;
```

Example 1: `F64vec2 R = F64vec2 A & F64vec2 B;`

```
[Fvec_Class] R = [operator]([Fvec_Class] A, [Fvec_Class] B);
```

Example 2: `F64vec2 R = andnot(F64vec2 A, F64vec2 B);`

```
[Fvec_Class] R [operator]= [Fvec_Class] A;
```

Example 3: `F64vec2 R &= F64vec2 A;`

where

[operator] is an operator (for example, `&`, `|`, or `^`)

[Fvec_Class] is any Fvec class (`F64vec2`, `F32vec4`, or `F32vec1`)

R, A, B are declared Fvec variables of the type indicated

Return Value Notation

Because the Fvec classes have packed elements, the return values typically follow the conventions presented in the Return Value Convention Notation Mappings table. `F32vec4` returns four single-

precision, floating-point values (R0, R1, R2, and R3); `F64vec2` returns two double-precision, floating-point values, and `F32vec1` returns the lowest single-precision floating-point value (R0).

Return Value Convention Notation Mappings

Example 1:	Example 2:	Example 3:	F32vec4	F64vec2	F32vec1
<code>R0 := A0 & B0;</code>	<code>R0 := A0 andnot B0;</code>	<code>R0 &= A0;</code>	x	x	x
<code>R1 := A1 & B1;</code>	<code>R1 := A1 andnot B1;</code>	<code>R1 &= A1;</code>	x	x	N/A
<code>R2 := A2 & B2;</code>	<code>R2 := A2 andnot B2;</code>	<code>R2 &= A2;</code>	x	N/A	N/A
<code>R3 := A3 & B3;</code>	<code>R3 := A3 andnot B3;</code>	<code>R3 &= A3;</code>	x	N/A	N/A

Data Alignment

Memory operations using the Streaming SIMD Extensions should be performed on 16-byte-aligned data whenever possible.

`F32vec4` and `F64vec2` object variables are properly aligned by default. Note that floating point arrays are not automatically aligned. To get 16-byte alignment, you can use the alignment `__declspec`:

```
__declspec( align(16) ) float A[4];
```

Conversions

All `Fvec` object variables can be implicitly converted to `__m128` data types. For example, the results of computations performed on `F32vec4` or `F32vec1` object variables can be assigned to `__m128` data types.

