

# **Fujitsu Software Compiler Package V1.0L21**

## **C++ User's Guide**

J2UL-2583-03ENZ0(04)  
October 2023

# Preface

---

## Purpose of This Manual

This manual describes how to use the C++ language processing system (called "this system" in this manual) intended for the PRIMEHPC FX700 system (called "FX system" in this manual).

It also details the language specifications based on the following C++ language standard, and OpenMP API, and which are expanded under this system.

- Programming languages -- C++(ISO/IEC 14882:2003)  
In this manual, this standard is called C++03.  
Specifications based on this standard are called C++03 specifications.
- Programming languages -- C++(ISO/IEC 14882:2011)  
In this manual, this standard is called C++11.  
Specifications based on this standard are called C++11 specifications.
- Programming languages -- C++(ISO/IEC 14882:2014)  
In this manual, this standard is called C++14.  
Specifications based on this standard are called C++14 specifications.
- Programming languages -- C++ (ISO/IEC 14882:2017)  
In this manual, this standard is called C++17.  
Specifications based on this standard are called C++17 specifications.
- Programming languages -- C(ISO/IEC 9899:1999)  
In this manual, this standard is called C99.  
Specifications based on this standard are called C99 specifications.
- OpenMP Application Program Interface Version 3.1 July 2011  
In this manual, specification based on this standard is called OpenMP 3.1.
- OpenMP Application Program Interface Version 4.0 July 2013  
In this manual, specification based on this standard is called OpenMP 4.0.
- OpenMP Application Programming Interface Version 4.5 November 2015  
In this manual, specification based on this standard is called OpenMP 4.5.
- OpenMP Application Programming Interface Version 5.0 November 2018  
In this manual, specification based on this standard is called OpenMP 5.0.

## Intended Readers

This manual assumes an understanding of the language specifications set forth in the C++ language standard. For details about standard specifications, refer to the C++ language standard or commercial books based on the standard.

This manual was written for people who write C++ programs and process them using this system. Readers should also know how to use Linux(R) commands, file manipulation and shell programming.

## Structure of This Manual

The structure of this manual is as follows:

### [Chapter 1 Overview](#)

Describes an overview of how to use this system to process C++ programs.

## [Chapter 2 From Compilation to Execution](#)

Describes the procedures from compiling to executing C++ programs.

## [Chapter 3 Optimization](#)

Describes how to use optimization functions that ensure programs execute as quickly as possible.

## [Chapter 4 Multiprocessing](#)

Describes how to use parallelization functions or OpenMP specifications that ensure programs execute as quickly as possible.

## [Chapter 5 Emitting Information](#)

Describes how to emit information in this system.

## [Chapter 6 Language Specifications](#)

Describes the language specifications that are extended in this system based on the C++ language standard, and the support status of the language specifications.

## [Chapter 7 Notes on Linking with Different Languages and Trad/Clang Modes](#)

Describes notes about linking with object programs in different programming languages and Trad/Clang Modes.

## [Chapter 8 Debugging Functions](#)

Describes the program debugging functions in this system.

## [Chapter 9 Clang Mode](#)

Describes the mode using a compiler that has been enhanced based on Clang/LLVM.

## [Appendix A Precompiled Headers](#)

Describes how to use precompiled headers in this system.

## [Appendix B Restrictions and Notes](#)

Describes the restrictions and notes of this system.

## [Appendix C Compatibility with GNU C++ Specifications](#)

Describes compatibility with GNU C++.

## [Appendix D Data and Memory Regions](#)

Describes the attributes of data and memory regions in this system.

## [Appendix E Code Coverage](#)

Describes the code coverage.

## [Appendix F Fujitsu Extended Functions](#)

Describes Fujitsu Extended Functions supported in this system.

## [Appendix G Runtime Information Output Function](#)

Describes the Runtime Information Output Function.

## [Appendix H Using High-Speed Facility](#)

Describes information on the compilation and execution processes for using the high-speed facility on FX system.

## [Appendix I Multi-Process Execution](#)

Describes how to execute a program in a multi-process and notes.

## [Appendix J Fujitsu OpenMP Library](#)

Describes how to perform parallel processing of a C++ program using Fujitsu OpenMP libraries.

## [Appendix K Large Page Library](#)

Describes the large page library of the HPC (High Performance Computing) extension function for FX system that extend the capabilities of standard Linux.

## Notes of This Manual

The code examples of optimization in this manual are conceptual source that complements each explanation of functions. When the code examples are compiled and executed, optimizations may not work as expected. This is because optimizations depend on compilation options and other conditions.

## Notation Used in This Manual

### Syntax Description Symbols

A syntax description symbol is a symbol that has specific meaning in syntax.

The following syntax symbols are used in this manual.

Symbol name	Symbol	Explanation
Selection symbols	{ }	Only one of the items enclosed in the braces must be selected (items are listed vertically).
		Multiple items are enumerated by this delimiter (items are listed horizontally).
Option symbol	[ ]	An item enclosed in brackets can be omitted. This symbol includes the meaning of the selection symbol "{ }".
Repeat symbol	...	The item immediately preceding the ellipsis can be specified repeatedly in the syntax.

### Parallel

The way of parallelization is thread parallelization unless otherwise described.

## Export Controls

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

## Trademark

- Arm is trademark or registered trademark of Arm Limited (or its subsidiaries) in the US and/or elsewhere.
- OpenMP is a trademark of OpenMP Architecture Review Board.
- Linux(R) is the registered trademark of Linus Torvalds in the U.S. and other countries.
- All other trademarks and product names are the property of their respective owners.
- The trademark notice symbol (TM,(R)) is not necessarily added in the system name and the product name, etc. published in this material.

## Date of Publication and Version

Version	Manual code
October 2023, Version 3.4	J2UL-2583-03ENZ0(04)
April 2023, Version 3.3	J2UL-2583-03ENZ0(03)
October 2022, Version 3.2	J2UL-2583-03ENZ0(02)
April 2022, Version 3.1	J2UL-2583-03ENZ0(01)
January 2022, 3rd Version	J2UL-2583-03ENZ0(00)
August 2021, Version 2.7	J2UL-2583-02ENZ0(07)
July 2021, Version 2.6	J2UL-2583-02ENZ0(06)
March 2021, Version 2.5	J2UL-2583-02ENZ0(05)
January 2021, Version 2.4	J2UL-2583-02ENZ0(04)

Version	Manual code
December 2020, Version 2.3	J2UL-2583-02ENZ0(03)
November 2020, Version 2.2	J2UL-2583-02ENZ0(02)
September 2020, Version 2.1	J2UL-2583-02ENZ0(01)
July 2020, 2nd Version	J2UL-2583-02ENZ0(00)
February 2020, 1st Version	J2UL-2583-01ENZ0(00)

## Copyright

Copyright FUJITSU LIMITED 2020-2023

## Update History

Changes	Location	Version
Improved the explanation of the following options: <ul style="list-style-type: none"> <li>- -K{fp_relaxed nofp_relaxed}</li> <li>- -K{preex nopreex}</li> <li>- -N{Rtrap Rnotrap}</li> </ul>	2.2.2.6	Version 3.4
Improved the explanation of the following optimization control specifiers: <ul style="list-style-type: none"> <li>- fission_point</li> <li>- fp_relaxed</li> <li>- nopreex</li> </ul>	3.4.1.2	
Improved the explanation of the following options: <ul style="list-style-type: none"> <li>- -O[n]</li> <li>- -f{fj-fast-matmul fj-no-fast-matmul}</li> <li>- -f{fj-optlib-string fj-no-optlib-string}</li> <li>- -f{lto no-lto}</li> </ul>	9.1.2.2.3	
Improved the explanation.	9.1.2.3.2	
Added the section "C++ Standard Library Header <filesystem>"	9.5.3	
Corrected description of supported OpenMP specifications.	1.1 4.3 4.3.2.2 I.3	Version 3.3
Added a note.	1.1	
Corrected the table "Return Values of the Compile Command".	2.1.3	
Corrected the explanation.	8.1.2	
Improved the explanation of the following options: <ul style="list-style-type: none"> <li>- -f{fj-swp fj-no-swp}</li> <li>- -f{fj-zfill[=N] fj-no-zfill}</li> <li>- -f{lto no-lto}</li> </ul>	9.1.2.2.3	
Deleted the following option: <ul style="list-style-type: none"> <li>- -f{fj-loop-interchange fj-no-loop-interchange}</li> </ul>	9.1.2.2.3 9.1.2.3.1	

Changes	Location	Version
Added the following optimization control specifiers: - #pragma fj loop swp - #pragma fj loop noswp	9.2.2.1.3	
Added the explanation.	9.5.1	
Corrected the explanation of the following option: - -fvisibility=internal	9.8.2	
Deleted the following GNU C++ compatible option of Trad Mode: - -fvisibility	C.2	
Improved figure design.	-	
Corrected the explanation.	6.1.6 9.1.5	Version 3.2
Added the following option: - -f{debug-info-for-profiling no-debug-info-for-profiling}	9.1.2.2.1	
Corrected the explanation of the following option: - -f{optimize-sibling-calls no-optimize-sibling-calls}	9.8.2	
Added the section "Restrictions and Notes".	9.10	
Corrected the explanation.	H.1	
Unified notation of "C++ Standard Library".	-	
Improved the explanation.	-	
Improved the explanation of the following options: - -O[n] - -msve-vector-bits={512 scalable} - -f{lto no-lto} - -f{omit-frame-pointer no-omit-frame-pointer} - -f{slp-vectorize no-slp-vectorize}	9.1.2.2.3	
Added the following option: - -m{omit-leaf-frame-pointer no-omit-leaf-frame-pointer}	9.1.2.1 9.1.2.2.3	
Added the explanation.	9.1.2.3.1	
Added the section "SIMD".	9.2.3	
Added the section "Restriction of registers with asm keyword and on an inline asm clobber lists".	B.2.13	
Improved the explanation.	-	
Improved the explanation of the following options: - -K{region_extension noregion_extension} - -K{striping[=N] nostriping}	2.2.2.6	3rd Version
Added some standard library functions to the list of optimization targets of the -Klib option.	2.2.2.6	
Added the section "Intrinsic functions that SIMD Extensions can be applied to".	3.2.7.5	
Improved the explanation of the following optimization control specifiers:	3.4.1.2	

Changes	Location	Version
- fission_point - striping		
Modified the table "Optimization Identifiers that can be specified for parallelization".	4.2.6.1	
Added the section "Restriction on Inheritance of a Constructor with Default Parameters Using "using" Inheritance between Inner Classes".	B.1.3	
Changed the description of the environment variable XOS_MMM_L_ARENA_LOCK_TYPE.	K.3.3	
Improved the explanation.	-	
Improved the description of "Predefined Macro Name".	6.1.6 9.1.5	Version 2.7
Added a note about SIMD built-in function documentation.	9.7	
Added a note.	1.2.1	Version 2.6
Added the following option: - -Kopenmp_loop_variable={private standard}	2.2.2.3 2.2.2.6	
Changed the -K{array_declaration_opt noarray_declaration_opt} option default to the -Knoarray_declaration_opt.	2.2.2.6	
Added the following environment variable: - FCOMP_LINK_FJOB	2.3	
Added the explanation.	3.2.7.4	
Improved the explanation.	3.3.9	
Changed the title.	Chapter 7 9.6	
Added the section "Compile Commands and Required Options at Linking".	7.1	
Added the following options: - -f{fj-regalloc-using-latency fj-no-regalloc-using-latency} - -f{fj-promote-licm-addressing fj-no-promote-licm-addressing} - -f{fj-sched-insn-contiguous fj-no-sched-insn-contiguous}	9.1.2.2.3	
Added a note.	9.4.1.2.4	
Added the section "Link Error (undefined reference to)".	B.2.12	
Deleted a note.	2.1.2	Version 2.5
Improved the explanation about the zfill optimization.	2.2.2.6 3.3.5 3.4.1.2	
Improved the explanation of the following options: - -K{fp_relaxed nofp_relaxed} - -K{ilfunc[={loop procedure}] noilfunc}	2.2.2.6	
Added the following environment variable: - FLIB_L1_SCCR_CNTL	3.5.2.2 H.2 J.4.3	
Added the explanation.	3.6.1	
Improved the explanation about supported specifications.	4.3	
Added the section "Creating Shared Libraries in Clang Mode".	4.3.4.5	

Changes	Location	Version
Added the section "Notes on Compilation".	9.1.1.4	
Added the following option: - <code>-msve-vector-bits={512 scalable}</code>	9.1.2.1 9.1.2.2.3	
Added the following options: - <code>-f{fj-interleave-loop-insns[=<i>N</i>] fj-no-interleave-loop-insns}</code> - <code>-f{fj-loop-fission fj-no-loop-fission}</code> - <code>-ffj-loop-fission-threshold=<i>N</i></code> - <code>-f{fj-swp fj-no-swp}</code> - <code>-f{fj-zfill[=<i>N</i>] fj-no-zfill}</code>	9.1.2.2.3 9.1.2.3.1	
Added an article.	9.1.2.3.4	
Added the section "Notes on Specified SVE Vector Register Size".	9.1.2.3.5	
Added the following optimization control specifiers: - <code>#pragma fj loop clone <i>var</i>==<i>n</i></code> - <code>#pragma fj loop loop_fission_target [<i>cl</i>]</code> - <code>#pragma fj loop loop_fission_threshold <i>n</i></code> - <code>#pragma fj loop zfill [<i>N</i>]</code> - <code>#pragma fj loop nozfill</code>	9.2.2.1.2 9.2.2.1.3	
Added the following optimization control specifiers: - <code>#pragma fj loop swp</code> - <code>#pragma fj loop noswp</code> - <code>#pragma clang loop vectorize_width(<i>n</i>, scalable)</code>	9.2.2.1.2	
Added the following information to the optimization information. - Loop fission - Software pipelining - Clone optimization	9.4.1.2.2	
Added the section "About Compile Time for Template Functions".	B.2.11	
Improved the explanation.	-	
Added a note.	4.2.2.4 4.3.2.1 J.2.2 J.3.2	Version 2.4
Added a table of commands and required options for interlanguage linking.	Chapter 7 9.6	
Added examples of compiling and linking.	7.2 9.6.3	
Added loop interchange to the optimizations applied by the <code>-O2</code> option or higher.	9.1.2.2.3	
Added the appendix "Multi-Process Execution".	Appendix I	
Added the section "Multi-Process Execution".	J.5	
Improved the explanation.	-	
Improved the notes on linking with the object of Trad Mode.	9.6.1	Version 2.3



Changes	Location	Version
Added a note.	2.1.2	Version 2.2
Corrected the errors in the explanations of the <code>-N{reordered_variable_stack noreordered_variable_stack}</code> option.	2.2.2.7 D.4	
Improved the explanation about the <code>-N{Rtrap Rnotrap}</code> option.	2.2.2.7	
Improved the explanation about loop unrolling.	3.2.4	
Corrected the errors in the explanations of detected errors.	8.2.1	
Corrected the errors in the explanations of the following options. - <code>-ffj-lst[={plt}]</code> - <code>-ffj-lst-out=<i>file</i></code>	9.1.2.2.1 9.4.1.2	
Added the section "Integer Division Exception when the Divisor is Zero".	B.2.10	
Improved the explanation.	-	
Improved code examples.	3.3.6 3.4.1.2 3.5.2.1 3.5.2.2	Version 2.1
Improved the explanation about the parameters aligned and unaligned for the <code>simd</code> specifier.	3.4.1.2 3.4.1.3	
Added the section "Notes of Wrong Erroneous Program".	3.6.4	
Added the following optimization control specifier: - <code>#pragma clang loop vectorize(assume_safety)</code>	9.2.2.1.2	
Added the section "Changing SVE Vector Register Size".	B.2.9	
Added the appendix "Large Page Library".	Appendix J	
Improved the explanation.	-	2nd Version
Added the supported language specifications.	1.1 6.2.4	
Changed the supported OpenMP specifications.	1.1 4.3 I.3	
Added the following option: - <code>{-help --help}</code>	2.2.1 2.2.2.2	
Added clone optimization to the optimizations applied by the <code>-O3</code> option.	2.2.2.3	
Added/improved the following options: - <code>-K{array_declaration_opt noarray_declaration_opt}</code> - <code>-K{extract_stride_store noextract_stride_store}</code> - <code>-K{fp_precision nofp_precision}</code> - <code>-K{loop_fission_stripmining[={M L1 L2}] loop_nofission_stripmining}</code> - <code>-Kloop_fission_threshold=<i>N</i></code> - <code>-K{loop_perfect_nest loop_noperfect_nest}</code> - <code>-K{prefetch_stride[={soft hard_auto hard_always}] prefetch_nostride}</code> - <code>-K{preload nopreload}</code> - <code>-K{rdconv[={1 2}] nordconv}</code>	2.2.2.3 2.2.2.6	

Changes	Location	Version
<ul style="list-style-type: none"> <li>- -K{simd_uncounted_loop simd_nouncounted_loop}</li> <li>- -K{simd_use_multiple_structures simd_nouse_multiple_structures}</li> <li>- -Kswp_policy={auto small large}</li> <li>- -Ktls_size={12 24 32 48}</li> </ul>		
Deleted the following option: <ul style="list-style-type: none"> <li>- -K{lto no to}</li> </ul>	2.2.2.3 2.2.2.6	
Added the following option: <ul style="list-style-type: none"> <li>- -std={c++17 gnu++17}</li> </ul>	2.2.2.4	
Changed the parameter default value of the -Kilfunc[={loop procedure}] option.	2.2.2.6	
Changed the target loops of the -Kloop_fission option.	2.2.2.6	
Added the following environment variables: <ul style="list-style-type: none"> <li>- FCCpx_trad_ENV</li> <li>- FCC_trad_ENV</li> <li>- FCOMP_UNRECOGNIZED_OPTION</li> </ul>	2.3	
Added the following compilation profile files: <ul style="list-style-type: none"> <li>- FCCpx_trad_PROF</li> <li>- FCC_trad_PROF</li> </ul>	2.4	
Added the following environment variable: FLIB_TRACEBACK_MEM_SIZE	2.5.1	
Deleted the section "Link Time Optimization".	3.3	
Added the section "Strip-mining".	3.3.10.1	
Enabled "#pragma fj" as the format of optimization control lines.	3.4.1.1	
Added the following optimization control specifiers: <ul style="list-style-type: none"> <li>- [no]array_declaration_opt</li> <li>- [no]extract_stride_store</li> <li>- [no]fullunroll_pre_simd [<i>n</i>]</li> <li>- loop_[no]fission_stripmining [<i>n</i>"L1" "L2"]</li> <li>- loop_fission_target [cl s]</li> <li>- loop_fission_threshold <i>n</i></li> <li>- loop_[no]perfect_nest</li> <li>- prefetch_cache_level <i>c-level</i></li> <li>- prefetch_[no]conditional</li> <li>- prefetch_[no]indirect</li> <li>- prefetch_[no]infer</li> <li>- prefetch_iteration <i>n</i></li> <li>- prefetch_iteration_L2 <i>n</i></li> <li>- prefetch_line <i>n</i></li> <li>- prefetch_line_L2 <i>n</i></li> </ul>	3.4.1.2	