```
__m128d mm = A & B; /* where A,B are F64vec2 object variables */
```

```
__m128 mm = A & B; /* where A,B are F32vec4 object variables */
```

```
__m128 mm = A & B; /* where A,B are F32vec1 object variables */
```

Constructors and Initialization

The following table shows how to create and initialize `F32vec` objects with the `Fvec` classes.

Constructors and Initialization for Fvec Classes

Example	Intrinsic	Returns
Constructor Declaration		
<code>F64vec2 A;</code> <code>F32vec4 B;</code> <code>F32vec1 C;</code>	N/A	N/A
__m128 Object Initialization		
<code>F64vec2 A(__m128d mm);</code> <code>F32vec4 B(__m128 mm);</code> <code>F32vec1 C(__m128 mm);</code>	N/A	N/A
Double Initialization		
<code>/* Initializes two doubles. */</code> <code>F64vec2 A(double d0, double d1);</code> <code>F64vec2 A = F64vec2(double d0, double d1);</code>	<code>_mm_set_pd</code>	<code>A0 := d0;</code> <code>A1 := d1;</code>
<code>F64vec2 A(double d0);</code> <code>/* Initializes both return values</code> <code>with the same double precision value */.</code>	<code>_mm_set1_pd</code>	<code>A0 := d0;</code> <code>A1 := d0;</code>
Float Initialization		
<code>F32vec4 A(float f3, float f2,</code> <code>float f1, float f0);</code> <code>F32vec4 A = F32vec4(float f3, float f2,</code> <code>float f1, float f0);</code>	<code>_mm_set_ps</code>	<code>A0 := f0;</code> <code>A1 := f1;</code> <code>A2 := f2;</code> <code>A3 := f3;</code>
<code>F32vec4 A(float f0);</code> <code>/* Initializes all return values</code> <code>with the same floating point value. */</code>	<code>_mm_set1_ps</code>	<code>A0 := f0;</code> <code>A1 := f0;</code> <code>A2 := f0;</code> <code>A3 := f0;</code>
<code>F32vec4 A(double d0);</code> <code>/* Initialize all return values with</code> <code>the same double-precision value. */</code>	<code>_mm_set1_ps(d)</code>	<code>A0 := d0;</code> <code>A1 := d0;</code> <code>A2 := d0;</code> <code>A3 := d0;</code>
<code>F32vec1 A(double d0);</code> <code>/* Initializes the lowest value of A</code> <code>with d0 and the other values with 0.*/</code>	<code>_mm_set_ss(d)</code>	<code>A0 := d0;</code> <code>A1 := 0;</code> <code>A2 := 0;</code> <code>A3 := 0;</code>
<code>F32vec1 B(float f0);</code> <code>/* Initializes the lowest value of B</code> <code>with f0 and the other values with 0.*/</code>	<code>_mm_set_ss</code>	<code>B0 := f0;</code> <code>B1 := 0;</code> <code>B2 := 0;</code> <code>B3 := 0;</code>

Example	Intrinsic	Returns
<pre>F32vec1 B(int I); /* Initializes the lowest value of B with f0, other values are undefined.*/</pre>	<code>_mm_cvtsi32_ss</code>	<pre>B0 := f0; B1 := {} B2 := {} B3 := {}</pre>

Arithmetic Operators

The following table lists the arithmetic operators of the `Fvec` classes and generic syntax. The operators have been divided into standard and advanced operations, which are described in more detail later in this section.

Fvec Arithmetic Operators

Category	Operation	Operators	Generic Syntax
Standard	Addition	+ +=	R = A + B; R += A;
	Subtraction	- -=	R = A - B; R -= A;
	Multiplication	* *=	R = A * B; R *= A;
	Division	/ /=	R = A / B; R /= A;
Advanced	Square Root	<code>sqrt</code>	R = <code>sqrt</code> (A);
	Reciprocal (Newton-Raphson)	<code>rcp</code> <code>rcp_nr</code>	R = <code>rcp</code> (A); R = <code>rcp_nr</code> (A);
	Reciprocal Square Root (Newton-Raphson)	<code>rsqrt</code> <code>rsqrt_nr</code>	R = <code>rsqrt</code> (A); R = <code>rsqrt_nr</code> (A);

Standard Arithmetic Operator Usage

The following two tables show the return values for each class of the standard arithmetic operators, which use the syntax styles described earlier in the Return Value Notation section.

Standard Arithmetic Return Value Mapping

R	A	Operators	B	F32vec4	F64vec2	F32vec1
R0 :=	A0	+ - * /	B0			
R1 :=	A1	+ - * /	B1			N/A
R2 :=	A2	+ - * /	B2		N/A	N/A
R3 :=	A3	+ - * /	B3		N/A	N/A

Arithmetic with Assignment Return Value Mapping

R	Operators				A	F32vec4	F64vec2	F32vec1
R0 :=	+=	-=	*=	/=	A0			
R1 :=	+=	-=	*=	/=	A1			N/A
R2 :=	+=	-=	*=	/=	A2		N/A	N/A
R3 :=	+=	-=	*=	/=	A3		N/A	N/A

This table lists standard arithmetic operator syntax and intrinsics.

Standard Arithmetic Operations for Fvec Classes

Operation	Returns	Example Syntax Usage	Intrinsic
Addition	4 floats	F32vec4 R = F32vec4 A + F32vec4 B; F32vec4 R += F32vec4 A;	_mm_add_ps
	2 doubles	F64vec2 R = F64vec2 A + F32vec2 B; F64vec2 R += F64vec2 A;	_mm_add_pd
	1 float	F32vec1 R = F32vec1 A + F32vec1 B; F32vec1 R += F32vec1 A;	_mm_add_ss
Subtraction	4 floats	F32vec4 R = F32vec4 A - F32vec4 B; F32vec4 R -= F32vec4 A;	_mm_sub_ps
	2 doubles	F64vec2 R = F64vec2 A - F32vec2 B; F64vec2 R -= F64vec2 A;	_mm_sub_pd
	1 float	F32vec1 R = F32vec1 A - F32vec1 B; F32vec1 R -= F32vec1 A;	_mm_sub_ss
Multiplication	4 floats	F32vec4 R = F32vec4 A * F32vec4 B; F32vec4 R *= F32vec4 A;	_mm_mul_ps
	2 doubles	F64vec2 R = F64vec2 A * F32vec2 B; F64vec2 R *= F64vec2 A;	_mm_mul_pd
	1 float	F32vec1 R = F32vec1 A * F32vec1 B; F32vec1 R *= F32vec1 A;	_mm_mul_ss
Division	4 floats	F32vec4 R = F32vec4 A / F32vec4 B; F32vec4 R /= F32vec4 A;	_mm_div_ps
	2 doubles	F64vec2 R = F64vec2 A / F64vec2 B; F64vec2 R /= F64vec2 A;	_mm_div_pd

Operation	Returns	Example Syntax Usage	Intrinsic
	1 float	F32vec1 R = F32vec1 A / F32vec1 B; F32vec1 R /= F32vec1 A;	_mm_div_ss

Advanced Arithmetic Operator Usage

The following table shows the return values classes of the advanced arithmetic operators, which use the syntax styles described earlier in the Return Value Notation section.

Advanced Arithmetic Return Value Mapping

R	Operators	A	F32vec4	F64vec2	F32vec1				
R0:=	sqrt	rcp	rsqrt	rcp_nr	rsqrt_nr	A0			
R1:=	sqrt	rcp	rsqrt	rcp_nr	rsqrt_nr	A1			N/A
R2:=	sqrt	rcp	rsqrt	rcp_nr	rsqrt_nr	A2		N/A	N/A
R3:=	sqrt	rcp	rsqrt	rcp_nr	rsqrt_nr	A3		N/A	N/A
f:=	add_horizontal			(A0 + A1 + A2 + A3)				N/A	N/A
d:=	add_horizontal			(A0 + A1)			N/A		N/A

This table shows examples for advanced arithmetic operators.

Advanced Arithmetic Operations for Fvec Classes

Returns	Example Syntax Usage	Intrinsic
Square Root		
4 floats	F32vec4 R = sqrt(F32vec4 A);	_mm_sqrt_ps
2 doubles	F64vec2 R = sqrt(F64vec2 A);	_mm_sqrt_pd
1 float	F32vec1 R = sqrt(F32vec1 A);	_mm_sqrt_ss
Reciprocal		
4 floats	F32vec4 R = rcp(F32vec4 A);	_mm_rcp_ps
2 doubles	F64vec2 R = rcp(F64vec2 A);	_mm_rcp_pd
1 float	F32vec1 R = rcp(F32vec1 A);	_mm_rcp_ss
Reciprocal Square Root		
4 floats	F32vec4 R = rsqrt(F32vec4 A);	_mm_rsqrt_ps

Returns	Example Syntax Usage	Intrinsic
2 doubles	<code>F64vec2 R = rsqrt(F64vec2 A);</code>	<code>_mm_rsqrt_pd</code>
1 float	<code>F32vec1 R = rsqrt(F32vec1 A);</code>	<code>_mm_rsqrt_ss</code>
Reciprocal Newton Raphson		
4 floats	<code>F32vec4 R = rcp_nr(F32vec4 A);</code>	<code>_mm_sub_ps</code> <code>_mm_add_ps</code> <code>_mm_mul_ps</code> <code>_mm_rcp_ps</code>
2 doubles	<code>F64vec2 R = rcp_nr(F64vec2 A);</code>	<code>_mm_sub_pd</code> <code>_mm_add_pd</code> <code>_mm_mul_pd</code> <code>_mm_rcp_pd</code>
1 float	<code>F32vec1 R = rcp_nr(F32vec1 A);</code>	<code>_mm_sub_ss</code> <code>_mm_add_ss</code> <code>_mm_mul_ss</code> <code>_mm_rcp_ss</code>
Reciprocal Square Root Newton Raphson		
4 float	<code>F32vec4 R = rsqrt_nr(F32vec4 A);</code>	<code>_mm_sub_pd</code> <code>_mm_mul_pd</code> <code>_mm_rsqrt_ps</code>
2 doubles	<code>F64vec2 R = rsqrt_nr(F64vec2 A);</code>	<code>_mm_sub_pd</code> <code>_mm_mul_pd</code> <code>_mm_rsqrt_pd</code>
1 float	<code>F32vec1 R = rsqrt_nr(F32vec1 A);</code>	<code>_mm_sub_ss</code> <code>_mm_mul_ss</code> <code>_mm_rsqrt_ss</code>
Horizontal Add		
1 float	<code>float f = add_horizontal(F32vec4 A);</code>	<code>_mm_add_ss</code> <code>_mm_shuffle_ss</code>
1 double	<code>double d = add_horizontal(F64vec2 A);</code>	<code>_mm_add_sd</code> <code>_mm_shuffle_sd</code>

Minimum and Maximum Operators

Compute the minimums of the two double precision floating-point values of A and B.