Changes	Location	Version
<ul style="list-style-type: none"> <li>- prefetch_nosequential</li> <li>- prefetch_[no]stride [soft hard_auto hard_always]</li> <li>- [no]preload</li> <li>- scache_isolate_assign <i>array1[,array2]...</i>, end_scache_isolate_assign</li> <li>- scache_isolate_way L2=<i>n1</i> [L1=<i>n2</i>], end_scache_isolate_way</li> <li>- simd_[no]use_multiple_structures</li> <li>- swp_policy {auto small large}</li> </ul>		
Deleted the following optimization control specifier: <ul style="list-style-type: none"> <li>- loop_[no]fission</li> </ul>	3.4.1.2	
Added the section "Optimizations Using the Hardware Stride Prefetcher".	3.6.3	
Changed the default value of the environment variable OMP_PROC_BIND in LLVM OpenMP Library.	4.2.2.4 4.3.2.1 4.3.3	
Added information about when an error occurred.	4.3.4.2	
Added and changed some predefined macros.	6.1.6 9.1.5	
Added the following option: <ul style="list-style-type: none"> <li>- --coverage</li> </ul>	9.1.2.1 9.1.2.2.1	
Added the following option: <ul style="list-style-type: none"> <li>- -mfj-tls-size={ 12 24 32 48}</li> </ul>	9.1.2.1 9.1.2.2.6	
Added the following options: <ul style="list-style-type: none"> <li>- -f{fj-line fj-no-line}</li> <li>- -fprofile-dir=<i>dir_name</i></li> </ul>	9.1.2.2.1	
Added the following options: <ul style="list-style-type: none"> <li>- -f{fj-fast-matmul fj-no-fast-matmul}</li> <li>- -f{fj-fp-precision fj-no-fp-precision}</li> <li>- -f{fj-fp-relaxed fj-no-fp-relaxed}</li> <li>- -f{fj-hpctag fj-no-hpctag}</li> <li>- -f{fj-loop-interchange fj-no-loop-interchange}</li> <li>- -f{fj-ocl fj-no-ocl}</li> <li>- -f{fj-optlib-string fj-no-optlib-string}</li> <li>- -f{fj-prefetch-strong fj-no-prefetch-strong}</li> <li>- -f{fj-prefetch-strong-L2 fj-no-prefetch-strong-L2}</li> </ul>	9.1.2.2.3	
Changed the parameter default value of the -ffj-ilfunc[={loop procedure}] option.	9.1.2.2.3	
Changed the libc++ version used by the option -stdlib=libc++.	9.1.2.2.4	
Improved the following options: <ul style="list-style-type: none"> <li>- -march=<i>arch[+features]</i>...</li> </ul>	9.1.2.2.5	

Changes	Location	Version
- <code>-mcpu=<i>cpuname</i>[+<i>features</i>]</code> ...		
Deleted the following option: - <code>-mtune=<i>cpu</i></code>	9.1.2.1 9.1.2.2.5	
Deleted the following option: - <code>-f{<i>plt</i> <i>no-plt</i>}</code>	9.1.2.2.6	
Added notes to "Notes on Compile Options"	9.1.2.3	
Added the following environment variables: - <code>FCCcpx_clang_ENV</code> - <code>FCC_clang_ENV</code>	9.1.3	
Added the following compilation profile files: - <code>FCCcpx_clang_PROF</code> - <code>FCC_clang_PROF</code>	9.1.4	
Added the section "Procedure of Execution".	9.1.6	
Added the following optimization control line: - <code>#pragma fj <i>Optimization-control-specifier</i></code>	9.2.2.1	
Added the following optimization control specifier: - <code>#pragma clang fp contract(fast)</code>	9.2.2.1.2	
Added "Information That Relates To The Register" to the optimization information.	9.4.1.2.2	
Added the section "SIMD Built-in Functions".	9.7	
Added the section "Code Coverage".	9.9	
Added the section "Debug using the debugger".	B.2.8	
Added the following attribute: - <code>aarch64_vector_pcs</code>	C.1	
Added the following GNU C++ compatible options: - <code>-ffp-contract=fast</code> - <code>-floop-parallelize-all</code> - <code>-fprofile-dir=<i>path</i></code> - <code>-funroll-loops</code> - <code>-f{unsafe-math-optimizations no-unsafe-math-optimizations}</code> - <code>-mmodel={small large}</code>	C.2	
Deleted the following GNU C++ compatible option: - <code>-frename-registers</code>	C.2	
Added the appendix "Using High-speed Facility".	Appendix H	
Improved the explanation.	-	
Changed the look according to product upgrades.	-	

All rights reserved.  
The information in this manual is subject to change without notice.

# Contents

---

Chapter 1 Overview.....	1
1.1 Configuration of the C++ Language Processing System.....	1
1.2 How to Use.....	2
1.2.1 Preparation.....	2
1.2.2 Compilation and Linking.....	3
1.2.3 Debugging.....	3
1.2.4 Tuning.....	4
Chapter 2 From Compilation to Execution.....	5
2.1 Compile Command.....	5
2.1.1 Syntax of the Compile Command.....	5
2.1.2 Input Files for the Compile Command.....	5
2.1.3 Return Value of the Compile Command.....	6
2.2 Compiler Options.....	6
2.2.1 Syntax of Compiler Options.....	6
2.2.2 Description of Compiler Options.....	7
2.2.2.1 General Options for Compiler.....	7
2.2.2.2 Options for Messages.....	11
2.2.2.3 Options for Optimization.....	12
2.2.2.4 Options for Language Specifications.....	14
2.2.2.5 Options for PCH.....	16
2.2.2.6 -K Option.....	16
2.2.2.7 -N Option.....	35
2.2.3 Notes for Using Compiler Options.....	41
2.3 Environment Variable for Compile Command.....	41
2.4 Compilation Profile File.....	44
2.5 Procedure of Execution.....	45
2.5.1 Environment Variable for Execution.....	45
2.5.2 Notes for Execution.....	46
2.5.2.1 Variable Allocation at Execution.....	46
Chapter 3 Optimization.....	47
3.1 Overview of Optimization.....	47
3.2 Standard Optimization.....	47
3.2.1 Elimination of Common Expressions.....	48
3.2.2 Movement of Invariant Expressions.....	48
3.2.3 Reducing Strength of Operators.....	48
3.2.4 Loop Unrolling.....	49
3.2.5 Loop Blocking.....	50
3.2.6 Software Pipelining.....	51
3.2.7 SIMD.....	52
3.2.7.1 Normal SIMD.....	52
3.2.7.2 SIMD Extension for Loop Containing "if" Statement.....	52
3.2.7.3 List Vector Conversion.....	52
3.2.7.4 SIMD with Redundant Executions for the SIMD Width.....	52
3.2.7.5 Math functions that SIMD Extensions can be applied to .....	54
3.2.8 Loop Unswitching.....	54
3.2.9 Inline Expansion.....	54
3.3 Extended Optimization.....	54
3.3.1 Optimization by Modifying Evaluation Methods.....	55
3.3.1.1 Advance Evaluation of Invariant Expressions.....	55
3.3.1.2 Arithmetic Evaluation Method Modification.....	55
3.3.2 Optimization of Pointers.....	56
3.3.3 Multi-Operation Function.....	57
3.3.3.1 About Calling of Multi-Operation Functions.....	57
3.3.3.2 Effects of Compiler Option -Kmfunc=3.....	60

3.3.4 Loop Striping.....	61
3.3.5 zfill.....	61
3.3.6 Loop Versioning.....	62
3.3.7 Clone Optimization.....	62
3.3.8 Unroll-and-Jam.....	63
3.3.9 Tree-Height-Reduction Optimization.....	63
3.3.10 Loop Fission.....	64
3.3.10.1 Strip-Mining.....	65
3.3.11 Strict Aliasing.....	66
3.3.12 Fujitsu Matrix Template Library.....	67
3.3.12.1 Common Definition.....	67
3.3.12.1.1 Class Constant.....	67
3.3.12.1.2 Member Function.....	68
3.3.12.2 matrix.....	69
3.3.12.3 banded_matrix.....	71
3.3.12.4 hermitian_matrix.....	73
3.3.12.5 symmetric_matrix.....	75
3.3.12.6 triangular_matrix.....	76
3.3.12.7 sparse_matrix.....	78
3.3.12.8 Usage Example.....	80
3.4 Using Optimization Functions.....	81
3.4.1 Using Optimization Control Line (OCL).....	81
3.4.1.1 Types of Optimization Control Lines.....	81
3.4.1.2 Optimization Control Specifier.....	82
3.4.1.3 Notes for Optimization Control Specifiers.....	120
3.5 Software Control of Sector Cache.....	121
3.5.1 Use of Sector Cache.....	121
3.5.2 Controlling Sector Cache via Software.....	121
3.5.2.1 Software Control with Optimization Control Lines.....	121
3.5.2.2 Software Control with Environment Variables and Optimization Control Line.....	122
3.5.2.3 Behavior when an Exceptional Value Is Specified.....	124
3.6 Notes.....	124
3.6.1 Side Effect of Optimizations for Floating-Point Operation.....	124
3.6.2 Notes on Specified SVE Vector Register Size.....	126
3.6.3 Optimizations Using the Hardware Stride Prefetcher.....	126
3.6.4 Notes of Wrong Erroneous Program.....	126
<b>Chapter 4 Multiprocessing.....</b>	<b>127</b>
4.1 Overview of Multiprocessing.....	127
4.1.1 What is Multiprocessing?.....	127
4.1.2 Effect of Multiprocessing.....	128
4.1.3 Requirements for Effective Multiprocessing.....	128
4.1.4 Parallelization.....	128
4.2 Automatic Parallelization.....	128
4.2.1 Compilation.....	129
4.2.1.1 Compiler Option for Automatic Parallelization.....	129
4.2.2 Execution Process.....	129
4.2.2.1 Number of Threads.....	129
4.2.2.2 Stack Size on Execution.....	129
4.2.2.3 Synchronization Process.....	129
4.2.2.4 CPU Binding for Thread.....	130
4.2.3 Example of Compilation and Execution.....	132
4.2.4 Performance Tuning.....	133
4.2.5 Feature Details on Automatic Parallelization.....	133
4.2.5.1 Targets for Automatic Parallelization.....	133
4.2.5.2 Loop Slicing.....	133
4.2.5.3 Automatic Loop Slicing by Compiler.....	133

4.2.5.4 Loop Interchange and Automatic Loop Slicing.....	134
4.2.5.5 Loop Distribution and Automatic Loop Slicing.....	134
4.2.5.6 Loop Fusion and Automatic Loop Slicing.....	135
4.2.5.7 Loop Reduction.....	135
4.2.5.8 Restrictions on Loop Slicing.....	136
4.2.5.9 Displaying the State of Automatic Parallelization.....	138
4.2.5.10 Parallel Region Extension.....	138
4.2.5.11 Block Distribution and Cyclic Distribution.....	139
4.2.6 Optimization Control Line.....	139
4.2.6.1 Optimization Control Specifier.....	139
4.2.6.2 Automatic Parallelization and Optimization Control Specifiers.....	141
4.2.6.3 Optimization Control Specifiers for Automatic Parallelization.....	141
4.2.7 Notes on Automatic Parallelization.....	151
4.2.7.1 Multiprocessing of Nested Loops.....	152
4.2.7.2 Side Effect of Using -kparallel,reduction.....	153
4.2.7.3 Examples of Invalid Usage of an Optimization Control Line.....	153
4.2.7.4 Standard Library Function References.....	155
4.2.7.5 Note on C++ Programs.....	155
4.3 Parallelization by OpenMP Specification.....	155
4.3.1 Compilation.....	156
4.3.1.1 Compiler Option for OpenMP C++ Program.....	156
4.3.1.2 Obtaining of Optimization Information.....	156
4.3.1.3 Restriction of OpenMP programs.....	157
4.3.2 Execution Process.....	157
4.3.2.1 Environment Variable at Execution.....	157
4.3.2.2 Environment Variable for OpenMP Specifications.....	160
4.3.2.3 Notes during Execution.....	160
4.3.3 Implementation-Dependent Specifications.....	161
4.3.4 Notes on OpenMP Programming.....	164
4.3.4.1 Implementation of parallel Region and Explicit task Region.....	164
4.3.4.2 Implementation of threadprivate Variable.....	164
4.3.4.3 Automatic Parallelization for OpenMP Programs.....	164
4.3.4.4 OpenMP Directive.....	165
4.3.4.5 Creating Shared Libraries in Clang Mode.....	165
4.4 Runtime Messages.....	165
<b>Chapter 5 Emitting Information.....</b>	<b>167</b>
5.1 Emitting Information at Compilation.....	167
5.1.1 Header.....	167
5.1.1.1 Output Format.....	167
5.1.2 Source List.....	167
5.1.2.1 Output Format.....	167
5.1.2.2 Information Included in Source List.....	168
5.1.2.2.1 Line Number of Source.....	168
5.1.2.2.2 Symbols for Parallelization.....	168
5.1.2.2.3 Symbols for Inline Expansion.....	169
5.1.2.2.4 Number of Loop Unrolling.....	169
5.1.2.2.5 Symbols for Using SIMD Extensions.....	169
5.1.2.2.6 Details Optimization Information.....	169
5.1.2.3 Example of Source Listing.....	172
5.1.2.4 Notes on Information at Compilation (Source List).....	174
5.1.3 Parallelization Messages.....	175
5.1.4 Statistics List.....	175
5.1.4.1 Output Format.....	176
5.1.4.2 Options Which Are Not Output.....	176
5.1.4.3 Options Which Are Interpreted by Compiler.....	176
5.2 Information at Execution.....	176



5.2.1 Runtime Messages.....	176
5.2.2 Trace Back Information.....	177
<b>Chapter 6 Language Specifications.....</b>	<b>178</b>
6.1 Language Specifications Extended in This System.....	178
6.1.1 long long Type.....	178
6.1.2 The pragma Directive.....	179
6.1.3 The ident Directive.....	179
6.1.4 The assert Directive.....	180
6.1.5 The unassert Directive.....	180
6.1.6 Predefined Macro Names.....	180
6.2 Support Status of Language Specifications.....	182
6.2.1 C++03 Specifications.....	183
6.2.2 C++11 Specifications.....	183
6.2.3 C++14 Specifications.....	185
6.2.4 C++17 Specifications.....	186
<b>Chapter 7 Notes on Linking with Different Languages and Trad/Clang Modes.....</b>	<b>187</b>
7.1 Compile Commands and Required Options when Linking.....	187
7.1.1 Linking C++ Trad Mode with C++ Clang Mode.....	191
7.1.2 Linking C++ Clang Mode with C++ Clang Mode.....	192
7.1.3 Linking MPI Object Programs and Use Profiler.....	192
7.1.4 Objects Generated in Clang Mode with -flto Option.....	192
7.2 Linking with C.....	192
7.3 Linking with Fortran.....	194
<b>Chapter 8 Debugging Functions.....</b>	<b>197</b>
8.1 Functions for Debugging.....	197
8.1.1 Subscript Range Checks (subchk Function).....	197
8.1.2 Heap Memory Checks (heapchk Function).....	197
8.1.2.1 Memory Release Check.....	198
8.1.2.2 Buffer Overrun Check.....	198
8.1.2.3 Memory Leak Check.....	198
8.2 Debugging Programs for Abend.....	198
8.2.1 Causes of Abend.....	198
8.2.2 Information Generated when an Abend Occurs.....	199
8.2.2.1 Information Generated for a General Abend.....	199
8.2.2.2 Information of SIGXCPU.....	200
8.2.2.3 Information when an Abend Occurs Again during Abend Processing.....	200
8.3 Hook Function.....	200
8.3.1 User-Defined Function Format.....	200
8.3.2 Notes on Hook Function.....	201
8.3.3 Calling the User-Defined Function from a Specified Location.....	201
8.3.3.1 Notes on Calling from a Specified Location.....	203
8.3.4 Calling the User-Defined Function at Regular Time Interval.....	203
8.3.4.1 Notes on Calling by Regular Time Interval.....	203
8.3.5 Calling from Any Location in the Program.....	203
<b>Chapter 9 Clang Mode.....</b>	<b>204</b>
9.1 From Compilation to Execution.....	204
9.1.1 Compile Command.....	204
9.1.1.1 Syntax of the Compile Command.....	204
9.1.1.2 Input Files for the Compile Command.....	204
9.1.1.3 Return Value of the Compile Command.....	205
9.1.1.4 Notes on Compilation.....	205
9.1.1.4.1 Stack Size on Compilation.....	205
9.1.2 Compiler Options.....	205
9.1.2.1 Syntax of the Compile Command.....	205

9.1.2.2 Description of Compiler Options.....	206
9.1.2.2.1 General Options for Compiler.....	206
9.1.2.2.2 Options for Messages.....	211
9.1.2.2.3 Options for Optimization.....	212
9.1.2.2.4 Options for Language Specifications.....	221
9.1.2.2.5 Options for CPU/Architecture.....	222
9.1.2.2.6 Options for Code Generation.....	223
9.1.2.3 Notes on Compile Options.....	224
9.1.2.3.1 Correspondences of Compile Options in Clang Mode and Trad Mode.....	224
9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation.....	228
9.1.2.3.3 Note on Using SVE.....	229
9.1.2.3.4 Notes on Using SIMD Built-in Functions.....	229
9.1.2.3.5 Notes on Specified SVE Vector Register Size.....	230
9.1.3 Environment Variable for Compile Command.....	230
9.1.4 Compilation Profile File.....	230
9.1.5 Predefined Macro Names.....	230
9.1.6 Procedure of Execution.....	232
9.1.6.1 Environment Variable for Execution.....	232
9.1.6.2 Notes for Execution.....	233
9.2 Optimization.....	233
9.2.1 Overview of Optimization.....	233
9.2.2 Using Optimization Functions.....	233
9.2.2.1 Using Optimization Control Line (Pragma Directives).....	233
9.2.2.1.1 Types of Optimization Control Lines.....	233
9.2.2.1.2 Optimization Control Specifier.....	234
9.2.2.1.3 Notes on Optimization Control Lines and Optimization Control Specifiers.....	248
9.2.3 SIMD.....	250
9.2.3.1 Math Functions that SIMD Extensions can be applied to.....	250
9.3 Multiprocessing.....	251
9.3.1 Overview of Multiprocessing.....	251
9.3.2 Parallelization by OpenMP Specification.....	251
9.4 Emitting Information.....	251
9.4.1 Emitting Information at Compilation.....	251
9.4.1.1 Header.....	252
9.4.1.2 Source List.....	252
9.4.1.2.1 Output Format.....	252
9.4.1.2.2 Information Included in Source List.....	253
9.4.1.2.3 Example of Source Listing.....	255
9.4.1.2.4 Notes on Information at Compilation (Source List).....	257
9.5 Language Specifications.....	258
9.5.1 Supported Language Standard.....	258
9.5.2 Half-Precision (16 bit) Floating-Point Type.....	259
9.5.3 C++ Standard Library Header <filesystem>.....	259
9.6 Notes on Linking with Different Languages and Trad/Clang Modes.....	259
9.6.1 Compile commands and required options when linking.....	260
9.6.2 Notes on Linking with the Object of Trad Mode.....	260
9.6.3 Linking with C.....	260
9.6.4 Linking with Fortran.....	262
9.7 SIMD Built-in Functions.....	263
9.8 Compatibility with GNU C++ Specifications.....	264
9.8.1 GNU C++ Extensions.....	264
9.8.2 GNU C++ Compatible Options.....	264
9.9 Code Coverage.....	266
9.9.1 How to Use Code Coverage.....	266
9.9.2 Necessary Files to Use Code Coverage.....	267
9.9.2.1 The ".gcno" File.....	267
9.9.2.2 The ".gcdc" File.....	268