```
F64vec2 R = simd_min(F64vec2 A, F64vec2 B)
R0 := min(A0,B0);
R1 := min(A1,B1);
Corresponding intrinsic: _mm_min_pd
```

Compute the minimums of the four single precision floating-point values of A and B.

```
F32vec4 R = simd_min(F32vec4 A, F32vec4 B)
R0 := min(A0,B0);
R1 := min(A1,B1);
```

```
R2 := min(A2,B2);
R3 := min(A3,B3);
Corresponding intrinsic: _mm_min_ps
```

Compute the minimum of the lowest single precision floating-point values of A and B.

```
F32vec1 R = simd_min(F32vec1 A, F32vec1 B)
R0 := min(A0,B0);
Corresponding intrinsic: _mm_min_ss
```

Compute the maximums of the two double precision floating-point values of A and B.

```
F64vec2 simd_max(F64vec2 A, F64vec2 B)
R0 := max(A0,B0);
R1 := max(A1,B1);
Corresponding intrinsic: _mm_max_pd
```

Compute the maximums of the four single precision floating-point values of A and B.

```
F32vec4 R = simd_max(F32vec4 A, F32vec4 B)
R0 := max(A0,B0);
R1 := max(A1,B1);
R2 := max(A2,B2);
R3 := max(A3,B3);
Corresponding intrinsic: _mm_max_ps
```

Compute the maximum of the lowest single precision floating-point values of A and B.