9.9.3 Notes on Code Coverage.....	269
9.10 Restrictions and Notes.....	269
9.10.1 Restrictions of Macros <code>CMPLX</code> , <code>CMPLXF</code> , and <code>CMPLXL</code> in <code>complex.h</code> .....	269
Appendix A Precompiled Headers.....	271
A.1 Automatic Precompiled Header Processing.....	271
A.2 Manual Precompiled Header Processing.....	273
A.3 Other Ways to Control Precompiled Headers.....	273
Appendix B Restrictions and Notes.....	274
B.1 Restriction.....	274
B.1.1 OpenMP Program.....	274
B.1.2 C++ Standard Library.....	275
B.1.3 Restriction on Inheritance of a Constructor with Default Parameters Using "using" Inheritance between Inner Classes.....	275
B.2 Notes.....	275
B.2.1 Variable Arguments of Function.....	275
B.2.2 Result of the Signed Integer Type Expression.....	275
B.2.3 Undefined Behavior.....	275
B.2.4 Arrays of Variable Length.....	275
B.2.5 Restriction of Register Definition for Floating Point Types with <code>asm</code> Keyword.....	276
B.2.6 Restriction of Floating Point Types Expression as a Register Constraint for <code>asm</code> Operands.....	276
B.2.7 Restriction of Random Access Iterator Type.....	276
B.2.8 Debug Using Debugger.....	276
B.2.9 Changing SVE Vector Register Size.....	276
B.2.10 Integer Division Exception when the Divisor Is Zero.....	276
B.2.11 About Compile Time for Template Functions.....	277
B.2.12 Link Error (undefined reference to).....	277
B.2.13 Restriction of registers with <code>asm</code> keyword and on an inline <code>asm</code> clobber lists.....	277
Appendix C Compatibility with GNU C++ Specifications.....	278
C.1 GNU C++ Extensions.....	278
C.1.1 Attributes.....	279
C.1.2 Built-in Functions.....	280
C.2 GNU C++ Compatible Options.....	283
Appendix D Data and Memory Regions.....	288
D.1 Data Size and Alignment.....	288
D.2 Memory Regions.....	288
D.3 Data Allocation.....	288
D.4 Data Allocation to Stack Region.....	289
Appendix E Code Coverage.....	292
E.1 How to Use Code Coverage.....	292
E.2 Necessary Files to Use Code Coverage.....	293
E.2.1 The ".gcn" File.....	293
E.2.2 The ".gcn" File.....	294
E.3 Notes on Code Coverage.....	295
Appendix F Fujitsu Extended Functions.....	296
F.1 How to Use Fujitsu Extended Functions.....	296
F.1.1 Header.....	296
F.1.2 Supported Functions.....	296
Appendix G Runtime Information Output Function.....	297
G.1 Usage of Runtime Information Output Function.....	297
G.1.1 Range of Obtainable Information.....	297
G.1.2 Runtime Environment Variables.....	298
G.2 Output of Runtime Information Output Function.....	298
G.2.1 Output Information.....	298

G.2.2 Output Format.....	299
G.2.3 Output Example.....	299
G.3 Notes on Runtime Information Output Function.....	301
Appendix H Using High-Speed Facility.....	303
H.1 Inter-Core Hardware Barrier.....	303
H.1.1 Compilation.....	303
H.1.2 Execution.....	303
H.2 Sector Cache.....	305
Appendix I Multi-Process Execution.....	307
I.1 Specifying CPU Affinity.....	307
Appendix J Fujitsu OpenMP Library.....	309
J.1 Overview of Multiprocessing.....	309
J.2 Automatic Parallelization.....	309
J.2.1 Compilation (Automatic Parallelization).....	309
J.2.2 Execution Process (Automatic Parallelization).....	310
J.2.3 Example of Compilation and Execution.....	313
J.2.4 Performance Tuning.....	313
J.2.5 Automatic Parallelization.....	313
J.2.6 Optimization Control Line.....	313
J.2.7 Notes on Automatic Parallelization.....	313
J.2.8 Linking with Other Multi-Thread Programs.....	313
J.3 Parallelization by OpenMP Specification.....	315
J.3.1 Compilation (OpenMP Parallelization).....	315
J.3.2 Execution Process (OpenMP Parallelization).....	316
J.3.3 Implementation-Dependent Specifications.....	321
J.3.4 Notes on OpenMP Programming.....	323
J.3.5 Linking with Other Multi-Thread Programs.....	324
J.3.6 Debugging OpenMP Programs.....	325
J.4 Using High-Speed Facility.....	325
J.4.1 CPU Binding.....	326
J.4.2 Inter-Core Hardware Barrier.....	326
J.4.3 Sector Cache.....	328
J.5 Multi-Process Execution.....	328
J.5.1 Specifying CPU Affinity.....	328
Appendix K Large Page Library.....	331
K.1 Overview of the Memory Allocation Function.....	331
K.1.1 Large Page Function.....	331
K.1.2 Paging Methods.....	333
K.2 Large Pages of the FX System.....	334
K.2.1 FX System Memory Areas and Large Page Applicability.....	334
K.2.2 Memory Areas for Placing Application Program Data.....	336
K.3 Environment Variables for Large Page Library Settings.....	337
K.3.1 Basic Settings of the Large Page Library.....	337
K.3.2 Paging Method Setting.....	339
K.3.3 Settings for Tuning.....	340
K.4 Messages.....	348
K.4.1 How to Read Messages.....	348
K.4.2 Messages.....	348

# Chapter 1 Overview

This chapter describes an overview and an outline of how to process C++ programs using this system.

## 1.1 Configuration of the C++ Language Processing System

This system is one of language processing programs. The processing system consists of the following components:

### Compile Command

The compile command is a program which calls the C++ compiler, the preprocessor, the assembler, and the linker. The compile command converts a program written in the C++ programming language into an executable file format through the compilation and linking.

There are two types of compile commands in this system,

Table 1.1 Compile Command

Kind	Command name	Description
Cross compiler	FCCpx	The command used on the login node
Native compiler	FCC	The command used on the compute node

In this manual, the FCCpx is used as the name of the compile command. Read the FCCpx as the FCC when you use the native compiler.

### C++ Compiler

Compile command calls C++ compiler. C++ compiler creates object programs from C++ source programs.

C++ compiler has two modes with different user interfaces. The mode to be used is specified by the option of the compile command.

Table 1.2 Two Modes of C++ Compiler

Mode	Description
Trad Mode	<p>This mode uses an enhanced compiler based on compilers for K computer and PRIMEHPC FX100 or earlier system.</p> <p>This mode is suitable for maintaining compatibility with the past Fujitsu compiler. Supported specifications are C++03/C++11/C++14/Part of C++17 and OpenMP 3.1/Part of OpenMP 4.0/Part of OpenMP 4.5.</p> <p>For specifications of Trad Mode, see "<a href="#">Chapter 2 From Compilation to Execution</a>" to "<a href="#">Chapter 8 Debugging Functions</a>" and "<a href="#">Appendix</a>".</p>
Clang Mode	<p>This mode uses an enhanced compiler based on Clang/LLVM.</p> <p>This mode is suitable for compiling programs using the latest language specification and open source software. Supported specifications are C++03/C++11/C++14/C++17 and OpenMP 4.5/Part of OpenMP 5.0.</p> <p>For specifications of Clang Mode, see "<a href="#">Chapter 9 Clang Mode</a>".</p>



Trad Mode and Clang Mode may have different implementation-dependent specifications since these modes are based on the different compilers.

### C++ Library

This system provides the following libraries:

- C++ Standard Library

This system provides libc++.

- Library for Multiprocessing

This system provides two libraries for multiprocessing.

Table 1.3 Two Libraries for Multiprocessing

Library Name	Description
LLVM OpenMP Library	This is the library for multiprocessing based on LLVM OpenMP Runtime Library which is an open source software. Supported specifications are OpenMP 4.5/Part of OpenMP 5.0. This library is available in Trad Mode and Clang Mode.  For specifications of LLVM OpenMP Library, see " <a href="#">Chapter 4 Multiprocessing</a> ".
Fujitsu OpenMP Library	This is the library for multiprocessing based on Fujitsu OpenMP Library for K computer and PRIMEHPC FX100 or earlier system. This library is suitable for maintaining compatibility with the past Fujitsu OpenMP Library. Supported specifications are OpenMP 3.1/Part of OpenMP 4.0/Part of OpenMP 4.5. This library is only available in Trad Mode.  For specifications of Fujitsu OpenMP Library, see " <a href="#">Appendix J Fujitsu OpenMP Library</a> ".

#### Online Manual

Online manual pages, accessed by the man command, provide information about FCCpx(1), FCC(1), FCCpx\_trad\_mode(7), FCC\_trad\_mode(7), FCCpx\_clang\_mode(7), and FCC\_clang\_mode(7).

## 1.2 How to Use

---

This section gives a simple description of how to use this system.

### 1.2.1 Preparation

---

Append the following values to user environment variable before use. "*installation\_path*" is the product/software installation path. For "*installation\_path*", contact the system administrator.

Environment Variable	Value
PATH	<i>/installation_path/bin</i>
LD_LIBRARY_PATH	<i>/installation_path/lib64</i>
MANPATH	<i>/installation_path/man</i>

Example:

```
$ PATH=/installation_path/bin:$PATH
$ export PATH
$ LD_LIBRARY_PATH=/installation_path/lib64:$LD_LIBRARY_PATH
$ export LD_LIBRARY_PATH
$ MANPATH=/installation_path/man:$MANPATH
$ export MANPATH
```



#### Note

When setting the environment variable LANG with a locale, ensure that the locale is installed on the node where the program is compiled.

A list of available locales can be found with the "locale -a" command.

Example:

```
$ locale -a
C
:
```

If the locale is not installed, a warning message may be output or a compilation error may occur. In this case, take one of the following actions:

- Specify C for the environment variable LANG
- Ask the system administrator to install the locale

The current locale can be found with the "locale" command.

Example:

```
$ locale
LANG=C
:
```

---

## 1.2.2 Compilation and Linking

---

Use the compile command to compile source code written in the C++ language.

A wealth of functions for optimizing the object code and other uses can be used at compilation time. These functions can be used by specifying the compiler option in the operands of the compile command.

The compile command supports linking by calling linker (ld command).

Use the compile command to generate the executable program by linking the C++ language library with the object programs which are generated during compilation.

In this system, the C++ language specifications used at compilation time are based on the GNU C++ compiler specification (GNU C++ Extensions). Part of the GNU++ extensions to the C++ language and part of the GNU C++ compiler options can be used.

For GNU extensions to the C++ language, refer to the Section "[C.1 GNU C++ Extensions](#)" in Trad Mode, "[9.8.1 GNU C++ Extensions](#)" in Clang Mode.

For GNU C++ compiler options, refer to the Section "[C.2 GNU C++ Compatible Options](#)" in Trad Mode, the Section "[9.8.2 GNU C++ Compatible Options](#)" in Clang Mode.

Here is an example of how to compile and link the program `sample.cc` with the cross compiler.

The sample program: `sample.cc`

```
#include <iostream>
using namespace std;
int main(void)
{
    cout << "Hello, world." << endl;
    return 0;
}
```

Compilation and linking:

```
$ FCCpx sample.cc
```

### Use ISO Conforming C++ Library

When all of the following conditions exist, all libraries must be linked dynamically.

- An ISO conformant C++ library is used.
- A user defined dynamic library contains a constructor, and an ISO conformant C++ library is used in this constructor.
- A Class object including a constructor is defined in externally.

---

## 1.2.3 Debugging

---

For C++ programs debugging, use the source-level debugger. Use `gdb` to execute source-level debugger.

To use the debugger effectively, compile the C++ source code using the `-g` option for the compile command.

Note that the -g option invalidates the PCH options in Tard Mode.

## 1.2.4 Tuning

---

To tune C++ programs, use a sampling function provided by the Profiler.

Refer to the "Profiler User's Guide" for the Profiler.



# Chapter 2 From Compilation to Execution

This chapter describes the procedure for compiling C++ source code.

Note that this chapter is for Trad Mode. See "[Chapter 9 Clang Mode](#)" about Clang Mode.

## 2.1 Compile Command

Use the compile command to compile source code and create an executable program.

The compile command analyzes various operands specified in the compile command line and calls the preprocessor, compiler, assembler (as command), and linker (ld command) as necessary.

### 2.1.1 Syntax of the Compile Command

You can specify options for the preprocessor, compiler, assembler (as command), and linker (ld command). You can also specify filenames.

Options for the preprocessor and compiler are called compiler options.

Options for the compiler (compile command), assembler (as command), and linker (ld command) can be viewed with the man command.

Table 2.1 Format of the Compile Command

Command Name		Operands
Cross compiler	FCCpx	[ <i>_option-list</i> ] <i>_filename-list</i>
Native compiler	FCC	[ <i>_option-list</i> ] <i>_filename-list</i>

\_: At least one blank is required.

#### Information

There is no restriction on the order in which *option-list* and *filename-list* are specified.

*option-list* and *filename-list* can be mixed.

### 2.1.2 Input Files for the Compile Command

The file types that can be specified as input files for the compile command are as follows:

Table 2.2 Format of Input Files

File Type	File Suffix	Passed To
Headers	.h	Preprocessor <sup>[a]</sup>
	.H	
C++ source file	.cc	Preprocessor and compiler
	.c	
	.C	
	.cpp	
	.cxx	
	.i	
Assembler source file	.s	Assembler
Assembler source file that must be preprocessed	.S	Preprocessor and assembler
Object file	.o	Linker

[a] When -E option is effective.

## 2.1.3 Return Value of the Compile Command

The possible return values of the compile command are as follows:

Table 2.3 Return Values of the Compile Command

Return Value	Meaning
0	Terminated normally.
Non-zero	An error occurred at compiling or linking.

## 2.2 Compiler Options

This section describes the format of compiler options and their meanings.

Compiler options may be specified as follows:

- Operand for the compile command
- Environment variable (Refer to "2.3 Environment Variable for Compile Command")
- Compilation Profile File (Refer to "2.4 Compilation Profile File")

Refer to "2.4 Compilation Profile File" for the priority of compiler options.

Usually, options that are unrecognizable to the compile command are ignored and warning message is output. The behavior for unrecognized options can be changed by specifying the environment variable FCOMP\_UNRECOGNIZED\_OPTION. Refer to "2.3 Environment Variable for Compile Command" for the environment variable FCOMP\_UNRECOGNIZED\_OPTION.

### 2.2.1 Syntax of Compiler Options

Compiler options can be specified as operands of the compile command. The format is given below.

- General compiler options

```
[ -# ] [ -### ] [ -A- ] [ -Aname[(tokens)] ] [ -B{dynamic|static} ] [ -C ] [ -Dname[=tokens] ] [ -E ] [ -H ] [ -Idir ] [ -Ldir ] [ -M ] [ -MD ] [ -MF filename ] [ -MM ] [ -MMD ] [ -MP ] [ -MT target ] [ -Nopt ] [ -P ] [ -Qc ] [ -S ] [ -SSL2 ] [ -SSL2BLAMP ] [ -Uname ] [ -Wtool, arg1[, arg2]... ] [ -Yitem, dir ] [ -c ] [ {-g|-g0} ] [ -lname ] [ -mt ] [ -o pathname ] [ -shared ]
```

- Message options

```
[ -V ] [ {-help|--help} ] [ -j ] [ -w ] [ {--brief_diagnostics|--no_brief_diagnostics} ] [ --diag_error tag[, tag]... ] [ --diag_suppress tag[, tag]... ] [ --display_error_number ] [ --error_limit number ] [ {--for_init_diff_warning|--no_for_init_diff_warning} ]
```

- Optimization options

```
[ -Kopt ] [ -O[n] ] [ -x- ] [ -x={quick|noquick} ] [ -xfunc1[, func2]... ] [ -xn ]
```

- Language specifications options

```
[ -std=level ] [ {--alternative_tokens|--no_alternative_tokens} ] [ --dollar ] [ --library={a|C|m|n} ] [ --linkcoarray ] [ --linkfortran ] [ --preinclude filename ]
```

- PCH options

```
[ --create_pch fname ] [ --pch ] [ --pch_dir dir ] [ {--pch_messages|--no_pch_messages} ] [ --use_pch fname ]
```

## 2.2.2 Description of Compiler Options

---

A description of the compiler options is given below.

### 2.2.2.1 General Options for Compiler

The following options are general options for the compiler.

**-#**

Shows commands passed by the compile command, but does not execute.

**-###**

Shows commands passed by the compile command.

**-A-**

Invalidates all predefined macros (except those starting with "\_") and predefined assertions.

**-Aname[(tokens)]**

Uses *name* as a predicate for the specified *tokens* in the same way as for an `#assert` preprocessing directive.

The predefined assertions are given below:

- `system(unix)`

**-B{dynamic|static}**

Specify the kind of the library at linking. `-Bdynamic` is set by default.

These options are passed to the linker.

By using these options together with the `-lname` option, you can specify the kind of each library at linking.

**-Bdynamic**

Specifies to link with dynamic library or static library.

When the `-lname` option is specified after the `-Bdynamic` option, the linker will first search for a file of the name `libname.so`, and then search for a file of the name `libname.a`.

**-Bstatic**

Specifies to link with static library.

When the `-lname` option is specified after the `-Bstatic` option, the linker will only search for `libname.a`.



### Example

When linking `libxxx.a` and `libyyy.so`

```
$ FCCpx a.cc -Bstatic -lxxx -Bdynamic -lyyy
```



### Note

If the `-Bstatic` option is specified, the linker will search only a static library for not only the static library specified with the `-lname` option, but also standard library (`-lc` etc.) and necessary libraries by the feature. And, the linker error may occur because the necessary library is not found.

Specify the `-Bdynamic` option at the end of the command line when specifying the `-Bstatic` option.

Example 1: When the `-Bdynamic` option is not specified at the end

```
$ FCCpx a.cc -Bstatic -lxxx
ld: cannot find -lc
```

Example 2: When the `-Bdynamic` option is specified at the end

```
$ FCCpx a.cc -Bstatic -lxxx -Bdynamic
```

-C

This option executes the preprocessing phase to retain all comments other than those on preprocessing directive lines.

-D*name*[=*tokens*]

Assigns *name* to the specified *tokens* as in the `#define` preprocessor directive.

If the same *name* is specified simultaneously in the `-D` and `-U` options, the *name* is defined following the order of the options.

The predefined macros are shown in "[6.1.6 Predefined Macro Names](#)".

-E

Performs only preprocessing on the specified C++ source code and sends the results to the standard output.

The output includes the preprocessor directives used in the next step of the processing system.

-H

Writes the pathnames of the valid headers in the current compilation to the standard error.

-I*dir*

Sets the search path for headers with names that do not start with `'/'` (i.e. headers specified with an absolute path name). The directory specified in *dir* is searched first and then the directory-installed standard headers are searched.

If multiple directories are specified in multiple `-I` options, the directories are searched in the specified order.

The directory that retrieves headers is performed in the following order:

- The files enclosed in double quotation marks (`"`) are searched in the following order:
  1. The current directory that contains the `#include` preprocessor directive
  2. Directories specified with the `-I` option
  3. Directories specified as the environment variable `CPATH`
  4. Directories specified with the `-isystem` option
  5. Directories specified as the environment variable `CPLUS_INCLUDE_PATH`
  6. The Install directory that contains headers provided by the compiler
  7. The directory-installed standard headers
  8. Directories specified with the `-idirafter` option
- The files enclosed in angle brackets (`<>`) are searched in the following order:
  1. Directories specified with the `-I` option
  2. Directories specified as the environment variable `CPATH`
  3. Directories specified with the `-isystem` option
  4. Directories specified as the environment variable `CPLUS_INCLUDE_PATH`
  5. The Install directory that contains headers provided by the compiler
  6. The directory-installed standard headers
  7. Directories specified with the `-idirafter` option

If a header is specified with an absolute path name, only the specified absolute path name is searched.

If the specified directory does not exist, this option is invalidated.

-L*dir*

Adds *dir* to the list of directories through which the linker searches for libraries.

This option is passed to the linker.

Libraries are searched in the following order:

1. Directories specified as the argument of the -L option
2. The directory installed libraries provided by the compiler
3. The directory installed standard libraries
4. Directories specified to the environment variable `LIBRARY_PATH`

**-M**

Specifies to output the dependency of the sources in the format recognized by the `make` command.

When this option is valid, only the preprocessing is executed.

**-MD**

Specifies to output the dependency of the sources in the format recognized by the `make` command.

By default, the result is output to the file with a ".d" suffix.

When the -E and -o options are specified, the result is output to the file specified by the -o option.

When the -MF option is specified simultaneously, the result is output to the file specified by the -MF option.

**-MF *filename***

Specifies the *filename* to output the result of the dependency of the source using the -M, -MM, -MD or -MMD options.

**-MM**

Specifies to output the dependency of the sources in the format recognized by the `make` command. However, a standard header is not contained in the output.

When this option is valid, only the preprocessing is executed.

**-MMD**

Specifies to output the dependency of the sources in the format to recognized by the `make` command. However, a standard header is not contained in the output.

By default, the result is output to the file with a ".d" suffix.

When the -E and -o options are specified, the result is output to the file specified by the -o option.

When the -MF option is specified simultaneously, the result is output to the file specified by the -MF option.

**-MP**

Specifies to append the phony target of the dependency to the output of the result produced by the -M, -MM, -MD or -MMD option.

**-MT *target***

Specifies to change the *target* of the dependency in the output of the result produced by the -M, -MM, -MD or -MMD option.

**-Nopt**

For *opt*, specify one of the following:

```
{ {Rtrap|Rnotrap} | {cancel_overtime_compilation|nocancel_overtime_compilation} | check_cache_arraysize | {clang|noclang} | {coverage|nocoverage} | {exceptions|noexceptions} | {fjcx|nofjcx} | {fjprof|nofjprof} | fmtl=fmtl_mode | {hook_func|nohook_func} | {hook_time|nohook_time} | {libomp|fjomplib} | {line|noline} | lst[={pt}] | lst_out=file | profile_dir=dir_name | quickdbg[={subchk|nosubchk|heapchk|noheapchk|inf_detail|inf_simple}] | {reordered_variable_stack|noreordered_variable_stack} | {rt_tune|rt_notune} | rt_tune_func | rt_tune_loop[={all|innermost}] | {setvalue[=setarg]|nosetvalue} | src | sta }
```

Multiple *opts* can be specified to "-N" option using commas to separate them. For example,

```
-Nsrc,sta
```

can be specified instead of specifying

-Nsrc -Nsta

For details, see Section "2.2.2.7 -N Option".

-P

This option specifies not to include linemarkers in the output of the -E option.

-Qc

In *c*, specify y or n. If -Qy is specified, information identifying the various compile tools called is written to the output file. If -Qn is specified, this information is not written. By default, -Qy is set.

-S

Compiles the specified C++ source file and leaves the assembly-language output in a corresponding file with the suffix ".s".

This option is invalidated when the -E option is specified simultaneously.

-SSL2

Specifies to link C-SSL II, C-SSL II thread-parallel capabilities, and BLAS/LAPACK.

-SSL2BLAMP

Specifies to link C-SSL II, C-SSL II thread-parallel capabilities, and BLAS/LAPACK thread-parallel versions.

This option just replaces the sequential BLAS/LAPACK from -SSL2 with the corresponding thread-parallel versions.

-Uname

Invalidates the definition of *name* as specified in #undef preprocessor directive.

If the same *name* is specified simultaneously in the -D and -U options, the *name* is invalidated following the order of the options.

-Wtool, arg1[, arg2]...

Directs that the specified arguments *arg1*[, *arg2*]... are to be passed as separate arguments to *tool*. An individual argument must be separated by a comma only.

*tool* can be one of the following characters:

p	Preprocessor
0	Compiler
a	Assembler
l	Linker

-Yitem, dir

Specifies the new directory *dir* as the position for *item*.

For *item*, specify one of the following characters:

0	Changes the path name of the compiler to <i>dir</i> /ccpcomp (cross compiler). Changes the path name of the compiler to <i>dir</i> /ccpcom (native compiler).
a	Changes the path name of the assembler to <i>dir</i> /as.
l (the lowercase letter L)	Changes the path name of the linker to <i>dir</i> /ld.
M	Changes the directory of the message files to <i>dir</i> .

If more than one -Y options are specified for an *item*, the last specified option is valid.

-c

Suppresses the linking phase. The object files generated up to then are not deleted.

This option is invalidated when the -E or -S option is specified simultaneously.

{-g|-g0}

The -g option generates additional information used by the debugger.

The -g0 option invalidates the -g option.

-g0 is set by default.

The -g option invalidates the PCH options.

Refer to "[B.2.8 Debug Using Debugger](#)" about the notes on the debugger in Trad Mode.

-lname

Specifies to link a library named *libname.so* or *libname.a*.

The position of this option in the command line is important, since libraries are searched for in the order in which the other libraries and object files appear in the command line. This option must be specified after source files.

This option is passed to the linker.

-mt

Creates a thread safe object.

The compiler option -mt is needed if an object program compiled with it exists in the command line as an input file.

This option is also needed if the program is multi-threaded.

-o *pathname*

Creates a file of the name specified in *pathname*.

- If the -c option is set, an object file is created with the name *pathname*.
- If the -shared option is set, a shared object is created with the name *pathname*.
- If the -S option is set, an assembler source file is created with the name *pathname*.
- In other cases, an executable program is created with the name *pathname*, instead of the default *a.out*.

-shared

Specifies to create a shared object. This option is passed to the linker.

## 2.2.2.2 Options for Messages

The options for messages are given below.

-V

Specifies to emit the version and release information of each command of the C++ compiler to the standard error.

{-help|--help}

Specifies to output help information.

-j

Suppresses warnings on local automatic variables that are used before their values are set.

-w

Suppresses the output of warning messages.

{--brief\_diagnostics|--no\_brief\_diagnostics}

Enable or disable 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 it is too long to fit on a single line. By default, --no\_brief\_diagnostics is set.

--diag\_error *tag*[, *tag*]...

Directs the compiler to treat warning message(s) specified by *tag* as error message(s). The *tag* can be specified by an error number.

`--diag_suppress tag[,tag]...`

Suppresses the warning message(s) specified by *tag*. The *tag* can be specified by an error number. The error message cannot be suppressed.

`--display_error_number`

Display the error message number in any diagnostic messages that are generated. The option may be used to determine the error number to be used when specifying `--diag_error` or `--diag_suppress`.

`--error_limit number`

Set the limit on the *number* of errors. The compiler will abandon compilation after detecting *number* of errors (remarks and warnings are not counted toward the limit). By default, the limit is 100.

`{--for_init_diff_warning|--no_for_init_diff_warning}`

Enable or disable a warning issued when programs compiled under the new for-init scoping rules would have had different behavior under the old rules. The diagnostic is only output when the new rules are used. By default, `--for_init_diff_warning` is set.

### 2.2.2.3 Options for Optimization

The optimization options are described below.

`-Kopt`

For *opt*, specify one of the following:

```
{ {PIC|pic} | {SVE|NOSVE} | {alias_const|noalias_const} | {align_loops[=N]|noalign_loops} | archi | {array_declaration_opt|noarray_declaration_opt} | {array_private|noarray_private} | assume={shortloop|noshortloop|memory_bandwidth|nomemory_bandwidth|time_saving_compilation|notime_saving_compilation} | cmodel={small|large} | cpu | {dynamic_iteration|nodynamic_iteration} | {eval|noeval} | {eval_concurrent|eval_noconcurrent} | {extract_stride_store|noextract_stride_store} | fast | {fast_matmul|nofast_matmul} | {fenv_access|nofenv_access} | {fp_contract|nofp_contract} | {fp_precision|nofp_precision} | {fp_relaxed|nofp_relaxed} | {fsimple|nofsimple} | {fz|nofz} | {hpctag|nohpctag} | {ifunc[={loop|procedure}]|noifunc} | instance=N | {largepage|nolargepage} | {lib|nolib} | {loop_blocking[=N]|loop_noblocking} | {loop_fission|loop_nofission} | {loop_fission_stripmining[={N|L1|L2}]|loop_nofission_stripmining} | loop_fission_threshold=N | {loop_fusion|loop_nofusion} | {loop_interchange|loop_nointerchange} | {loop_part_parallel|loop_nopart_parallel} | {loop_part_simd|loop_nopart_simd} | {loop_perfect_nest|loop_noperfect_nest} | {loop_versioning|loop_noversioning} | looptype={f|n|s} | {memalias|nomemalias} | {mfunc[={1|2|3}]|nomfunc} | noprefetch | {ocl|noocl} | {omitfp|noomitfp} | {openmp|noopenmp} | {openmp_assume_norecurrence|openmp_noassume_norecurrence} | {openmp_collapse_except_innermost|openmp_nocollapse_except_innermost} | openmp_loop_variable={private|standard} | {openmp_ordered_reduction|openmp_noordered_reduction} | {openmp_simd|noopenmp_simd} | {optlib_string|nooptlib_string} | {optmsg[={1|2}]|nooptmsg} | {parallel|noparallel} | {parallel_fp_precision|parallel_nofp_precision} | parallel_iteration=N | parallel_strong | {pc_relative_literal_loads|nopc_relative_literal_loads} | {plt|noplt} | {preex|nopreex} | prefetch_cache_level={1|2|all} | {prefetch_conditional|prefetch_noconditional} | {prefetch_indirect|prefetch_noindirect} | {prefetch_infer|prefetch_noinfer} | prefetch_iteration=N | prefetch_iteration_L2=N | prefetch_line=N | prefetch_line_L2=N | {prefetch_sequential[={auto|soft}]|prefetch_nosequential} | {prefetch_stride[={soft|hard_auto|hard_always}]|prefetch_nostride} | {prefetch_strong|prefetch_nostrong} | {prefetch_strong_L2|prefetch_nostrong_L2} | {preload|nopreload} | {rdconv[={1|2}]|nordconv} | {reduction|noreduction} | {region_extension|noregion_extension} | {remove_inlinefunction|noremove_inlinefunction} | {restp[={all|arg}]|norestp} | {sch_post_ra|nosch_post_ra} | {sch_pre_ra|nosch_pre_ra} | {sibling_calls|nosibling_calls} | {simd[={1|2|auto}]|nosimd} | {simd_packed_promotion|simd_nopacked_promotion} | {simd_reduction_product|simd_noreduction_product} | simd_reg_size={128|256|512|agnostic} | {simd_uncounted_loop|simd_nouncounted_loop} | {simd_use_multiple_structures|simd_nouse_multiple_structures} | strict_aliasing|nostrict_aliasing} | {striping[=N]|nostriping} | {swp|noswp} | {swp_freq_rate=N|swp_ireg_rate=N|swp_preg_rate=N} | swp_policy={auto|small|large} | swp_strong | swp_weak | tls_size={12|24|32|48} | {unroll[=N]|nounroll} | {unroll_and_jam[=N]|nounroll_and_jam} | visimpact | {zfill[=N]|nozfill} }
```

Multiple *opts* can be specified to "-K" option using commas to separate them. For example,

```
-Kfast,parallel
```

can be specified instead of specifying

```
-Kfast -Kparallel
```



For more details, see Section ["2.2.2.6 -K Option"](#).

## **-O[n]**

In *n*, specify the level of optimization as 0, 1, 2 or 3. If the -O option is specified without *n*, -O2 is set. As the optimization level increases, the execution time decreases and the compilation time increases. The higher levels of optimization functionally include the lower levels of optimization.

If the -O[n] option is not specified, -O0 is set.

This option is effective to C++ source files only. For information about C++ source files, see Section ["2.1.2 Input Files for the Compile Command"](#).

### **- Optimization level 0**

No optimization is performed.

### **- Optimization level 1**

Basic optimization is performed. Moreover, the -Knoalias\_const option is set. This may reduce the object size and the execution time compared with optimization level 0.

### **- Optimization level 2**

In addition to the optimization of level 1, the following optimizations are performed:

- Loop unrolling (Equivalent to specifying -Kunroll)
- Software pipelining (Equivalent to specifying -Kswp)
- Loop blocking (Equivalent to specifying -Kloop\_blocking)
- Loop fusion (Equivalent to specifying -Kloop\_fusion)
- Loop fission (Equivalent to specifying -Kloop\_fission)
- Loop interchange (Equivalent to specifying -Kloop\_interchange)
- Using prefetch instructions (Equivalent to specifying -Kprefetch\_sequential,prefetch\_cache\_level=all)
- SIMD extend instruction conversion (Equivalent to specifying -Ksimd=auto)
- Adjusting loop alignment (Equivalent to specifying -Kalign\_loops)
- Optimization of the sibling call (Equivalent to specifying -Ksibling\_calls)
- Repeated application of optimization functions  
Repeated application of optimization functions means that the optimization functions performed in optimization level 1 are repeatedly performed until there is no room for further optimization.
- Optimization is performed on the assumption that expressions of signed integer type whose size is 4-byte or less do not overflow when the expressions increase or decrease by a constant value as the loop repeats (Equivalent to specifying -Krdconv=1)

This may result in an increased object program size.

This level is specified to pursue the execution performance improvement of the object program.

### **- Optimization level 3**

In addition to the optimization of level 2, the following optimizations are performed:

- Unrolling of nested loops
- Transforming loops into perfectly nested loops to promote optimization such as loop interchange (Equivalent to specifying -Kloop\_perfect\_nest)
- Loop unswitching
- Clone optimization
- Inline expansion (Equivalent to specifying -x-)

This level is specified to pursue the execution performance improvement of the object program.

-x-

Inline expansion instead of function calls is performed for all defined functions in which the number of initialized declarations and executable statements after macro expansion is not greater than that allowed by the compiler ( $n$ ).

This option is effective only if the `-O1` option or higher is set.

If both this option and the `-xn` option are specified, the one specified last is effective.

This option determines whether to perform inline expansion by using the value of  $n$ . Therefore, the inline function does not become the target of the inline expansion when the compiler judges that inline expansion should not be performed.

See Section "3.2.9 Inline Expansion" for about inline expansion.

-x={quick|noquick}

`-x=quick` option applies optimizations which do not enlarge the object size. For example, if the source program contains many functions, the functions which have less executive instructions are inline expanded with priority. The source program may be compiled more quickly using less memories.

The `-x=noquick` option invalidates the `-x=quick` option. `-x=quick` option is set by default.

The `-x=quick` option and `-x=noquick` option require that the `-x-` option is set.

-xfunc1[,func2]...

Inline expansion of function calls is performed for the functions `func1[,func2]...` defined in the C++ source code.

This option is effective only if the `-O1` option or higher is set.

-xn

Inline expansion is performed for functions in which the number of initialized declarations and executed statements is equal to or less than the specified number  $n$ .

$n$  can be a number from 0 to 2147483647.

If both this option and the `-x-` option are specified, the one specified last is effective.

When the `-x0` option is specified, the `-x-`, `-xn`, and `-xfunc1[,func2]...` options are ignored, so inline expansion is invalidated. However, the one specified last is effective.

This option is effective only if the `-O1` option or higher is set.

The inline function becomes the target of the inline expansion regardless of the value of  $n$ .

## 2.2.2.4 Options for Language Specifications

This section explains the options relating to language specifications.

-std=*level*

This option specifies the level of language specifications to be interpreted by the compiler (including the preprocessor). For *level*, specify `c++98`, `c++03`, `c++11`, `c++14`, `c++17`, `gnu++98`, `gnu++03`, `gnu++11`, `gnu++14`, or `gnu++17`. `-std=gnu++14` is set by default.

`-std=c++98`

`-std=c++03`

`-std=gnu++98`

`-std=gnu++03`

In addition to the C++03 specifications, GNU C++ Extensions is used at compilation time.

`-std=c++11`

`-std=gnu++11`

In addition to the C++11 specifications, GNU C++ Extensions is used at compilation time.

`-std=c++14`

`-std=gnu++14`

In addition to the C++14 specifications, GNU C++ Extensions is used at compilation time.

`-std=c++17`  
`-std=gnu++17`

In addition to the C++17 specifications, GNU C++ Extensions is used at compilation time.

Some of the predefined macros are changed by the `-std=level` option. For predefined macro names, see "[6.1.6 Predefined Macro Names](#)".

`{--alternative_tokens|--no_alternative_tokens}`

This option specifies whether to recognize digraph tokens and an operator keyword. By default, `--no_alternative_tokens` is set.

The following digraph tokens are available:

<code>&lt;%</code>	Two character notation of "{ "
<code>%&gt;</code>	Two character notation of "} "
<code>&lt;:</code>	Two character notation of "[ "
<code>:&gt;</code>	Two character notation of "] "
<code>#: </code>	Two character notation of "# "
<code>%%: </code>	Two character notation of "## "

The following operator keywords are available:

<code>and</code>	Operator keyword of "&&"
<code>and_eq</code>	Operator keyword of "&="
<code>bitand</code>	Operator keyword of "&"
<code>bitor</code>	Operator keyword of "  "
<code>compl</code>	Operator keyword of "~"
<code>not</code>	Operator keyword of "!"
<code>not_eq</code>	Operator keyword of "!="
<code>or</code>	Operator keyword of "   "
<code>or_eq</code>	Operator keyword of " ="
<code>xor</code>	Operator keyword of "^"
<code>xor_eq</code>	Operator keyword of "^="

`--dollar`

Accept \$ signs in identifiers.

`--library={a|C|m|n}`

Specifies library use. If a library is specified with the `--library` option, the proper options are passed to the compiler and linker in order to use (or not use) a specified library. By default, `--library=a` is set.

<code>a</code>	Use all libraries, in the order libstd, libm
<code>C</code>	Do not use libstd
<code>m</code>	Do not use libm
<code>n</code>	Do not use C++ libraries

`--linkcoarray`

Directs to search for a library required for linkage with Fortran that does use the COARRAY specification.