```
F32vec1 simd_max(F32vec1 A, F32vec1 B)
R0 := max(A0,B0);
Corresponding intrinsic: _mm_max_ss
```

Logical Operators

The following table lists the logical operators of the Fvec classes and generic syntax. The logical operators for F32vec1 classes use only the lower 32 bits.

Fvec Logical Operators Return Value Mapping

Bitwise Operation	Operators	Generic Syntax
AND	& &=	R = A & B; R &= A;
OR	 =	R = A B; R = A;
XOR	^ ^=	R = A ^ B; R ^= A;
andnot	andnot	R = andnot(A);

The following table lists standard logical operators syntax and corresponding intrinsics. Note that there is no corresponding scalar intrinsic for the F32vec1 classes, which accesses the lower 32 bits of the packed vector intrinsics.

Logical Operations for Fvec Classes

Operation	Returns	Example Syntax Usage	Intrinsic
AND	4 floats	F32vec4 R = F32vec4 A & F32vec4 B; F32vec4 R &= F32vec4 A;	_mm_and_ps
	2 doubles	F64vec2 R = F64vec2 A & F32vec2 B; F64vec2 R &= F64vec2 A;	_mm_and_pd
	1 float	F32vec1 R = F32vec1 A & F32vec1 B; F32vec1 R &= F32vec1 A;	_mm_and_ps
OR	4 floats	F32vec4 R = F32vec4 A F32vec4 B; F32vec4 R = F32vec4 A;	_mm_or_ps
	2 doubles	F64vec2 R = F64vec2 A F32vec2 B; F64vec2 R = F64vec2 A;	_mm_or_pd
	1 float	F32vec1 R = F32vec1 A F32vec1 B; F32vec1 R = F32vec1 A;	_mm_or_ps
XOR	4 floats	F32vec4 R = F32vec4 A ^ F32vec4 B; F32vec4 R ^= F32vec4 A;	_mm_xor_ps
	2 doubles	F64vec2 R = F64vec2 A ^ F364vec2 B; F64vec2 R ^= F64vec2 A;	_mm_xor_pd
	1 float	F32vec1 R = F32vec1 A ^ F32vec1 B; F32vec1 R ^= F32vec1 A;	_mm_xor_ps
ANDNOT	2 doubles	F64vec2 R = andnot(F64vec2 A, F64vec2 B);	_mm_andnot_pd

Compare Operators

The operators described in this section compare the single precision floating-point values of A and B. Comparison between objects of any `Fvec` class return the same class being compared.

The following table lists the compare operators for the `Fvec` classes.

Compare Operators and Corresponding Ininsics

Compare For:	Operators	Syntax
Equality	<code>cmpeq</code>	<code>R = cmpeq(A, B)</code>
Inequality	<code>cmpneq</code>	<code>R = cmpneq(A, B)</code>
Greater Than	<code>cmpgt</code>	<code>R = cmpgt(A, B)</code>
Greater Than or Equal To	<code>cmpge</code>	<code>R = cmpge(A, B)</code>
Not Greater Than	<code>cmpngt</code>	<code>R = cmpngt(A, B)</code>
Not Greater Than or Equal To	<code>cmpnge</code>	<code>R = cmpnge(A, B)</code>
Less Than	<code>cmplt</code>	<code>R = cmplt(A, B)</code>
Less Than or Equal To	<code>cmple</code>	<code>R = cmple(A, B)</code>
Not Less Than	<code>cmpnlt</code>	<code>R = cmpnlt(A, B)</code>
Not Less Than or Equal To	<code>cmpnle</code>	<code>R = cmpnle(A, B)</code>

Compare Operators

The mask is set to `0xffffffff` for each floating-point value where the comparison is true and `0x00000000` where the comparison is false. The following table shows the return values for each class of the compare operators, which use the syntax described earlier in the Return Value Notation section.

Compare Operator Return Value Mapping

R	A0	For Any Operators	B	If True	If False	F32vec4	F64vec2	F32vec1
R0:=	(A1 !(A1	cmp[eq lt le gt ge] cmp[ne nlt nle ngt nge]	B1) B1)	0xffffffff	0x00000000	X	X	X
R1:=	(A1 !(A1	cmp[eq lt le gt ge] cmp[ne nlt nle ngt nge]	B2) B2)	0xffffffff	0x00000000	X	X	N/A
R2:=	(A1 !(A1	cmp[eq lt le gt ge] cmp[ne nlt nle ngt nge]	B3) B3)	0xffffffff	0x00000000	X	N/A	N/A
R3:=	A3	cmp[eq lt le gt ge] cmp[ne nlt nle ngt nge]	B3) B3)	0xffffffff	0x00000000	X	N/A	N/A

The following table shows examples for arithmetic operators and intrinsics.

Compare Operations for Fvec Classes

Returns	Example Syntax Usage	Intrinsic
Compare for Equality		
4 floats	F32vec4 R = cmpeq(F32vec4 A);	_mm_cmpeq_ps
2 doubles	F64vec2 R = cmpeq(F64vec2 A);	_mm_cmpeq_pd
1 float	F32vec1 R = cmpeq(F32vec1 A);	_mm_cmpeq_ss
Compare for Inequality		
4 floats	F32vec4 R = cmpneq(F32vec4 A);	_mm_cmpneq_ps
2 doubles	F64vec2 R = cmpneq(F64vec2 A);	_mm_cmpneq_pd
1 float	F32vec1 R = cmpneq(F32vec1 A);	_mm_cmpneq_ss
Compare for Less Than		
4 floats	F32vec4 R = cmplt(F32vec4 A);	_mm_cmplt_ps
2 doubles	F64vec2 R = cmplt(F64vec2 A);	_mm_cmplt_pd
1 float	F32vec1 R = cmplt(F32vec1 A);	_mm_cmplt_ss
Compare for Less Than or Equal		
4 floats	F32vec4 R = cmple(F32vec4 A);	_mm_cmple_ps
2 doubles	F64vec2 R = cmple(F64vec2 A);	_mm_cmple_pd
1 float	F32vec1 R = cmple(F32vec1 A);	_mm_cmple_ss
Compare for Greater Than		
4 floats	F32vec4 R = cmpgt(F32vec4 A);	_mm_cmpgt_ps
2 doubles	F64vec2 R = cmpgt(F64vec2 A);	_mm_cmpgt_pd
1 float	F32vec1 R = cmpgt(F32vec1 A);	_mm_cmpgt_ss
Compare for Greater Than or Equal To		
4 floats	F32vec4 R = cmpge(F32vec4 A);	_mm_cmpge_ps
2 doubles	F64vec2 R = cmpge(F64vec2 A);	_mm_cmpge_pd
1 float	F32vec1 R = cmpge(F32vec1 A);	_mm_cmpge_ss
Compare for Not Less Than		
4 floats	F32vec4 R = cmpnlt(F32vec4 A);	_mm_cmpnlt_ps
2 doubles	F64vec2 R = cmpnlt(F64vec2 A);	_mm_cmpnlt_pd
1 float	F32vec1 R = cmpnlt(F32vec1 A);	_mm_cmpnlt_ss
Compare for Not Less Than or Equal		

4 floats	F32vec4 R = cmpnle(F32vec4 A);	_mm_cmpnle_ps
2 doubles	F64vec2 R = cmpnle(F64vec2 A);	_mm_cmpnle_pd
1 float	F32vec1 R = cmpnle(F32vec1 A);	_mm_cmpnle_ss
Compare for Not Greater Than		
4 floats	F32vec4 R = cmpngt(F32vec4 A);	_mm_cmpngt_ps
2 doubles	F64vec2 R = cmpngt(F64vec2 A);	_mm_cmpngt_pd
1 float	F32vec1 R = cmpngt(F32vec1 A);	_mm_cmpngt_ss
Compare for Not Greater Than or Equal		
4 floats	F32vec4 R = cmpnge(F32vec4 A);	_mm_cmpnge_ps
2 doubles	F64vec2 R = cmpnge(F64vec2 A);	_mm_cmpnge_pd
1 float	F32vec1 R = cmpnge(F32vec1 A);	_mm_cmpnge_ss

Conditional Select Operators for Fvec Classes

Each conditional function compares single-precision floating-point values of A and B. The C and D parameters are used for return value. Comparison between objects of any Fvec class returns the same class.

Conditional Select Operators for Fvec Classes

Conditional Select for:	Operators	Syntax
Equality	select_eq	R = select_eq(A, B)
Inequality	select_neq	R = select_neq(A, B)
Greater Than	select_gt	R = select_gt(A, B)
Greater Than or Equal To	select_ge	R = select_ge(A, B)
Not Greater Than	select_gt	R = select_gt(A, B)
Not Greater Than or Equal To	select_ge	R = select_ge(A, B)
Less Than	select_lt	R = select_lt(A, B)
Less Than or Equal To	select_le	R = select_le(A, B)
Not Less Than	select_nlt	R = select_nlt(A, B)
Not Less Than or Equal To	select_nle	R = select_nle(A, B)

Conditional Select Operator Usage

For conditional select operators, the return value is stored in C if the comparison is true or in D if false. The following table shows the return values for each class of the conditional select operators, using the Return Value Notation described earlier.

Compare Operator Return Value Mapping

R	A0	Operators	B	C	D	F32vec4	F64vec2	F32vec1
R0:=	(A1 !(A1	select_[eq lt le gt ge] select_[ne nlt nle ngt nge]	B0) B0)	C0 C0	D0 D0	X	X	X
R1:=	(A2 !(A2	select_[eq lt le gt ge] select_[ne nlt nle ngt nge]	B1) B1)	C1 C1	D1 D1	X	X	N/A
R2:=	(A2 !(A2	select_[eq lt le gt ge] select_[ne nlt nle ngt nge]	B2) B2)	C2 C2	D2 D2	X	N/A	N/A
R3:=	(A3 !(A3	select_[eq lt le gt ge] select_[ne nlt nle ngt nge]	B3) B3)	C3 C3	D3 D3	X	N/A	N/A

The following table shows examples for conditional select operations and corresponding intrinsics.

Conditional Select Operations for Fvec Classes

Returns	Example Syntax Usage	Intrinsic
Compare for Equality		
4 floats	F32vec4 R = select_eq(F32vec4 A);	_mm_cmpeq_ps
2 doubles	F64vec2 R = select_eq(F64vec2 A);	_mm_cmpeq_pd
1 float	F32vec1 R = select_eq(F32vec1 A);	_mm_cmpeq_ss
Compare for Inequality		
4 floats	F32vec4 R = select_neq(F32vec4 A);	_mm_cmpneq_ps
2 doubles	F64vec2 R = select_neq(F64vec2 A);	_mm_cmpneq_pd
1 float	F32vec1 R = select_neq(F32vec1 A);	_mm_cmpneq_ss
Compare for Less Than		
4 floats	F32vec4 R = select_lt(F32vec4 A);	_mm_cmplt_ps
2 doubles	F64vec2 R = select_lt(F64vec2 A);	_mm_cmplt_pd

1 float	F32vec1 R = select_lt(F32vec1 A);	_mm_cmplt_ss
Compare for Less Than or Equal		
4 floats	F32vec4 R = select_le(F32vec4 A);	_mm_cmple_ps
2 doubles	F64vec2 R = select_le(F64vec2 A);	_mm_cmple_pd
1 float	F32vec1 R = select_le(F32vec1 A);	_mm_cmple_ss
Compare for Greater Than		
4 floats	F32vec4 R = select_gt(F32vec4 A);	_mm_cmpgt_ps
2 doubles	F64vec2 R = select_gt(F64vec2 A);	_mm_cmpgt_pd
1 float	F32vec1 R = select_gt(F32vec1 A);	_mm_cmpgt_ss
Compare for Greater Than or Equal To		
4 floats	F32vec4 R = select_ge(F32vec4 A);	_mm_cmpge_ps
2 doubles	F64vec2 R = select_ge(F64vec2 A);	_mm_cmpge_pd
1 float	F32vec1 R = select_ge(F32vec1 A);	_mm_cmpge_ss
Compare for Not Less Than		
4 floats	F32vec4 R = select_nlt(F32vec4 A);	_mm_cmpnlt_ps
2 doubles	F64vec2 R = select_nlt(F64vec2 A);	_mm_cmpnlt_pd
1 float	F32vec1 R = select_nlt(F32vec1 A);	_mm_cmpnlt_ss
Compare for Not Less Than or Equal		
4 floats	F32vec4 R = select_nle(F32vec4 A);	_mm_cmpnle_ps
2 doubles	F64vec2 R = select_nle(F64vec2 A);	_mm_cmpnle_pd
1 float	F32vec1 R = select_nle(F32vec1 A);	_mm_cmpnle_ss
Compare for Not Greater Than		
4 floats	F32vec4 R = select_ngt(F32vec4 A);	_mm_cmpngt_ps
2 doubles	F64vec2 R = select_ngt(F64vec2 A);	_mm_cmpngt_pd
1 float	F32vec1 R = select_ngt(F32vec1 A);	_mm_cmpngt_ss
Compare for Not Greater Than or Equal		
4 floats	F32vec4 R = select_nge(F32vec4 A);	_mm_cmpnge_ps
2 doubles	F64vec2 R = select_nge(F64vec2 A);	_mm_cmpnge_pd
1 float	F32vec1 R = select_nge(F32vec1 A);	_mm_cmpnge_ss

Cacheability Support Operations

Stores (non-temporal) the two double-precision, floating-point values of A. Requires a 16-byte aligned address.