If this option is not specified, search processing is not performed, enabling a reduction in compilation time.

--linkfortran

Directs to search for a library required for linkage with Fortran that does not use the COARRAY specification.

If this option is not specified, search processing is not performed, enabling a reduction in compilation time.

--preinclude *filename*

Directs to include the source code of the indicated file at the beginning of the compilation. This can be used to establish standard macro definitions. *filename* is searched for in the directories on header search list.

## 2.2.2.5 Options for PCH

This section explains the options specified when a PCH (Pre Compiled Headers) function is used.

--create\_pch *fname*

Create a precompiled header file with the specified *fname*. If the --pch or --use\_pch option appears on the command line following this option, this option is invalidated.

--pch

Automatically use and/or create a precompiled header file. If the --use\_pch or --create\_pch option appears on the command line following this option, this option is invalidated.

--pch\_dir *dir*

The directory in which to search for and/or create a precompiled header file. This option is effective only if the --pch, --create\_pch or --use\_pch option is also specified.

{--pch\_messages|--no\_pch\_messages}

Enable or disable the display of a message indicating that a precompiled header file was created or used in the current compilation. By default, --pch\_messages is set.

--use\_pch *fname*

Use a precompiled header file of the specified *fname* as part of the current compilation. If the --pch or --create\_pch option appears on the command line following this option, this option is invalidated.

## 2.2.2.6 -K Option

-K{PIC|pic}

Specifies to generate position-independent code (PIC). The -KPIC and -Kpic options are equivalent.

Note that the -K{PIC|pic} and -Kcmodel=large options cannot be set simultaneously.

This option takes effect only during compilation.

-K{SVE|NOSVE}

-KSVE specifies to output object files using SVE, which is an Armv8-A extension. -KSVE is set by default.

Specify the -KNOSVE option to generate object files for a processor that does not support SVE.

-K{alias\_const|noalias\_const}

-Knoalias\_const specifies data optimization pointed to by const pointers. This assumes that the const pointers are not defined by other pointers. The -Kalias\_const option invalidates the -Knoalias\_const option. -Knoalias\_const is set by default.

The -Kalias\_const and -Knoalias\_const options require that the -O1 option or higher is set.

-K{align\_loops[=*N*]|noalign\_loops}

-Kalign\_loops specifies to align loops to a power of two byte boundary. The -Knoalign\_loops option invalidates the -Kalign\_loops option. *N* is the number of bytes for the alignment boundary. The value that can be specified for *N* is a power of two (from 1 to 32768) or 0. If a value is not specified or 0 is specified for *N*, the compiler will automatically determine a value. The -Knoalign\_loops is set by default when the -O1 option is set. The -Kalign\_loops option is set by default when the -O2 option or higher is set.

Specifying -Knoalign\_loops is equivalent to specifying -Kalign\_loops=1.

The -Kalign\_loops and -Knoalign\_loops options require that the -O1 option or higher is set.

## -Karchi

Specifies the architecture. In *archi*, specify one of ARMV8\_A, ARMV8\_1\_A, ARMV8\_2\_A, or ARMV8\_3\_A. -KARMV8\_3\_A is set by default.

### -KARMV8\_A

This option specifies to generate object files using instructions in Armv8-A.

### -KARMV8\_1\_A

This option specifies to generate object files using instructions in Armv8-A and Armv8.1-A.

### -KARMV8\_2\_A

This option specifies to generate object files using instructions in Armv8-A, Armv8.1-A, and Armv8.2-A.

### -KARMV8\_3\_A

This option specifies to generate object files using instructions in Armv8-A, Armv8.1-A, Armv8.2-A, and Armv8.3-A.

## -K{array\_declaration\_opt|noarray\_declaration\_opt}

-Karray\_declaration\_opt specifies to perform the optimization such as SIMDization assuming that the array subscript does not exceed the range of the array declaration. The -Knoarray\_declaration\_opt option invalidates the -Karray\_declaration\_opt option.

The -Knoarray\_declaration\_opt is set by default.

The result of the execution cannot be guaranteed if the -Karray\_declaration\_opt option is set when the above assumption is not based.

The -Karray\_declaration\_opt and -Knoarray\_declaration\_opt options require that the -O1 option or higher is set.

## -K{array\_private|noarray\_private}

-Karray\_private specifies to localize arrays in a loop. The -Knoarray\_private option invalidates the -Karray\_private option. -Knoarray\_private is set by default.

The -Karray\_private and -Knoarray\_private options require that the -Kparallel option is set.

## -Kassume={shortloop|noshortloop|memory\_bandwidth|nomemory\_bandwidth|time\_saving\_compilation|notime\_saving\_compilation}

These options specify whether or not to control optimization considering features of the program. Multiple -Kassume options can be specified at the same time.

-Kassume=noshortloop,assume=nomemory\_bandwidth,assume=notime\_saving\_compilation is set by default. The compiler performs optimizations under the following policy.

- Assume that the iteration count is large when the iteration count of the innermost loop is unknown at compilation.
- Solve the bottleneck of CPU operation preferentially.
- Give priority to increase the speed of the executable program rather than decrease of the compilation time.

These options require that the -O1 option or higher is set.

### -Kassume={shortloop|noshortloop}

-Kassume=shortloop specifies to assume that iteration counts are small when the iteration count of the innermost loop in the program is unknown at compilation. -Kassume=noshortloop specifies to assume that iteration counts are not small. -Kassume=noshortloop is set by default.

When -Kassume=shortloop option is specified, optimizations such as automatic parallelization, loop unrolling, and software pipelining may be controlled or invalidated.

### -Kassume={memory\_bandwidth|nomemory\_bandwidth}

-Kassume=memory\_bandwidth specifies to assume that innermost loops in the program has a memory bandwidth bottleneck. -Kassume=nomemory\_bandwidth specifies to assume that the innermost loop in the program does not have a memory bandwidth bottleneck. -Kassume=nomemory\_bandwidth is set by default.

When -Kassume=memory\_bandwidth option is specified, optimizations such as zfill optimization promotion and software pipelining may be controlled or invalidated.

`-Kassume={time_saving_compilation|notime_saving_compilation}`

`-Kassume=time_saving_compilation` specifies to control optimization for decreasing compilation time of the program. `-Kassume=notime_saving_compilation` specifies to optimize with giving priority to the speed of the executable program rather than decrease of the compilation time. `-Kassume=notime_saving_compilation` is set by default.

`-Kcmodel={small|large}`

This option specifies the possible largest size of the text area and the static data area in an executable program or a shared object. `-Kcmodel=small` is set by default.

Note that the `-K{PIC|pic}` and `-Kcmodel=large` options cannot be set simultaneously.

`-Kcmodel=small`

The total size of the text area and the static data area is limited to 4GB at linking. This option creates an efficient object program.

`-Kcmodel=large`

Only the size of the text area is limited to 4GB at linking. This option is used when the static data area is large and an error occurs at linking.

`-Kcpu`

Specifies the target processor. In *cpu*, specify A64FX or GENERIC\_CPU. `-KA64FX` is set by default.

`-KA64FX`

`-KA64FX` specifies to output object files for the A64FX processor. The `-Khpctag` option is effective when the `-KA64FX` option is set.

`-KGENERIC_CPU`

`-KGENERIC_CPU` specifies to output the object file for the Arm processor.

`-K{dynamic_iteration|nodynamic_iteration}`

In a multiloop composed of parallel executable loops, the `-Kdynamic_iteration` option specifies to select the loop dynamically (at execution) depending on the execution environment. The nest level of the loop that can be dynamically selected is up to three from the outside of the multiloop. The `-Knodynamic_iteration` option specifies to select the outermost loop of the multiloop to execute in parallel. `-Knodynamic_iteration` is set by default.

The `-Kdynamic_iteration` and `-Knodynamic_iteration` options require that the `-Kparallel` option is set.

`-K{eval|noeval}`

`-Keval` performs the optimization that changes the method of operator evaluation. The `-Knoeval` option invalidates the `-Keval` option. `-Knoeval` is set by default.

If the `-Keval` option is set, side effects (calculation errors or runtime exceptions) may occur in the execution results and produce unexpected results. See Section "[3.6.1 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

If the `-Keval` option is set, the `-Kfsimple` option also takes effect. Moreover, the `-Kreduction` option also takes effect if `-Kparallel` is set.

If `-Keval` and `-Ksimd[={1|2|auto}]` options are set, the `-Kfsimple` and `-Ksimd_reduction_product` options also takes effect.

If the `-Knoeval` option is set, the `-Knofsimple`, `-Knoreduction`, and `-Ksimd_noreduction_product` options are set.

The `-Keval` and `-Knoeval` options require that the `-O1` option or higher is set.

Before using the `-Keval` option, see Section "[3.3.1.2 Arithmetic Evaluation Method Modification](#)".

`-K{eval_concurrent|eval_noconcurrent}`

`-Keval_concurrent` option specifies to give priority to the instruction-level parallelism in the tree-height-reduction optimization. `-Keval_noconcurrent` option specifies to suppress the instruction-level parallelism and specifies to give priority to utilizing FMA instructions in the tree-height-reduction optimization. `-Keval_noconcurrent` option is set by default.

The `-Keval_concurrent` and `-Keval_noconcurrent` options require that the `-O1` option or higher is effective and the `-Keval` option is effective.

The effect of the `-Keval_concurrent` option is expected when the loop iteration count is small and therefore the software pipelining is not applied.

See section [3.3.9 Tree-Height-Reduction Optimization](#) for the tree-height-reduction optimization.

#### `-K{extract_stride_store|noextract_stride_store}`

`-Kextract_stride_store` option specifies to use scalar instructions for stride access store in the target loop when using SIMD extensions. The `-Knoextract_stride_store` option specifies to use SIMD instructions for stride access store in the target loop when using SIMD extensions. `-Kextract_stride_store` option is set by default.

When the `-Kextract_stride_store` option is specified, the function of "[3.2.7.4 SIMD with Redundant Executions for the SIMD Width](#)" is suppressed for loops to which this function is applied.

Note that the size of the object program and compile time may increase if the `-Kextract_stride_store` option is specified.

The `-Kextract_stride_store` and `-Knoextract_stride_store` options require that the `-O2` option or higher option is set.

#### `-Kfast`

Specifies the optimizations for high speed execution on the target machine.

This option is equivalent to the following:

`-O3 -Keval,fast_matmul,fp_contract,fp_relaxed,fz,ilfunc,mfunc,omitfp,simd_packed_promotion`

This option must be set at both compilation and linking.

When this option is set, side effects may occur in the execution results due to effect of the `-Keval`, `-Kfast_matmul`, `-Kfp_contract`, `-Kfp_relaxed`, `-Kilfunc`, and `-Kmfunc` options. See Section "[3.6.1 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

#### `-K{fast_matmul|nofast_matmul}`

The compiler option `-Kfast_matmul` performs the optimization that converts the loop of matrix multiplication into a high speed library call. The compiler option `-Knofast_matmul` invalidates the compiler option `-Kfast_matmul`. The compiler option `-Knofast_matmul` is set by default.

When the compiler option `-Kfast_matmul` is set, side effects (calculation errors) may occur in the execution results. See Section "[3.6.1 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

The compiler options `-Kfast_matmul` and `-Knofast_matmul` require that the compiler option `-O2` or higher option is set.

The compiler option `-Kfast_matmul` is needed if an object program compiled with it exists in the command line as an input file.

#### `-K{fenv_access|nofenv_access}`

`-Kfenv_access` specifies that a program may access the floating-point environment to test flags or run under non-default modes. The `-Knofenv_access` option invalidates the `-Kfenv_access` option. `-Knofenv_access` is set by default.

The `-Kfenv_access` option disables some floating-point optimization.

#### `-K{fp_contract|nofp_contract}`

`-Kfp_contract` specifies to perform optimization using Floating-Point Multiply-Add/Subtract instructions. The `-Knofp_contract` option invalidates the `-Kfp_contract` option. `-Knofp_contract` is set by default.

When the `-Kfp_contract` option is set, side effects (calculation errors in rounding error extent) may occur in the execution results and produce unexpected results. See Section "[3.6.1 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

The `-Kfp_contract` and `-Knofp_contract` options require that the `-O1` option or higher is set.

#### `-K{fp_precision|nofp_precision}`

`-Kfp_precision` specifies to induce the combination of optimization options that do not cause calculation errors in floating-point operations. The `-Knofp_precision` option invalidates the `-Kfp_precision` option. `-Knofp_precision` is set by default.

The `-Kfp_precision` option is equivalent to replacing the `-Kfp_precision` option with the following:

`-Knoeval,nofast_matmul,nofp_contract,nofp_relaxed,nofz,noilfunc,nomfunc,parallel_fp_precision`

If the `-Kfp_precision` option is effective, some optimizations are limited and the execution performance may decrease.

The `-Knofp_precision` option does not affect the individual options induced by `-Kfp_precision` option, even if they are specified together.

### -K{fp\_relaxed|nofp\_relaxed}

-Kfp\_relaxed specifies to convert floating point division operations and `sqrt` functions into the reciprocal approximation operation with the reciprocal approximation instructions and the Floating-Point Multiply-Add/Subtract instructions. This conversion is applied to the single-precision and double-precision real type calculation. The -Knofp\_relaxed option invalidates the -Kfp\_relaxed option. -Knofp\_relaxed is set by default.

When the -Kfp\_relaxed option is set, side effects may occur in the execution results. Also, this optimization may result in generating a floating-point exception whether the -NRtrap option is valid or not. See Section "3.6.1 Side Effect of Optimizations for Floating-Point Operation" for the side effect of optimization.

The -Kfp\_relaxed and -Knofp\_relaxed options require that the -O1 option or higher is set.

When the -NRtrap option and either of the -Knosimd option or the -KNOSVE option are effective, the optimization to convert a `sqrt` function into reciprocal approximation instructions is suppressed even though the -Kfp\_relaxed option is set. Therefore, the execution performance may decrease as compared with when the -NRnotrap option is effective. Note that the reciprocal approximation instructions generated by the -Kfp\_relaxed option may be evaluated in advance even if the -Knopreex option is enabled. And a floating-point exception may occur when the -Knopreex, -Kfp\_relaxed, and -NRtrap options are valid.

### -K{fsimple|nofsimple}

-Kfsimple specifies to simplify floating-point operation. For example,  $x*0$  operation is simplified to 0. The -Knofsimple option invalidates the -Kfsimple option. -Knofsimple is set by default.

When -Kfsimple is specified, side effects may occur in the execution results. See Section "3.6.1 Side Effect of Optimizations for Floating-Point Operation" for the side effect of optimization.

The -Kfsimple and -Knofsimple options require that the -O1 option or higher is set.

### -K{fz|nofz}

-Kfz option specifies to use flush-to-zero mode. The -Knofz option specifies not to use flush-to-zero mode. The default is -Knofz.

When -Kfz option is specified, flush-to-zero mode replaces a denormalized number with zero with the same sign if a result or a source operand is a denormalized number. Moreover, side effects may occur in the execution results. See Section "3.6.1 Side Effect of Optimizations for Floating-Point Operation" for the side effect of optimization. Some programs may get more performance.

-Kfz and -Knofz options must be set at linking.

-Kfz is effective only when the -O1 option or higher is set.

### -K{hpctag|nohpctag}

-Khpctag specifies to use the HPC tag address override function of the A64FX processor. The -Knohpctag option invalidates the -Khpctag option. -Khpctag is set by default.

Using the HPC tag address override function, the Sector cache and the hardware prefetch assistance function become effective.

The -Khpctag and -Knohpctag options require that the -KA64FX option is set.

The -Khpctag and -Knohpctag options should be specified when programs are compiled and linked.

### -K{ilfunc[={loop|procedure}]}|noilfunc}

-Kilfunc performs the inline expansion for the math functions. The -Knoilfunc option invalidates the -Kilfunc option.

The following shows the functions for which inline expansion can be performed.

- Single-precision and double-precision real type function: `atan`, `atan2`, `cos`, `exp`, `log`, `log10`, `pow`, `sin`, and `tan`
- Single-precision and double-precision complex type function: `abs` and `exp`

The -Klib option is required for the -Kilfunc option to take effect.

If the argument `={loop|procedure}` is omitted, `-Kilfunc=procedure` is set.

-Knoilfunc is set by default.

If the compiler determines that the optimization is not promoted due to such as data type and the presence of function calls, inline expansion is not applied.



Inline expansion is applied by the `-Kilfunc` option, side effects may occur in the execution result. Also, this optimization may result in generating a floating-point exception when `-NRtrap` is not valid. See Section "3.6.1 Side Effect of Optimizations for Floating-Point Operation" for the side effect of optimization.

The `-Kilfunc` and `-Knoifunc` options require that the `-O1` option or higher is set.

`-Kilfunc=loop`

Specifies to perform the inline expansion for the math functions in a loop.

`-Kilfunc=procedure`

Specifies to perform the inline expansion for the math functions in a function.

`-Kinstance=N`

Specifies the number of threads for parallel execution as  $N$ .  $N$  should be from 2 to 512.

This option requires that the `-Kparallel` option is set.

When the number of threads at run-time does not equal the value of  $N$ , the following message is output and execution is discontinued.

```
jwe1040i-s This program cannot be executed, because the number of threads specified by -Kinstance=N
option and the actual number of threads are not equal.
```

Before using this option, see Section "4.2.7.1 Multiprocessing of Nested Loops".

`-K{largepage|nolargepage}`

`-Klargepage` creates an executable program which uses the large page function. The `-Knolargepage` option invalidates the `-Klargepage` option. `-Klargepage` is set by default.

The `-Klargepage` and `-Knolargepage` option must be set at linking.

See Section "Appendix K Large Page Library" for details.

`-K{lib|nolib}`

`-Klib` promotes optimization by recognizing the operation of standard library functions. The `-Knolib` option invalidates the `-Klib` option. `-Knolib` is set by default when the `-O0` option is set. `-Klib` option is set by default when the `-O2` option or higher is set.

If a user-defined function with the same name as a standard library function is used, unexpected results may occur.

The `-Klib` and `-Knolib` options require that the `-O1` option or higher is set.

The recognized standard library functions are as follows:

```
abort, abs, acos, acosf, acosh, acoshf, asin, asinf, asinh, asinhf, atan, atan2, atan2f, atanf,
atanh, atanhf, calloc, cbrt, cbrtf, ceil, ceilf, clearerr, copysign, copysignf, cos, cosf, cosh,
coshf, csqrt, csqrtf, erf, erfc, erfcf, erff, exit, exp, exp2, exp2f, expf, expmlf, expmlf, fabs,
fabsf, fclose, fdim, fdimf, feof, ferror, fflush, fgetc, fgetpos, fgets, floor, floorf, fma, fmaf,
fmax, fmaxf, fmin, fminf, fmod, fmodf, fputc, fputs, free, frexp, frexpf, fseek, fwrite, getenv,
hypot, hypotf, ilogb, ilogbf, isalnum, isalpha, iscntrl, isdigit, isgraph, islower, ispunct,
isspace, isupper, isxdigit, ldexp, ldexpf, lgamma, lgammaf, llrint, llrintf, llround, llroundf,
log, log10, log10f, loglp, loglpf, log2, log2f, logb, logbf, logf, lrint, lrintf, lround, lroundf,
malloc, memchr, memcmp, memcpy, memmove, memset, modf, nearbyint, nearbyintf, nextafter,
nextafterf, nexttoward, nexttowardf, perror, pow, powf, printf, putchar, rand, realloc, remainder,
remainderf, remove, remquo, remquof, rename, rint, rintf, round, roundf, scalbln, scalblnf, scalbn,
scalbnf, scanf, setvbuf, sin, sinf, sinh, sinh, sprintf, sqrt, sqrtf, srand, sscanf, strcat,
strchr, strcmp, strcpy, strcspn, strerror, strlen, strncat, strncmp, strncpy, strpbrk, strrchr,
strspn, strstr, strtod, strtok, strtol, strtoul, system, tan, tanf, tanh, tanhf, tgamma, tgammaf,
tolower, toupper, trunc, truncf, vprintf, vsprintf
```

`-K{loop_blocking[=M]|loop_noblocking}`

`-Kloop_blocking` performs optimization by loop blocking. The `-Kloop_noblocking` option invalidates the `-Kloop_blocking` option. `-Kloop_noblocking` is always applied when the `-O1` option is effective. `-Kloop_blocking` is set by default when the `-O2` option or higher is set.

$N$  indicates the block size which should be from 2 to 10000. If  $N$  is omitted, the compiler automatically determines a suitable value for  $N$ .

The `-Kloop_blocking` and `-Kloop_noblocking` options require that the `-O1` option or higher is set.

For details about the loop blocking, see Section "[3.2.5 Loop Blocking](#)".

`-K{loop_fission|loop_nofission}`

`-Kloop_fission` specifies to perform the loop fission optimization for promotion of the software pipelining, improvement of the cache memory efficiency, and the resolution of register shortage. The `-Kloop_nofission` option invalidates the `-Kloop_fission` option. `-Kloop_nofission` is always applied when the `-O1` option is set. `-Kloop_fission` is set by default when the `-O2` option or higher is set.

The target loops and fission points of the loop fission optimization are specified by the following optimization control specifiers.

- `loop_fission_target` specifier

This optimization control specifier fissions the specified loop automatically.

- `fission_point` specifier

This optimization control specifier fissions the loop at the specified executable statement.

See Section "[3.4.1 Using Optimization Control Line \(OCL\)](#)" for information about `loop_fission_target` and `fission_point` specifiers. Note that you must specify the `-Kocl` option to validate optimization control specifiers.

The `-Kloop_fission` and `-Kloop_nofission` options require that the `-O1` option or higher is set.

The `-Kloop_fission` and `-Kloop_nofission` options work only when used in combination with above optimization control specifiers.

For details about this function, see Section "[3.3.10 Loop Fission](#)".

`-K{loop_fission_stripmining[={M|L1|L2}]|loop_nofission_stripmining}`

`-Kloop_fission_stripmining` specifies to perform optimization of strip-mining at the automatic loop fission. The `-Kloop_fission_stripmining` option invalidates the `-Kloop_nofission_stripmining` option.

If the argument `={M|L1|L2}` is omitted, the compiler choose the strip length automatically. `-Kloop_nofission_stripmining` is set by default.

By this optimization, improvement of the cache memory efficiency is expected for the data accessed between fissioned loops.

The `-Kloop_fission_stripmining` option requires that the `-Kloop_fission` option, the `-Kocl` option, and the `-O2` option or higher are set and the `loop_fission_target` specifier is specified in the optimization control line.

For details about this function, see Section "[3.3.10.1 Strip-Mining](#)".

`-Kloop_fission_stripmining=N`

*N* indicates the strip length. *N* should be an integer value from 2 to 100000000.

`-Kloop_fission_stripmining=L1`

The strip length is adjusted to the size of the level 1 cache for the cache memory efficiency.

`-Kloop_fission_stripmining=L2`

The strip length is adjusted to the size of the level 2 cache for the cache memory efficiency.

`-Kloop_fission_threshold=N`

Specifies the threshold *N* to decide the granularity of loops after automatic loop fission. *N* is an integer value from 1 to 100. If the value of *N* is reduced, the fissioned loops tend to become small and the number of the fissioned loops tends to increase. `-Kloop_fission_threshold=50` is set by default.

This option requires that the `-Kloop_fission` option, the `-Kocl` option, and the `-O2` option or higher are set and the `loop_fission_target` specifier is specified in the optimization control line.

`-K{loop_fusion|loop_nofusion}`

`-Kloop_fusion` fuses neighboring loops. The `-Kloop_nofusion` option invalidates the `-Kloop_fusion` option. `-Kloop_nofusion` is always applied when the `-O1` option is set. `-Kloop_fusion` is set by default when the `-O2` option or higher is set.

The `-Kloop_fusion` and `-Kloop_nofusion` options require that the `-O1` option or higher is set.

#### -K{loop\_interchange|loop\_nointerchange}

-Kloop\_interchange exchanges loops. The -Kloop\_nointerchange option invalidates the -Kloop\_interchange option. -Kloop\_nointerchange is always applied when the -O1 option is set. -Kloop\_interchange is set by default when the -O2 option or higher is set.

The -Kloop\_interchange and -Kloop\_nointerchange options require that the -O1 option or higher is set.

#### -K{loop\_part\_parallel|loop\_nopart\_parallel}

-Kloop\_part\_parallel specifies that when a loop contains statements to which parallelization can be applied and statements to which parallelization cannot be applied, the loop is divided into two or more loops and automatic parallelization is applied to one of the loops. This optimization is applied to the innermost loop. The -Kloop\_nopart\_parallel option invalidates the -Kloop\_part\_parallel option. -Kloop\_nopart\_parallel is set by default.

Note that compile time and execution time may increase.

The -Kloop\_part\_parallel and -Kloop\_nopart\_parallel options require that the -Kparallel option is set.

#### -K{loop\_part\_simd|loop\_nopart\_simd}

-Kloop\_part\_simd specifies that when a loop contains instructions to which SIMD extensions can be applied and instructions to which SIMD extensions cannot be applied, the loop is divided into two or more loops and SIMD extensions are applied to one of the loops. This optimization is applied to the innermost loop. The -Kloop\_nopart\_simd option invalidates the -Kloop\_part\_simd option. -Kloop\_nopart\_simd is set by default.

Note that compile time and execution time may increase.

The -Kloop\_part\_simd and -Kloop\_nopart\_simd options require that the -Ksimd[={1|2|auto}] option is set.

#### -K{loop\_perfect\_nest|loop\_noperfect\_nest}

-Kloop\_perfect\_nest performs optimization which fissions the imperfectly nested loop into the perfectly nested loops. The -Kloop\_noperfect\_nest option invalidates the -Kloop\_perfect\_nest option. The -Kloop\_noperfect\_nest is set by default when the -O2 option is set, and the -Kloop\_perfect\_nest is set by default when the -O3 option is set.

This optimization may promote optimizations such as the loop interchange and loop collapsing.

The -Kloop\_perfect\_nest and -Kloop\_noperfect\_nest options require that the -O2 option or higher is set.

#### -K{loop\_versioning|loop\_noversioning}

-Kloop\_versioning performs optimization by the loop versioning. The -Kloop\_noversioning option invalidates the -Kloop\_versioning option. -Kloop\_noversioning is set by default.

The loop versioning may promote optimizations such as SIMD extension, software pipelining, or automatic parallelization.

Note that the size of the object program and compile time may increase because the loop versioning generates two loops. Moreover, the execution performance of the program may decrease due to the overhead of judgement for the choice of generated loops.

The -Kloop\_versioning and -Kloop\_noversioning options require that the -O2 option or higher is set.

For details about the loop versioning, see Section "[3.3.6 Loop Versioning](#)".

#### -Klooptype={f|n|s}

Specifies loop continuity conditions and direction. This option may promote optimizations. -Klooptype=f is set by default.

##### -Klooptype=f

Declares that the program does not contain an infinite loop but may contain a wrap-around loop.

##### -Klooptype=n

Declares that the program may contain an infinite loop or wrap-around loop.

##### -Klooptype=s

Declares that the program does not contain an infinite or wrap-around loop.

#### -K{memalias|nomemalias}

-Kmemalias specifies that in the case of indirect memory access through a pointer, when the accessing types are different, no memory alias is set. The -Knomemalias option invalidates the -Kmemalias option. -Knomemalias is set by default.

The -Kmemalias option may promote optimizations.

The -Kmemalias and -Knomemalias options require that the -O1 option or higher is set.

The -Kmemalias and -Knomemalias options are deprecated and will be removed in the future.

#### -K{mfunc={1|2|3}}|nomfunc}

-Kmfunc specifies to apply the optimization which changes a function into a multi-operation function. A multi-operation function is a function achieving improved execution performance using one tuned function call instead of performing multiple functions. If the argument ={1|2|3} is omitted, -Kmfunc=1 is set. The -Knomfunc option invalidates the -Kmfunc option. -Knomfunc is set by default.

When the -Kmfunc option is set, side effects may occur in the execution result. See Section "[3.6.1 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

The -Kmfunc and -Knomfunc options require that the -O2 option or higher is set.

The -Klib option is required for the -Kmfunc option to take effect.

The following table shows the function of the target.

Functions	Float Type	Double Type
acos	* (a)	* (a)
asin	* (a)	* (a)
atan	*	*
atan2	*	*
cos	*	*
erf	* (a)	* (a)
erfc	* (a)	* (a)
exp	*	*
log	*	*
log10	*	*
sin	*	*
pow	*	*

\*: To be a target

-: Not to be a target

a) These multi-operation functions execute sequential instructions internally. The execution performance improves because optimizations, such as branch optimization, are applied to the internal instructions. The improvement is inferior to that of the normal multi-operation functions.

#### -Kmfunc=1

A multi-operation function of multiplicities which is the same as the SIMD width is used.

#### -Kmfunc=2

In addition to the function of -Kmfunc=1, a multi-operation function of multiplicities automatically determined by the compiler is also used. This function needs a large stack area.

#### -Kmfunc=3

In addition to the function of -Kmfunc=2, a multi-operation function of multiplicities automatically determined by the compiler is also used in a loop which contains an "if" statement or similar. This function needs a large stack area. The execution performance may be worse than with -Kmfunc=1 or -Kmfunc=2 when the true ratio of the "if" statement is low.

Moreover, this optimization has side effects which may cause abend. See Section "[3.3.3.2 Effects of Compiler Option -Kmfunc=3](#)" for more information about the side effects.

## -Knoprefetch

Specifies to generate an object without a prefetch instruction.

## -K{ocl|noocl}

-Kocl specifies that the optimization control line (OCL) is effective. The -Knoocl option invalidates the -Kocl option. -Kocl is set by default.

The -Kocl and -Knoocl options require that the -O1 option or higher is set.

## -K{omitfp|noomitfp}

The -Komitfp option performs the optimization that the frame pointer register is not kept for a function call. The -Knoomitfp option invalidates the -Komitfp option. -Knoomitfp is set by default.

Note that when the -Komitfp option is specified, the trace back information is not kept.

The -Komitfp and -Knoomitfp options require that the -O1 option or higher is set.

## -K{openmp|noopenmp}

-Kopenmp specifies to enable the Specification of OpenMP Application Program Interface. The -Knoopenmp option invalidates the -Kopenmp option. -Knoopenmp is set by default.

When the -Kopenmp option is specified, -mt is assumed.

The -Kopenmp option is needed if an object program compiled with it exists in the command line as an input file.

When the -O1 option or less is effective, SIMD Extensions are not used even if `simd construct` or `declare simd construct` is effective.

## -K{openmp\_assume\_norecurrence|openmp\_noassume\_norecurrence}

These options direct whether or not to promote optimizations by assuming that array elements of the chunk size have no data dependency over iteration for the loop with the OpenMP "for" directive. -Kopenmp\_noassume\_norecurrence is set by default.

When -Kopenmp\_assume\_norecurrence is effective, SIMD extensions, software pipelining, and other optimizations are promoted.

The -Kopenmp\_assume\_norecurrence and -Kopenmp\_noassume\_norecurrence options are applied to the innermost loops with the OpenMP "for" directive.

The -Kopenmp\_assume\_norecurrence and -Kopenmp\_noassume\_norecurrence options are effective only when the -Kopenmp option and -O2 option or higher is set.

## -K{openmp\_collapse\_except\_innermost|openmp\_nocollapse\_except\_innermost}

In the loop construct of OpenMP, this option specifies whether to remove the loop from the object of collapse when it meets all the following conditions:

- The `collapse` clause that includes the innermost loop is specified.
- When the compiler can judge that there is a possibility that the execution performance decreases by collapse that includes the innermost loop at the compile time.

When the -Kopenmp\_collapse\_except\_innermost is specified, the loop is removed from the object of collapse. -Kopenmp\_nocollapse\_except\_innermost is set by default.

The decrease of execution performance by collapse might be able to be prevented by -Kopenmp\_collapse\_except\_innermost option.

The -Kopenmp\_collapse\_except\_innermost and -Kopenmp\_nocollapse\_except\_innermost options are effective only when the -Kopenmp option is set.

The loop removed from the object of collapse can be known from the diagnostic message by setting the -Koptmsg=2 option at the compile time.

## -Kopenmp\_loop\_variable={private|standard}

Indicates whether the sequential loop control variable in the OpenMP parallel syntax should be treated as `private` or `shared`. The -Kopenmp\_loop\_variable=private option is default.

This option is effective only when the -Kopenmp option is set.

`-Kopenmp_loop_variable=private`

Indicates the sequential loop control variable in the OpenMP parallel syntax is treated as `private`.

The handling of the loop control variable does not conform to the standard.

`-Kopenmp_loop_variable=standard`

Indicates the sequential loop control variable in the OpenMP parallel syntax is treated as `shared`.

The handling of the loop control variable conforms to the standard.

`-K{openmp_ordered_reduction|openmp_noordered_reduction}`

`-Kopenmp_ordered_reduction` specifies to fix the order of the reduction operations same as in the numerical order of the threads at the end of the region for which the `reduction` clause of OpenMP was specified. The `-Kopenmp_noordered_reduction` option invalidates the `-Kopenmp_ordered_reduction` option. `-Kopenmp_noordered_reduction` is set by default.

If the number of threads used is identical, by fixing the order of the reduction operations same as in the numerical order of the threads, identical results will be always obtained. However, rounding errors may occur when the operation order is changed by the effect of the scheduling of a loop construct with a dynamic or a guided schedule kind, or a `sections` construct. Note that the execution performance may decrease compared to when the `-Kopenmp_ordered_reduction` option is invalidated.

The `-Kopenmp_ordered_reduction` and `-Kopenmp_noordered_reduction` options require that the `-Kopenmp` option is set.

`-K{openmp_simd|noopenmp_simd}`

`-Kopenmp_simd` specifies that only the `simd` construct and the `declare simd` construct of OpenMP are enabled.

`-Knoopenmp_simd` is set by default.

When the `-Kopenmp_simd` option is specified, only the SIMD Extensions based on the OpenMP specifications are applied, and the parallelization is not applied.

When the `-Knoopenmp_simd` option is specified, the `-Kopenmp_simd` option is disabled.

When the `-O1` option or less is enabled, SIMD Extensions are not used even if `simd` construct or `declare simd` construct is enabled.

`-K{optlib_string|nooptlib_string}`

`-Koptlib_string` specifies the compiler to link the library of the optimized version statically about the string handling functions. The `-Knooptlib_string` option invalidates the `-Koptlib_string` option. `-Knooptlib_string` is set by default.

`-Koptlib_string` option is effective only when the `-KSVE` option and `-KA64FX` option are set.

These options must be set at linking.

The following shows the string handling functions for which the library of the optimized version can be linked.

```
bcopy, bzero, memchr, memcmp, memccpy, memcpy, memmove, memset, strcat, strcmp, strcpy, strlen,
strncmp, strncpy, strncat
```

`-K{optmsg[={1|2}]|nooptmsg}`

`-Koptmsg` specifies to output the messages of applied optimizations. If the argument `={1|2}` is omitted, `-Koptmsg=1` is set. The `-Knooptmsg` option invalidates the `-Koptmsg` option. `-Knooptmsg` is set by default.

The `-Koptmsg` and `-Knooptmsg` options require that the `-O1` option or higher is set.

`-Koptmsg=1`

This option outputs messages about applied optimizations with side effects on the execution result.

`-Koptmsg=2`

In addition to the output of `-Koptmsg=1`, messages are output about optimizations such as automatic parallelization, SIMD extend instruction conversion, and loop unrolling.

`-K{parallel|noparallel}`

The compiler option `-Kparallel` specifies to perform automatic parallelization. However, if the effect of parallel execution is not expected, automatic parallelization is not performed. The compiler option `-Knoparallel` invalidates the compiler option `-Kparallel`. The compiler option `-Knoparallel` is set by default.

If the compiler option `-Kparallel` is set, the compiler options `-O2`, `-Kregion_extension`, `-Kloop_part_parallel`, `-Kloop_perfect_nest`, and `-mt` also take effect. Note that if the compiler option `-O3` is set at the same time, the compiler option `-O3` takes effect.

The compiler option `-Kparallel` is invalidated if the compiler option `-O0` or `-O1` is set.

The compiler option `-Kparallel` is needed if an object program compiled with it exists in the command line as an input file.

`-K{parallel_fp_precision|parallel_nofp_precision}`

`-Kparallel_fp_precision` specifies that the compiler controls to apply optimizations considering calculation error of a floating type or a complex type operation that is caused by the difference of the parallel number of threads. The `-Kparallel_nofp_precision` option invalidates the `-Kparallel_fp_precision` option. `-Kparallel_nofp_precision` is set by default.

When the `-Kparallel_fp_precision` option is specified, the compiler applies optimizations within a range in which no calculation errors occur.

The `-Kparallel_fp_precision` and `-Kparallel_nofp_precision` options require that the `-Kparallel` or `-Kopenmp` option is set.

When the `-Kparallel_fp_precision` option and `-Kopenmp` option are set, `-Kopenmp_ordered_reduction` option is valid.

When the `-Kparallel_fp_precision` option is set, the execution performance may decrease because the part of optimization is restricted.

Note that the calculation error of a floating type or a complex type operation that is caused by the difference of the parallel number of threads even if the `-Kparallel_fp_precision` option is valid when the reduction clause of OpenMP is specified.

`-Kparallel_iteration=N`

Specifies that parallelization targets only the loop which has proved that the iteration counts is  $N$  or more at compilation.  $N$  can be specified from 1 to 2147483647.

This option requires that the `-Kparallel` option is set.

`-Kparallel_strong`

Specifies to perform automatic parallelization for all loops that can be made parallel without estimating parallelization effects.

This option sets the `-Keval` and `-Kpreex` options.

Aside from this point, the function and notes are same as for the `-Kparallel` option.

`-K{pc_relative_literal_loads|nopc_relative_literal_loads}`

`-Kpc_relative_literal_loads` option specifies to limit the size of text data area in the function to 1MB, and access to the literal pool by one instruction. The `-Knopc_relative_literal_loads` option invalidates the `-Kpc_relative_literal_loads` option. `-Knopc_relative_literal_loads` is set by default.

`-Kpc_relative_literal_loads` and `-Knopc_relative_literal_loads` options are effective at compilation time only.