```
void store_nta(double *p, F64vec2 A);
Corresponding intrinsic: _mm_stream_pd
```

Stores (non-temporal) the four single-precision, floating-point values of A. Requires a 16-byte aligned address.

```
void store_nta(float *p, F32vec4 A);
Corresponding intrinsic: _mm_stream_ps
```

Debugging

The debug operations do not map to any compiler intrinsics for MMX(TM) technology or Streaming SIMD Extensions. They are provided for debugging programs only. Use of these operations may result in loss of performance, so you should not use them outside of debugging.

Output Operations

The two single, double-precision floating-point values of A are placed in the output buffer and printed in decimal format as follows:

```
cout << F64vec2 A;
"[1]:A1 [0]:A0"
Corresponding intrinsics: none
```

The four, single-precision floating-point values of A are placed in the output buffer and printed in decimal format as follows:

```
cout << F32vec4 A;
"[3]:A3 [2]:A2 [1]:A1 [0]:A0"
Corresponding intrinsics: none
```

The lowest, single-precision floating-point value of A is placed in the output buffer and printed.

```
cout << F32vec1 A;
Corresponding intrinsics: none
```

Element Access Operations

```
double d = F64vec2 A[int i]
```

Read one of the two, double-precision floating-point values of A without modifying the corresponding floating-point value. Permitted values of *i* are 0 and 1. For example:

If DEBUG is enabled and *i* is not one of the permitted values (0 or 1), a diagnostic message is printed and the program aborts.

```
double d = F64vec2 A[1];
Corresponding intrinsics: none
```

Read one of the four, single-precision floating-point values of A without modifying the corresponding floating point value. Permitted values of `i` are 0, 1, 2, and 3. For example:

```
float f = F32vec4 A[int i]
```

If DEBUG is enabled and `i` is not one of the permitted values (0-3), a diagnostic message is printed and the program aborts.