`-K{plt|noplt}`

`-Kplt` specifies to use Procedure Linkage Table (PLT) for accessing a global symbol in the position-independent code (PIC). `-Knoplt` option invalidates the `-Kplt` option. `-Kplt` is set by default.

`-Kplt` and `-Knoplt` options are effective only if the `-KPIC` option or the `-Kpic` option is set.

`-K{preex|nopreex}`

`-Kpreex` specifies to evaluate invariant expressions in advance. The `-Knopreex` option invalidates the `-Kpreex` option. `-Knopreex` is set by default.

When the `-Kpreex` option is set, side effects (calculation error or runtime exceptions) may occur in the execution results because instructions that may not to be executed based on the logic of the program are likely to be executed. For detail of the side effects, see Section "[3.3.1.1 Advance Evaluation of Invariant Expressions](#)".

The `-Kpreex` and `-Knopreex` options require that the `-O1` option or higher is set. Note that a floating-point exception may occur when the `-Knopreex`, `-Kfp_relaxed`, and `-NRtrap` options are valid.

`-Kprefetch_cache_level={1|2|all}`

Specifies the data cache-level of prefetch instructions. By default, `-Kprefetch_cache_level=all` is set.

This option requires that at least one of the `-Kprefetch_indirect`, `-Kprefetch_sequential`, or `-Kprefetch_stride` options is set.

`-Kprefetch_cache_level=1`

Data is prefetched in the first level cache. Normal prefetch instructions are used.

`-Kprefetch_cache_level=2`

Data is prefetched in the second level cache.

`-Kprefetch_cache_level=all`

Both `-Kprefetch_cache_level=1` and `-Kprefetch_cache_level=2` functions are effective. By using two levels of prefetch instructions, the prefetch function becomes more sophisticated.

`-K{prefetch_conditional|prefetch_noconditional}`

`-Kprefetch_conditional` specifies to generate prefetch instructions for array data in the block included in an "if" statement and a "switch" statement. The `-Kprefetch_noconditional` option invalidates the `-Kprefetch_conditional` option. `-Kprefetch_noconditional` is set by default.

The `-Kprefetch_conditional` and `-Kprefetch_noconditional` options require that at least one of the `-Kprefetch_indirect`, `-Kprefetch_sequential`, or `-Kprefetch_stride` options is set.

`-K{prefetch_indirect|prefetch_noindirect}`

`-Kprefetch_indirect` specifies to generate prefetch instructions for array data accessed indirectly (list access) within a loop. The `-Kprefetch_noindirect` option invalidates the `-Kprefetch_indirect` option. `-Kprefetch_noindirect` is set by default.

The `-Kprefetch_indirect` and `-Kprefetch_noindirect` options require that the `-O1` option or higher is set.

`-K{prefetch_infer|prefetch_noinfer}`

`-Kprefetch_infer` specifies to generate the prefetch instructions under the assumption that the memory access is continuous. The `-Kprefetch_noinfer` option invalidates the `-Kprefetch_infer` option. `-Kprefetch_noinfer` is set by default.

The `-Kprefetch_infer` and `-Kprefetch_noinfer` options require that at least one of the `-Kprefetch_indirect`, `-Kprefetch_sequential`, or `-Kprefetch_stride` options is set.

`-Kprefetch_iteration=N`

Specifies to prefetch data which is referred to or defined in a loop after  $N$  iteration(s).  $N$  should be from 1 to 10000. If the optimization using SIMD extensions is applied to a loop, specify the loop iteration count after using SIMD extensions for  $N$ .

This option specifies to prefetch data to only the first level cache.

This option requires that at least one of the `-Kprefetch_indirect`, `-Kprefetch_sequential`, or `-Kprefetch_stride` options is set, and the `-Kprefetch_cache_level=1` or `-Kprefetch_cache_level=all` is set.

This option cannot be specified simultaneously with the `-Kprefetch_line` option.

`-Kprefetch_iteration_L2=N`

Specifies to prefetch data which is referred to or defined in a loop after  $N$  iteration(s).  $N$  should be from 1 to 10000. If the optimization using SIMD extensions is applied to a loop, specify the loop iteration count after using SIMD extensions for  $N$ .

This option specifies to prefetch data to only the second level cache.

This option requires that at least one of the `-Kprefetch_indirect`, `-Kprefetch_sequential`, or `-Kprefetch_stride` options is set, and `-Kprefetch_cache_level=2` or `-Kprefetch_cache_level=all` is set.

This option cannot be specified simultaneously with the `-Kprefetch_line_L2` option.

`-Kprefetch_line=N`

Specifies to prefetch data which is referred to or defined in a line after  $N$  line(s).  $N$  should be from 1 to 100.

This option specifies to prefetch data to only the first level cache.

This option requires that at least one of the `-Kprefetch_indirect` or `-Kprefetch_sequential` options is set, and `-Kprefetch_cache_level=1` or `-Kprefetch_cache_level=all` is set.

This option cannot be specified simultaneously with the `-Kprefetch_iteration` option.



#### `-Kprefetch_line_L2=N`

Specifies to prefetch data which is referred to or defined in a line after  $N$  line(s).  $N$  should be from 1 to 100. This option specifies to prefetch data to only the second level cache.

This option requires that at least one of the `-Kprefetch_indirect` or `-Kprefetch_sequential` options is set, and `-Kprefetch_cache_level=2` or `-Kprefetch_cache_level=all` is set.

This option cannot be specified simultaneously with the `-Kprefetch_iteration_L2` option.

#### `-K{prefetch_sequential[={auto|soft}]]prefetch_nosequential}`

`-Kprefetch_sequential` specifies to generate prefetch instructions for prefetch array data accessed sequentially within a loop. The `-Kprefetch_nosequential` option invalidates the `-Kprefetch_sequential` option. When the `-O1` option is set, `-Kprefetch_nosequential` is set by default. When the `-O2` option or higher is set, `-Kprefetch_sequential` is set by default.

If the argument `={auto|soft}` is omitted, `-Kprefetch_sequential=auto` is assumed.

The `-Kprefetch_sequential` and `-Kprefetch_nosequential` options require that the `-O1` option or higher is set.

#### `-Kprefetch_sequential=auto`

The compiler automatically selects whether to use the hardware prefetch function or to output the prefetch instruction for array data accessed sequentially within a loop.

#### `-Kprefetch_sequential=soft`

The compiler specifies to output the prefetch instruction for array data accessed sequentially within a loop without using the hardware prefetch function.

#### `-K{prefetch_stride[={soft|hard_auto|hard_always}]]prefetch_nostride}`

`-Kprefetch_stride` specifies to perform the prefetch for array data with a stride larger than the cache line size in its loop. This includes loops with prefetch addresses not defined at compilation. The `-Kprefetch_nostride` option invalidates the `-Kprefetch_stride` option. `-Kprefetch_nostride` is set by default.

If `"={soft|hard_auto|hard_always}"` of the `-Kprefetch_stride` option is omitted, the `-Kprefetch_stride=soft` option is set by default.

The `-Kprefetch_stride=soft` and `-Kprefetch_nostride` options require that the `-O1` option or higher is set.

The `-Kprefetch_stride=hard_auto` and `-Kprefetch_stride=hard_always` options require that the `-Khpctag` option and the `-O1` option or higher are set.

#### `-Kprefetch_stride=soft`

The `-Kprefetch_stride=soft` option specifies to generate prefetch instructions and perform the prefetch.

#### `-Kprefetch_stride=hard_auto`

The `-Kprefetch_stride=hard_auto` option specifies to perform the prefetch using the hardware stride prefetcher.

This option sets the stride prefetcher to prefetch only data that is not on cache.

See Section ["3.6.3 Optimizations Using the Hardware Stride Prefetcher"](#) for the prefetch using the hardware stride prefetcher.

#### `-Kprefetch_stride=hard_always`

The `-Kprefetch_stride=hard_always` option specifies to perform the prefetch using the hardware stride prefetcher.

In contrast with the `-Kprefetch_stride=hard_auto` option, this option sets the stride prefetcher to always prefetch.

See Section ["3.6.3 Optimizations Using the Hardware Stride Prefetcher"](#) for the prefetch using the hardware stride prefetcher.

#### `-K{prefetch_strong|prefetch_nostrong}`

`-Kprefetch_strong` specifies to generate strong prefetch instructions. The `-Kprefetch_nostrong` option invalidates the `-Kprefetch_strong` option. `-Kprefetch_strong` is set by default.

The `-Kprefetch_strong` option specifies to prefetch data to only the first level cache.

The `-Kprefetch_strong` and `-Kprefetch_nostrong` options require that at least one of the `-Kprefetch_indirect`, `-Kprefetch_sequential`, or `-Kprefetch_stride` options is set, and the `-Kprefetch_cache_level=1` or `-Kprefetch_cache_level=all` option is set.

The `-Kprefetch_nostrong` option requires that the `-Khpctag` option is set.

#### -K{prefetch\_strong\_L2|prefetch\_nostrong\_L2}

-Kprefetch\_strong\_L2 specifies to generate strong prefetch instructions. The -Kprefetch\_nostrong\_L2 option invalidates the -Kprefetch\_strong\_L2 option. -Kprefetch\_strong\_L2 is set by default.

The -Kprefetch\_strong\_L2 option specifies to prefetch data to only the second level cache.

The -Kprefetch\_strong\_L2 and -Kprefetch\_nostrong\_L2 options require that at least one of the -Kprefetch\_indirect, -Kprefetch\_sequential, or -Kprefetch\_stride options is set, and the -Kprefetch\_cache\_level=2 or -Kprefetch\_cache\_level=all option is set.

The -Kprefetch\_nostrong\_L2 option requires that the -Khpctag option is set.

#### -K{preload|nopreload}

-Kpreload specifies to perform the speculative execution of load instructions. By specifying the -Kpreload option, the optimization of instruction scheduling is promoted, and the execution performance is expected to improve.

The -Knopreload option invalidates the -Kpreload option. The -Knopreload is set by default.

Note that this optimization may cause interruption of execution of the program by the execution of load instructions referring to illegal areas that are not to be executed logically in the program.

The -Kpreload option is effective only when the -O1 option or higher is set.

#### -K{rdconv[={1|2}]|nordconv}

-Krdconv specifies to perform on the assumption that expressions of integer type whose size is 4-byte or less do not overflow or wrap around when the expressions increase or decrease by a constant value as the loop repeats.

The -Knordconv option invalidates the -Krdconv option. If "{1|2}" of -Krdconv is omitted, the -Krdconv=1 option is assumed to be set.

-Krdconv=1 is set by default when the -O2 or higher option is set.

If the expression overflows or wrap around, the result may differ depending on the presence or absence of the -Krdconv option.

These options require that the -O2 option or higher is set.

##### -Krdconv=1

This option specifies to perform optimizations on the assumption that expressions of signed integer type whose size is 4-byte or less do not overflow when the expressions increase or decrease by a constant value as the loop repeats.

##### -Krdconv=2

In addition to the features of the -Krdconv=1 option, this option specifies to perform on the assumption that expressions of unsigned integer type whose size is 4-byte or less do not wrap around when the expressions increases or decreases by a constant value as the loop repeats.

#### -K{reduction|noreduction}

-Kreduction specifies to perform the reduction optimization. The -Knoreduction option invalidates the -Kreduction option. -Knoreduction is set by default.

When the -Kreduction option is set, side effects (calculation errors) may occur in the execution results. See Section ["3.6.1 Side Effect of Optimizations for Floating-Point Operation"](#) for the side effect of optimization.

The -Kreduction and -Knoreduction options require that the -Kparallel option is set.

Before using the -Kreduction option, see Section ["4.2.1.1 Compiler Option for Automatic Parallelization"](#).

#### -K{region\_extension|noregion\_extension}

-Kregion\_extension specifies to perform the optimization that extends the parallel region to reduce the overhead caused by automatic parallelization.

If the -Kregion\_extension option is specified, the execution performance may be degraded on loops with a small parallelization effect.

The -Kregion\_extension and -Knoregion\_extension options require that the -Kparallel option is set.

For details, see Section ["4.2.5.10 Parallel Region Extension"](#).

`-K{remove_inlinefunction|noremove_inlinefunction}`

`-Kremove_inlinefunction` specifies to remove unreferenced extern inline function definition for object size reduction. The `-Knoremove_inlinefunction` option invalidates the `-Kremove_inlinefunction` option. `-Kremove_inlinefunction` is set by default.

The `-Kremove_inlinefunction` and `-Knoremove_inlinefunction` options require that the `-O1` or higher option is set.

`-K{restp[={all|arg}]}|norestp}`

`-Krestp` specifies to perform optimization of restricted pointers and references. If the argument `={all|arg}` is omitted, `-Krestp=arg` is applied. The `-Knorestp` option invalidates the `-Krestp` option. `-Knorestp` is set by default.

If the `-Krestp` option is set, side effects (calculation errors or runtime exceptions) may occur in the execution results.

The `-Krestp` and `-Knorestp` options require that the `-O1` option or higher is set.

Before using the `-Krestp` option, see Section "[3.3.2 Optimization of Pointers](#)".

`-Krestp=all`

Specifies to perform optimization of restricted pointers assuming that the `restrict` qualifier in the C99 is specified for all pointers and references.

`-Krestp=arg`

Specifies to perform optimization of restricted pointers assuming that the `restrict` qualifier in the C99 is specified for the pointers, references and the `"this"` pointers in arguments.

`-K{sch_post_ra|nosch_post_ra}`

`-Ksch_post_ra` specifies to perform instruction scheduling after register allocation. The `-Knosch_post_ra` option invalidates the `-Ksch_post_ra` option. When the `-O0` option is set, `-Knosch_post_ra` is set by default. When the `-O1` option or higher is set, `-Ksch_post_ra` is set by default.

The compiler rearranges execution instructions to improve the execution performance when the `-Ksch_post_ra` option is effective. This optimization does not increase saving and restoring instructions for registers to and from the memory.

`-K{sch_pre_ra|nosch_pre_ra}`

`-Ksch_pre_ra` specifies to perform instruction scheduling before register allocation. The `-Knosch_pre_ra` option invalidates the `-Ksch_pre_ra` option. When the `-O0` option is set, the default is `-Knosch_pre_ra`. When the `-O1` option or higher is set, the default is `-Ksch_pre_ra`.

The compiler rearranges execution instructions to improve the execution performance when the `-Ksch_pre_ra` option is effective. Note that the execution performance may decrease by increasing saving and restoring instructions for registers to and from the memory.

`-K{sibling_calls|nosibling_calls}`

`-Ksibling_calls` specifies to optimize sibling calls. The `-Knosibling_calls` option invalidates the `-Ksibling_calls` option. `-Knosibling_calls` is set by default when the `-O1` option is set. `-Ksibling_calls` is set by default when the `-O2` or higher option is set.

When the `-Ksibling_calls` option is set, the trace back information is not kept.

The `-Ksibling_calls` and `-Knosibling_calls` options require that the `-O1` option or higher is set.

`-K{simd[={1|2|auto}]}|nosimd}`

`-Ksimd` specifies to perform optimization using the SIMD Extensions for loop.

If the argument `={1|2|auto}` is omitted, `-Ksimd=auto` is assumed. The `-Knosimd` option invalidates the `-Ksimd` option.

When the `-O2` option or higher is set, `-Ksimd=auto` is set by default.

`-Kloop_part_simd` option is also effective when the `-Ksimd[={1|2|auto}]` option is set.

The `-Ksimd` and `-Knosimd` options require that the `-O2` option or higher is set.

For details about SIMD, see Section "[3.2.7 SIMD](#)".

`-Ksimd=1`

Specifies optimization that uses the SIMD Extensions for loop excluding conditional branch (like `"if"` statement).

## -Ksimd=2

In addition to the function of -Ksimd=1, -Ksimd=2 specifies optimization that uses the SIMD Extensions for loop including conditional branch.

The generated code has redundant instructions in conditional branches, so execution time may increase according to the true-ratio of the condition.

-Ksimd=2 may produce side effects (calculation errors or runtime exceptions) in the execution results because instructions that may not be executed based on the logic of the program are likely to be executed, as the result of the speculative execution for the expressions in the "if" statement.

## -Ksimd=auto

The compiler automatically determines whether to use SIMD Extensions for the loop. SIMD Extensions are promoted for loops that contain "if" statement.

## -K{simd\_packed\_promotion|simd\_nopacked\_promotion}

-Ksimd\_packed\_promotion specifies to perform optimization promoting packed SIMD assuming that index calculations of array elements of both single-precision floating-point type and 4-byte integer type do not exceed 4-byte range. -Ksimd\_nopacked\_promotion is set by default.

When the -Ksimd\_packed\_promotion is effective and the index calculations exceed 4-byte range, the program execution may be aborted or the execution result may be incorrect by referring illegal area.

-Ksimd\_packed\_promotion and -Ksimd\_nopacked\_promotion options require that the -Ksimd={1|2|auto} option is set.

## -K{simd\_reduction\_product|simd\_noreduction\_product}

-Ksimd\_reduction\_product specifies to perform SIMD extensions to the reduction operation of product. The -Ksimd\_noreduction\_product option invalidates the -Ksimd\_reduction\_product option. -Ksimd\_noreduction\_product is set by default.

When the -Ksimd\_reduction\_product option is set, side effects (calculation errors) may occur in the execution results. See Section "[3.6.1 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

The -Ksimd\_reduction\_product and -Ksimd\_noreduction\_product options require that the -Ksimd[={1|2|auto}] option is set.

This option cannot be specified with the -Ksimd\_reg\_size=agnostic option at the same time.

## -Ksimd\_reg\_size={128|256|512|agnostic}

Specifies the size of the SVE vector register. Units are bits.

When the -Ksimd\_reg\_size={128|256|512} option is specified, optimizations are performed on the assumption that the vector register size is a fixed value specified as the option at compilation time. Therefore, optimizations are promoted and the improvement of the execution performance is expected. However, the generated executable program works normally only on CPU architecture which has the same size of the SVE vector register as the size specified at compilation time. For details, see Section "[3.6.2 Notes on Specified SVE Vector Register Size](#)".

When specifies -Ksimd\_reg\_size=agnostic, the SVE vector register is not considered to be a specific size, and the executable program decides the vector register size at execution time. The executable program does not depend on the SVE vector register size on the CPU architecture. Note that the execution performance might decrease compared with the case of -Ksimd\_reg\_size={128|256|512}.

-Ksimd\_reg\_size=512 is set by default.

This option is effective when the -KSVE option is set.

## -K{simd\_uncounted\_loop|simd\_nouncounted\_loop}

-Ksimd\_uncounted\_loop specifies to create objects that use SIMD extension instructions for statements in a while loop, do-while loop, if-goto loop, and a "for" loop with jumps out of the loop. The -Ksimd\_nouncounted\_loop option invalidates the -Ksimd\_uncounted\_loop option. -Ksimd\_nouncounted\_loop is set by default.

This option is effective when the -Ksimd[={1|2|auto}] and -KSVE options are set.

## -K{simd\_use\_multiple\_structures|simd\_nouse\_multiple\_structures}

-Ksimd\_use\_multiple\_structures specifies to use the SVE Load Multiple Structures and SVE Store Multiple Structures instructions when using SIMD extensions. The execution performance is improved by utilizing the above instructions for the load and store when

using SIMD extensions. Note that performance may decrease depending on data alignment. The `-Ksimd_nouse_multiple_structures` option invalidates the `-Ksimd_use_multiple_structures` option.