```
float f = F32vec4 A[2];
Corresponding intrinsics: none
```

Element Assignment Operations

```
F64vec4 A[int i] = double d;
```

Modify one of the two, double-precision floating-point values of A. Permitted values of `int i` are 0 and 1. For example:

```
F32vec4 A[1] = double d;
F32vec4 A[int i] = float f;
```

Modify one of the four, single-precision floating-point values of A. Permitted values of `int i` are 0, 1, 2, and 3. For example:

If DEBUG is enabled and `int i` is not one of the permitted values (0-3), a diagnostic message is printed and the program aborts.

```
F32vec4 A[3] = float f;
Corresponding intrinsics: none.
```

Load and Store Operators

Loads two, double-precision floating-point values, copying them into the two, floating-point values of A. No assumption is made for alignment.

```
void loadu(F64vec2 A, double *p)
Corresponding intrinsic: _mm_loadu_pd
```

Stores the two, double-precision floating-point values of A. No assumption is made for alignment.

```
void storeu(float *p, F64vec2 A);
Corresponding intrinsic: _mm_storeu_pd
```

Loads four, single-precision floating-point values, copying them into the four floating-point values of A. No assumption is made for alignment.

```
void loadu(F32vec4 A, double *p)
Corresponding intrinsic: _mm_loadu_ps
```

Stores the four, single-precision floating-point values of A. No assumption is made for alignment.

```
void storeu(float *p, F32vec4 A);
Corresponding intrinsic: _mm_storeu_ps
```

Unpack Operators for Fvec Operators

Selects and interleaves the lower, double-precision floating-point values from A and B.

```
F64vec2 R = unpack_low(F64vec2 A, F64vec2 B);
Corresponding intrinsic: _mm_unpacklo_pd(a, b)
```

Selects and interleaves the higher, double-precision floating-point values from A and B.

```
F64vec2 R = unpack_high(F64vec2 A, F64vec2 B);
Corresponding intrinsic: _mm_unpackhi_pd(a, b)
```

Selects and interleaves the lower two, single-precision floating-point values from A and B.

```
F32vec4 R = unpack_low(F32vec4 A, F32vec4 B);
Corresponding intrinsic: _mm_unpacklo_ps(a, b)
```

Selects and interleaves the higher two, single-precision floating-point values from A and B.

```
F32vec4 R = unpack_high(F32vec4 A, F32vec4 B);
Corresponding intrinsic: _mm_unpackhi_ps(a, b)
```

Move Mask Operator

Creates a 2-bit mask from the most significant bits of the two, double-precision floating-point values of A, as follows:

```
int i = move_mask(F64vec2 A)
i := sign(a1)<<1 | sign(a0)<<0
Corresponding intrinsic: _mm_movemask_pd
```

Creates a 4-bit mask from the most significant bits of the four, single-precision floating-point values of A, as follows:

```
int i = move_mask(F32vec4 A)
i := sign(a3)<<3 | sign(a2)<<2 | sign(a1)<<1 | sign(a0)<<0
Corresponding intrinsic: _mm_movemask_ps
```

Classes Quick Reference

This appendix contains tables listing the class, functionality, and corresponding intrinsics for each class in the Intel® C++ Class Libraries for SIMD Operations. The following table lists all Intel C++ Compiler intrinsics that are not implemented in the C++ SIMD classes.

Logical Operators: Corresponding Intrinsics and Classes

Operators	Corresponding Intrinsic	I128vec1, I64vec2, I32vec4, I16vec8, I8vec16	I64vec, I32vec, I16vec, I8vec8	F64vec2	F32vec4	F32vec1
&, &=	<code>_mm_and_[x]</code>	si128	si64	pd	ps	ps
, =	<code>_mm_or_[x]</code>	si128	si64	pd	ps	ps
^, ^=	<code>_mm_xor_[x]</code>	si128	si64	pd	ps	ps
Andnot	<code>_mm_andnot_[x]</code>	si128	si64	pd	N/A	N/A

Arithmetic: Corresponding Intrinsics and Classes, Part 1

Operators	Corresponding Intrinsic	I64vec2	I32vec4	I16vec8	I8vec16
+, +=	<code>_mm_add_[x]</code>	epi64	epi32	epi16	epi8
-, -=	<code>_mm_sub_[x]</code>	epi64	epi32	epi16	epi8
*, *=	<code>_mm_mullo_[x]</code>	N/A	N/A	epi16	N/A
/, /=	<code>_mm_div_[x]</code>	N/A	N/A	N/A	N/A
mul_high	<code>_mm_mulhi_[x]</code>	N/A	N/A	epi16	N/A
mul_add	<code>_mm_madd_[x]</code>	N/A	N/A	epi16	N/A
sqrt	<code>_mm_sqrt_[x]</code>	N/A	N/A	N/A	N/A
rcp	<code>_mm_rcp_[x]</code>	N/A	N/A	N/A	N/A
rcp_nr	<code>_mm_rcp_[x]</code> <code>_mm_add_[x]</code> <code>_mm_sub_[x]</code> <code>_mm_mul_[x]</code>	N/A	N/A	N/A	N/A
rsqrt	<code>_mm_rsqrt_[x]</code>	N/A	N/A	N/A	N/A
rsqrt_nr	<code>_mm_rsqrt_[x]</code> <code>_mm_sub_[x]</code> <code>_mm_mul_[x]</code>	N/A	N/A	N/A	N/A

Arithmetic: Corresponding Intrinsics and Classes, Part 2

Operators	Corresponding Intrinsic	I32vec2	I16vec4	I8vec8	F64vec2	F32vec4	F32vec1
+, +=	_mm_add_[x]	pi32	pi16	pi8	pd	ps	ss
-, -=	_mm_sub_[x]	pi32	pi16	pi8	pd	ps	ss
*, *=	_mm_mullo_[x]	N/A	pi16	N/A	pd	ps	ss
/, /=	_mm_div_[x]	N/A	N/A	N/A	pd	ps	ss
mul_high	_mm_mulhi_[x]	N/A	pi16	N/A	N/A	N/A	N/A
mul_add	_mm_madd_[x]	N/A	pi16	N/A	N/A	N/A	N/A
sqrt	_mm_sqrt_[x]	N/A	N/A	N/A	pd	ps	ss
rcp	_mm_rcp_[x]	N/A	N/A	N/A	pd	ps	ss
rcp_nr	_mm_rcp_[x] _mm_add_[x] _mm_sub_[x] _mm_mul_[x]	N/A	N/A	N/A	pd	ps	ss
rsqrt	_mm_rsqrt_[x]	N/A	N/A	N/A	pd	ps	ss
rsqrt_nr	_mm_rsqrt_[x] _mm_sub_[x] _mm_mul_[x]	N/A	N/A	N/A	pd	ps	ss

Shift Operators: Corresponding Intrinsics and Classes, Part 1

Operators	Corresponding Intrinsic	I128vec1	I64vec2	I32vec4	I16vec8	I8vec16
>>, >>=	_mm_srl_[x] _mm_srli_[x] _mm_sra_[x] _mm_srai_[x]	N/A N/A N/A N/A	epi64 epi64 N/A N/A	epi32 epi32 epi32 epi32	epi16 epi16 epi16 epi16	N/A N/A N/A N/A
<<, <<=	_mm_sll_[x] _mm_slli_[x]	N/A N/A	epi64 epi64	epi32 epi32	epi16 epi16	N/A N/A

Shift Operators: Corresponding Intrinsics and Classes, Part 2

Operators	Corresponding Intrinsic	l64vec1	l32vec2	l16vec4	l8vec8
>>, >>=	_mm_srl_[x]	si64	pi32	pi16	N/A
	_mm_srli_[x]	si64	pi32	pi16	N/A
	_mm_sra_[x]	N/A	pi32	pi16	N/A
	_mm_srai_[x]	N/A	pi32	pi16	N/A
<<, <<=	_mm_sll_[x]	si64	pi32	pi16	N/A
	_mm_slli_[x]	si64	pi32	pi16	N/A

Comparison Operators: Corresponding Intrinsics and Classes, Part 1

Operators	Corresponding Intrinsic	l32vec4	l16vec8	l8vec16	l32vec2	l16vec4	l8vec8
cmpeq	_mm_cmpeq_[x]	epi32	epi16	epi8	pi32	pi16	pi8
cmpneq	_mm_cmpeq_[x] _mm_andnot_[y]*	epi32 sil28	epi16 sil28	epi8 sil28	pi32 si64	pi16 si64	pi8 si64
cmpgt	_mm_cmpgt_[x]	epi32	epi16	epi8	pi32	pi16	pi8
cmpge	_mm_cmpge_[x] _mm_andnot_[y]*	epi32 sil28	epi16 sil28	epi8 sil28	pi32 si64	pi16 si64	pi8 si64
cmplt	_mm_cmplt_[x]	epi32	epi16	epi8	pi32	pi16	pi8
cmple	_mm_cmple_[x] _mm_andnot_[y]*	epi32 sil28	epi16 sil28	epi8 sil28	pi32 si64	pi16 si64	pi8 si64
cmpngt	_mm_cmpngt_[x]	epi32	epi16	epi8	pi32	pi16	pi8
cmpnge	_mm_cmpnge_[x]	N/A	N/A	N/A	N/A	N/A	N/A
cmnpnlt	_mm_cmnpnlt_[x]	N/A	N/A	N/A	N/A	N/A	N/A
cmnpnle	_mm_cmnpnle_[x]	N/A	N/A	N/A	N/A	N/A	N/A

* Note that `_mm_andnot_[y]` intrinsics do not apply to the `fvec` classes.

Comparison Operators: Corresponding Intrinsics and Classes, Part 2

Operators	Corresponding Intrinsic	F64vec2	F32vec4	F32vec1
cmpeq	_mm_cmpeq_[x]	pd	ps	ss
cmpneq	_mm_cmpeq_[x] _mm_andnot_[y]*	pd	ps	ss
cmpgt	_mm_cmpgt_[x]	pd	ps	ss
cmpge	_mm_cmpge_[x] _mm_andnot_[y]*	pd	ps	ss
cmplt	_mm_cmplt_[x]	pd	ps	ss
cmple	_mm_cmple_[x] _mm_andnot_[y]*	pd	ps	ss
cmpngt	_mm_cmpngt_[x]	pd	ps	ss
cmpnge	_mm_cmpnge_[x]	pd	ps	ss
cmnpnlt	_mm_cmnpnlt_[x]	pd	ps	ss
cmnpnle	_mm_cmnpnle_[x]	pd	ps	ss

Conditional Select Operators: Corresponding Intrinsics and Classes, Part 1

Operators	Corresponding Intrinsic	I32vec4	I16vec8	I8vec16	I32vec2	I16vec4	I8vec8
select_eq	_mm_cmpeq_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	epi32 si128 si128 si128	epi16 si128 si128 si128	epi8 si128 si128 si128	pi32 si64 si64 si64	pi16 si64 si64 si64	pi8 si64 si64 si64
select_neq	_mm_cmpeq_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	epi32 si128 si128 si128	epi16 si128 si128 si128	epi8 si128 si128 si128	pi32 si64 si64 si64	pi16 si64 si64 si64	pi8 si64 si64 si64
select_gt	_mm_cmpgt_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	epi32 si128 si128 si128	epi16 si128 si128 si128	epi8 si128 si128 si128	pi32 si64 si64 si64	pi16 si64 si64 si64	pi8 si64 si64 si64
select_ge	_mm_cmpge_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	epi32 si128 si128 si128	epi16 si128 si128 si128	epi8 si128 si128 si128	pi32 si64 si64 si64	pi16 si64 si64 si64	pi8 si64 si64 si64
select_lt	_mm_cmplt_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	epi32 si128 si128 si128	epi16 si128 si128 si128	epi8 si128 si128 si128	pi32 si64 si64 si64	pi16 si64 si64 si64	pi8 si64 si64 si64
select_le	_mm_cmple_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	epi32 si128 si128 si128	epi16 si128 si128 si128	epi8 si128 si128 si128	pi32 si64 si64 si64	pi16 si64 si64 si64	pi8 si64 si64 si64

Operators	Corresponding Intrinsic	I32vec4	I16vec8	I8vec16	I32vec2	I16vec4	I8vec8
select_ngt	_mm_cmpgt_[x]	N/A	N/A	N/A	N/A	N/A	N/A
select_nge	_mm_cmpge_[x]	N/A	N/A	N/A	N/A	N/A	N/A
select_nlt	_mm_cmplt_[x]	N/A	N/A	N/A	N/A	N/A	N/A
select_nle	_mm_cmple_[x]	N/A	N/A	N/A	N/A	N/A	N/A

* Note that `_mm_andnot_[y]` intrinsics do not apply to the `fvec` classes.

Conditional Select Operators: Corresponding Intrinsics and Classes, Part 2

Operators	Corresponding Intrinsic	F64vec2	F32vec4	F32vec1
select_eq	_mm_cmpeq_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	pd	ps	ss
select_neq	_mm_cmpeq_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	pd	ps	ss
select_gt	_mm_cmpgt_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	pd	ps	ss
select_ge	_mm_cmpge_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	pd	ps	ss
select_lt	_mm_cmplt_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	pd	ps	ss
select_le	_mm_cmple_[x] _mm_and_[y] _mm_andnot_[y]* _mm_or_[y]	pd	ps	ss
select_ngt	_mm_cmpgt_[x]	pd	ps	ss
select_nge	_mm_cmpge_[x]	pd	ps	ss
select_nlt	_mm_cmplt_[x]	pd	ps	ss
select_nle	_mm_cmple_[x]	pd	ps	ss

Packing and Unpacking Operators: Corresponding Intrinsics and Classes, Part 1

Operators	Corresponding Intrinsic	I64vec2	I32vec4	I16vec8	I8vec16	I32vec2
unpack_high	_mm_unpackhi_[x]	epi64	epi32	epi16	epi8	pi32
unpack_low	_mm_unpacklo_[x]	epi64	epi32	epi16	epi8	pi32
pack_sat	_mm_packs_[x]	N/A	epi32	epi16	N/A	pi32
packu_sat	_mm_packus_[x]	N/A	N/A	epi16	N/A	N/A
sat_add	_mm_adds_[x]	N/A	N/A	epi16	epi8	N/A
sat_sub	_mm_subs_[x]	N/A	N/A	epi16	epi8	N/A

Packing and Unpacking Operators: Corresponding Intrinsics and Classes, Part 2

Operators	Corresponding Intrinsic	I16vec4	I8vec8	F64vec2	F32vec4	F32vec1
unpack_high	_mm_unpackhi_[x]	pi16	pi8	pd	ps	N/A
unpack_low	_mm_unpacklo_[x]	pi16	pi8	pd	ps	N/A
pack_sat	_mm_packs_[x]	pi16	N/A	N/A	N/A	N/A
packu_sat	_mm_packus_[x]	pu16	N/A	N/A	N/A	N/A
sat_add	_mm_adds_[x]	pi16	pi8	pd	ps	ss
sat_sub	_mm_subs_[x]	pi16	pi8	pi16	pi8	pd

Conversions Operators: Corresponding Intrinsics and Classes

Operators	Corresponding Intrinsic
F64vec2ToInt	_mm_cvttss_sd
F32vec4ToF64vec2	_mm_cvtps_pd
F64vec2ToF32vec4	_mm_cvtpd_ps
IntToF64vec2	_mm_cvtsi32_sd
F32vec4ToInt	_mm_cvtt_ss2si
F32vec4ToI32vec2	_mm_cvttps_pi32
IntToF32vec4	_mm_cvtsi32_ss
I32vec2ToF32vec4	_mm_cvtpi32_ps

Programming Example

This sample program uses the `F32vec4` class to average the elements of a 20 element floating point array.

```
// Include Streaming SIMD Extension Class Definitions
#include <fvec.h>

// Shuffle any 2 single precision floating point from a
// into low 2 SP FP and shuffle any 2 SP FP from b
// into high 2 SP FP of destination
#define SHUFFLE(a,b,i) (F32vec4)_mm_shuffle_ps(a,b,i)
#include <stdio.h>
#define SIZE 20

// Global variables
float result;
_MM_ALIGN 16 float array[SIZE];

//*****
// Function: Add20ArrayElements
// Add all the elements of a 20 element array
//*****

void Add20ArrayElements (F32vec4 *array, float *result)
{
    F32vec4 vec0, vec1;
    vec0 = _mm_load_ps ((float *) array); // Load array's first
    4 floats

    //*****
    // Add all elements of the array, 4 elements at a time
    //*****

    vec0 += array[1]; // Add elements 5-8
    vec0 += array[2]; // Add elements 9-12
    vec0 += array[3]; // Add elements 13-16
    vec0 += array[4]; // Add elements 17-20

    //*****
    // There are now 4 partial sums.
    // Add the 2 lowers to the 2 raises,
    // then add those 2 results together
    //*****

    vec1 = SHUFFLE(vec1, vec0, 0x40);
    vec0 += vec1;
    vec1 = SHUFFLE(vec1, vec0, 0x30);
    vec0 += vec1;
    vec0 = SHUFFLE(vec0, vec0, 2);
    _mm_store_ss (result, vec0); // Store the final sum
}

void main(int argc, char *argv[])
{
    int i;
    // Initialize the array
    for (i=0; i < SIZE; i++)
    {
        array[i] = (float) i;
    }
}
```

```
// Call function to add all array elements
Add20ArrayElements (array, &result);

// Print average array element value
printf ("Average of all array values = %f\n", result/20.);
printf ("The correct answer is %f\n\n\n", 9.5);
}
```