`-Ksimd_use_multiple_structures` is set by default.

This option is effective when the `-Ksimd[={1|2|auto}]` and the `-KSVE` options are set.

`-K{strict_aliasing|nostrict_aliasing}`

The `-Kstrict_aliasing` option specifies to allow optimizations with considering overlaps of memory regions according to the strict aliasing rules defined by language standard. `-Knostrict_aliasing` option is set by default.

The `-Kstrict_aliasing` and `-Knostrict_aliasing` options require that the `-O2` option or higher is set.

See section [3.3.11 Strict Aliasing](#) when specifying the `-Kstrict_aliasing` option.

`-K{striping[=N]|nostriping}`

`-Kstriping` specifies to apply loop striping.  $N$  is the striping size which should be from 2 to 100. If the argument  $=N$  is omitted, `-Kstriping=2` is assumed. The `-Knostriping` option invalidates the `-Kstriping` option. `-Knostriping` is always applied when the `-O1` option is set. The `-Knostriping` option is set by default when the `-O2` or higher option is set.

The `-Kstriping` and `-Knostriping` options requires that the `-O1` option or higher is set.

For details about loop striping, see Section "[3.3.4 Loop Striping](#)".

`-K{swp|noswp}`

`-Kswp` specifies to apply software pipelining. However, if the effect of software pipelining is not expected, software pipelining is not performed. The `-Knoswp` option invalidates the `-Kswp` option. `-Knoswp` is always applied when the `-O1` option is set. `-Kswp` is set by default when the `-O2` or higher option is set.

The `-Kswp` and `-Knoswp` options require that the `-O1` option or higher is set.

If the `-Kswp` option is specified with the `-Kswp_weak` or `-Kswp_strong` option, the one specified last is effective.

For details about software pipelining, see Section "[3.2.6 Software Pipelining](#)".

`-K{swp_freq_rate=N|swp_ireg_rate=N|swp_preg_rate=N}`

Specifies the rate (percentage) about the following registers that can be used by software pipelining.

- Floating-point register and SVE vector register
- Integer register
- SVE predicate register

$N$  should be an integer value number from 1 to 1000. `-Kswp_freq_rate=100,swp_ireg_rate=100,swp_preg_rate=100` is set by default.

The application of the software pipelining can be adjusted by changing the condition of the number of registers. If the software pipelining is not applied due to shortage of register, it may be applied by specifying integer values that are larger than 100.

Specifying this option may increase saving and restoring instructions for registers to and from the memory, and the execution performance may decrease.

This option is effective only if the `-O2` option or higher is set.

`-Kswp_freq_rate=N`

Specifies that  $N\%$  of the floating-point register and the SVE vector register is available when applying the software pipelining.

`-Kswp_ireg_rate=N`

Specifies that  $N\%$  of the integer register is available when applying the software pipelining.

`-Kswp_preg_rate=N`

Specifies that  $N\%$  of the SVE predicate register is available when applying the software pipelining.

`-Kswp_policy={auto|small|large}`

Specifies a policy to select an instruction scheduling algorithm used in software pipelining.

Software pipelining is performed by the `-Kswp` option, `-Kswp_weak` option, `-Kswp_strong` option, or the corresponding optimization control lines.

-Kswp\_policy=auto is set by default.

-Kswp\_policy=auto

The compiler automatically selects a fit algorithm for each loop.

-Kswp\_policy=small

An algorithm fit for a small loop, such as a loop with low register pressure, is used.

-Kswp\_policy=large

An algorithm fit for a large loop, such as a loop with high register pressure, is used.

-Kswp\_strong

Specifies to perform software pipelining for more loops by easing its applicable condition.

This option may increase requirement of time and memory significantly at compilation time.

If this option is specified with the -Kswp or -Kswp\_weak option, the one specified last is effective.

Aside from this point, the functions and notes are the same as for the -Kswp option.

-Kswp\_weak

Specifies to adjust software pipelining for the target loop and reduce overlap of instructions in the loop.

When the loop iteration count is uncertain and small, the effect of the optimization is expected because the loop iteration count required to execute the software-pipelined loop becomes small.

This option may decrease the execution performance because the overlaps of instructions become small.

If this option is specified with the -Kswp or -Kswp\_strong option, the one specified last is effective.

Aside from this point, the functions and notes are the same as for the -Kswp option.

-Ktls\_size={12|24|32|48}

Specifies the size of an offset necessary for the access to Thread-Local Storage. Units are bits.

As the size of Thread-Local Storage, -Ktls\_size=12 (4K bytes), -Ktls\_size=24 (16M bytes), -Ktls\_size=32 (4G bytes) or -Ktls\_size=48 (256T bytes) can be specified.

When the size of the Thread-Local Storage exceeds the range of the offset, an error occurs at link time.

These options are invalidated when the -KPIC or -Kpic option is specified simultaneously.

-K{unroll[=*N*]|nounroll}

-Kunroll specifies to apply loop unrolling. *N* is the upper limit of the unrolling expansion number which should be from 2 to 100. If =*N* is omitted, the compiler automatically determines a suitable value for *N*. The -Knounroll option invalidates the -Kunroll option. -Knounroll is set by default when the -O1 option is set. -Kunroll is set by default when the -O2 or higher option is set.

The -Kunroll and -Knounroll options require that the -O1 option or higher is set.

For details about loop unrolling, see Section "[3.2.4 Loop Unrolling](#)".

-K{unroll\_and\_jam[=*N*]|nounroll\_and\_jam}

-Kunroll\_and\_jam specifies to apply unroll-and-jam. The -Knounroll\_and\_jam option invalidates the -Kunroll\_and\_jam option. However, the optimization is suppressed in the following case:

- Assuming that the optimization is not effective.
- Assuming that there is a data dependency over iterations of the loops.

*N* is the upper limit of the unrolling expansion number which should be an integer value from 2 to 100. If =*N* is omitted, the compiler automatically determines a value for *N*. -Knounroll\_and\_jam option is set by default.

The -Kunroll\_and\_jam and -Knounroll\_and\_jam options require that the -O2 option or higher is set.

Unroll-and-jam promotes removing common expression and improves the execution performance. However, the increase of data stream and change of the order of data access may decrease the cache efficiency and the execution performance.

This optimization is not expected to be effective for all the loops in the program, so it is recommended to use the optimization control specifier `unroll_and_jam` or `unroll_and_jam_force` for each loop rather than the `-Kunroll_and_jam` option for the entire program.

For details about `unroll-and-jam`, see Section "3.3.8 Unroll-and-Jam".

#### -Kvisimpact

Specifies to generate an optimized object for VISIMPACT (Virtual Single Processor by Integrated Multicore Parallel Architecture).

This option is equivalent to the following options:

`-Kfast,parallel`

This option must be set at both compilation and linking.

See Section "3.6.1 Side Effect of Optimizations for Floating-Point Operation" for the side effect of optimization.

#### -K{zfill[=*N*]|nozfill}

The `-Kzfill` option specifies to perform `zfill` optimization. The `zfill` optimization speeds up write operations for array data that is only written in a loop, by using an instruction that allocates space on the cache for writing (DC ZVA) without loading data from the memory. The `zfill` optimization works on the data `N`blocks ahead of the address pointed to by the target store instruction where one block is 256 byte-long and `N`is an integer value between 1 and 100. If a value is not specified for `N`, the compiler will automatically determine a value. The `-Knozfill` option invalidates the `-Kzfill` option. `-Knozfill` is always applied when the `-O1` option is effective. `-Knozfill` is set by default when the `-O2` option or higher is set.

The `-Kzfill` option requires that the `-KA64FX` option and the `-O2` option or higher are set.

Note that if an object program compiled with the `-Kzfill` option is executed on a CPU other than the one on which a single cache write operation of a DC ZVA instruction is 256 bytes, the execution may be terminated abnormally or an incorrect result may occur. The amount of the cache write operation for a DC ZVA instruction on an A64FX processor is 256 bytes.

Performance may also be reduced under the following conditions:

- The program is not affected by memory bandwidth bottleneck.
- When the loop iteration count is too few.
- When the block size is explicitly specified with the `-Kzfill=N` option and the memory size where the data is written in the loop is smaller than `N`blocks.

This optimization is not expected to be effective for all the loops in the program, so it is recommended to use the optimization control specifier `zfill` for each loop rather than the `-Kzfill` option for the entire program.

For details about `zfill`, see Section "3.3.5 zfill".

## 2.2.2.7 -N Option

#### -N{Rtrap|Rnotrap}

`-NRtrap` specifies to catch signals. The `-NRnotrap` option invalidates the `-NRtrap` option. `-NRnotrap` is set by default. When the `-NRtrap` option is specified, the following signals are caught and the messages are output.

Signal	Message
SIGILL, SIGBUS, SIGSEGV	jwe0019i-u
SIGXCPU	jwe0017i-u
SIGFPE	jwe1017i-u

The floating-point underflow exception is caught when the environment variable `FLIB_EXCEPT=u` is specified.

Note that the integer divide by zero is not detected in the Fujitsu CPU A64FX with Arm architecture. See Section "B.2.10 Integer Division Exception when the Divisor Is Zero" for details.

For details about each signal, see Section "8.2.1 Causes of Abend".

If the `-NRtrap` option is used together with `-Kpreex` option and `-Ksimd=2` option, the speculative execution may cause exceptions that would not normally occur.

When the compiler options `-Kfp_relaxed` and `-NRtrap` are effective and either of the compiler option `-Knosimd` or `-KNOSVE` is effective, a floating-point exception may occur. However, to avoid a floating-point divide-by-zero exception for `sqrt(0.0)`, the optimization to convert to reciprocal approximation instructions is suppressed for the `sqrt` function. Therefore, the execution performance may decrease as compared with when the compiler option `-NRnotrap` is effective.

To enable the `-NRtrap` option, it must be set at both compilation and linking.

#### `-N{cancel_overtime_compilation|nocancel_overtime_compilation}`

These options specify whether to cancel the compilation if the compiler forecasts that it takes a long time (24 hours or more as a guide) to compile the program. When the `-Ncancel_overtime_compilation` option is specified, the compilation is canceled. The `-Ncancel_overtime_compilation` is set by default.

#### `-Ncheck_cache_arraysize`

The `-Ncheck_cache_arraysize` option inspects the sizes of arrays during the compilation and issue a warning message if the sizes could be a cause of an impact to execution performance.

This function decides that an array could be a cause of a cache conflict if its size is a multiple of that of second level cache. The array can be rewritten so as not to have a size of multiple of that of the cache in order to avoid the impact.

This function cannot judge correctly when the Sector cache is active.

#### `-N{clang|noclang}`

`-Nclang` option specifies that the objects are generated using Clang Mode. `-Nnoclang` option invalidates `-Nclang` option. `-Nnoclang` is set by default.

For details about Clang Mode, see "[Chapter 9 Clang Mode](#)".

#### `-N{coverage|nocoverage}`

The `-Ncoverage` option specifies to generate the information for the code coverage. `-Ncoverage` option should be specified when programs are compiled and linked. `-Nnocoverage` is set by default.

If the `-Ncoverage` option specifies, the execution performance may decrease due to the instructions that measures the execution time are added in the object program.

For details about the code coverage, see section "[Appendix E Code Coverage](#)".

#### `-N{exceptions|noexceptions}`

The `-Nexceptions` option specifies to enable exception handling. Also, `-Nexceptions` defines the `__EXCEPTIONS` macro. The `-Nnoexceptions` option invalidates the `-Nexceptions` option. `-Nexceptions` is set by default.

The `-Nexceptions` option is necessary when a program handles exception.

If `-Nnoexceptions` is specified and the program uses exception handling syntax, compilation is terminated.

#### `-N{fjcex|nofjcex}`

`-Nfjcex` specifies to use the Fujitsu Extended Functions in this system. The `-Nnofjcex` option invalidates the `-Nfjcex` option. `-Nnofjcex` is set by default.

The `-Nfjcex` option must be set at both compilation and linking.

For details about the extended functions, see "[Appendix F Fujitsu Extended Functions](#)".

#### `-N{fjprof|nofjprof}`

`-Nfjprof` option specifies to enable the Profiler function. The `-Nnofjprof` option invalidates the `-Nfjprof` option. These options must be set at linking. The default is `-Nfjprof`.

See the "Profiler User's Guide" for information on the Profiler.

#### `-Nfmtl=fmtl_mode`

Specifies to use Fujitsu Matrix Template Library (hereinafter called "FMTL"). For *fmtl\_mode*, specify "serial", "SSL2", "parallel", or "SSL2BLAMP".

This option should to be specified at both compilation and linking.

FMTL cannot be used in the parallel region. For more detail, see "[3.3.12 Fujitsu Matrix Template Library](#)".



**-Nfmtl=serial**

Specifies to use FMTL (serial execution version). When this option is specified, predefined macro "`__FMTL_SERIAL`" is defined.

**-Nfmtl=SSL2**

Specifies to use FMTL (serial high speed execution version using SSL2 library). When this option is specified, predefined macro "`__FMTL_SSL2`" is defined.

**-Nfmtl=parallel**

Specifies to use FMTL (thread parallel execution version). When this option is specified, predefined macro "`__FMTL_PARALLEL`" is defined. This library is executed in parallel in accordance with environment variables for OpenMP specifications. For more detail of environment variables for OpenMP specifications, see "[4.3.2.2 Environment Variable for OpenMP Specifications](#)".

**-Nfmtl=SSL2BLAMP**

Specifies to use FMTL (high speed thread parallel execution version using SSL2 library). When this option is specified, predefined macro "`__FMTL_SSL2BLAMP`" is defined. This library is executed in parallel in accordance with environment variables for OpenMP specifications. For more detail of environment variables for OpenMP specifications, see "[4.3.2.2 Environment Variable for OpenMP Specifications](#)".

**-N{hook\_func|nohook\_func}**

`-Nhook_func` specifies to use the hook function that is called from a specified location. The `-Nnohook_func` option invalidates the `-Nhook_func` option. `-Nnohook_func` is set by default.

If the `-Nhook_func` option is set, the user-defined function is called from the following locations:

- Program entry and exit
- Function entry and exit
- Parallel region (OpenMP or automatic parallelization) entry and exit

This option must be set when compiling and linking programs.

See Section "[8.3 Hook Function](#)", for information about the hook function.

**-N{hook\_time|nohook\_time}**

`-Nhook_time` specifies to use a hook function that is called at regular time interval. The `-Nnohook_time` option invalidates the `-Nhook_time` option. `-Nnohook_time` is set by default.

If the `-Nhook_time` option is set, the user-defined function is called at regular time interval.

The interval time to call the user-defined function is specified by environment variable `FLIB_HOOK_TIME`. If environment variable `FLIB_HOOK_TIME` is not specified, the user-defined function is called at the interval of a minute.

For details about the environment variable `FLIB_HOOK_TIME`, see Section "[2.5.1 Environment Variable for Execution](#)".

This option must be set at linking.

See Section "[8.3 Hook Function](#)", for information about the hook function.

**-N{libomp|fjomplib}**

Specifies the library for multiprocessing. `-Nlibomp` option is set by default. If `-Nclang` option is set, `-Nfjomplib` option is disabled and `-Nlibomp` option is enabled. These options must be set at linking.

**-Nlibomp**

Specifies to use LLVM OpenMP Library for multiprocessing. See "[Chapter 4 Multiprocessing](#)" for LLVM OpenMP Library.

**-Nfjomplib**

Specifies to use Fujitsu OpenMP Library for multiprocessing. See "[Appendix J Fujitsu OpenMP Library](#)" for Fujitsu OpenMP Library.

**-N{line|noline}**

`-Nline` generates additional information required to use the execution time sampling function provided with the Profiler. The `-Nnoline` option invalidates the `-Nline` option.

-Nline is set by default.

**-Nlst[={p|t}]**

Specifies to output compilation information to file(s) whose name is with suffix ".lst". When more than one C++ source files are specified, the file names are "each-file-name.lst".

When the -Nlst option is specified without arguments, -Nlst=p is set.

For details on the output format on statements specified with the OpenMP C++ directives, see ["Chapter 4 Multiprocessing"](#).

**-Nlst=p**

Specifies to write compilation information that consists of source list and statistics list.

The source list includes annotation of optimization which shows the situation of automatic parallelization, inline expansion, loop unrolling, etc.

**-Nlst=t**

Specifies to write detail optimization information to source list in addition to the function of -Nlst=p.

**-Nlst\_out=file**

Specifies the filename to output compilation information.

When this option is specified, the -Nlst=p option is also effective.

**-Nprofile\_dir=dir\_name**

Specifies storage location directory of the .gcca file which is necessary for the use of the code coverage. For *dir\_name*, the storage location directory name is specified by the relative path or the absolute path.

The .gcca file is generated when the executable program linked with object programs containing information for coverage is executed. If the directory specified as *dir\_name* does not exist at the execution time, the directory is generated.

This option requires that the -Ncoverage option and the -S or -c option are set.

For details about the code coverage, see section ["Appendix E Code Coverage"](#).

**-Nquickdbg={subchk|nosubchk|heapchk|noheapchk|inf\_detail|inf\_simple}}**

The -Nquickdbg option is used to debug C++ programs. The debugging function includes embedded information for debugging to a C++ object file, and checks are performed automatically during execution. Checks are performed if the -Nquickdbg=subchk or -Nquickdbg=heapchk option is set.

Multiple debugging functions can be used in combination by specifying multiple -Nquickdbg values. The -Nquickdbg=inf\_detail and -Nquickdbg=inf\_simple options are enabled with at least one of the -Nquickdbg, -Nquickdbg=subchk, or -Nquickdbg=heapchk options.

Omission of an argument is equivalent to specifying -Nquickdbg=subchk,quickdbg=heapchk.

If the -Nquickdbg=subchk or -Nquickdbg=heapchk option is enabled, the -Nquickdbg=inf\_detail option is also enabled.

This option must be set when compiling and linking programs.

See Section ["8.1 Functions for Debugging"](#) for details about this function.

**-Nquickdbg={subchk|nosubchk}**

Specifies whether to check the validity of the referenced subscript range against the declared array size when arrays are referenced. See Section ["8.1.1 Subscript Range Checks \(subchk Function\)"](#) for details.

**-Nquickdbg=subchk**

The subscript range is checked. If an error is detected, a diagnostic message is output.

**-Nquickdbg=nosubchk**

The subscript range is not checked.

**-Nquickdbg={heapchk|noheapchk}**

Specifies whether to check the heap for improper use of memory, buffer overrun, and memory leaks. See Section ["8.1.2 Heap Memory Checks \(heapchk Function\)"](#) for details.

`-Nquickdbg=heapchk`

The heap is checked. If an error is detected, a diagnostic message is output.

If allocation and release of a heap memory exist in different compilation units, the `-Nquickdbg=heapchk` option must be specified for both compilation units.

`-Nquickdbg=noheapchk`

The heap is not checked.

`-Nquickdbg={inf_detail|inf_simple}`

Specifies the information to be included in diagnostic messages output when errors are detected. See Section "[8.1 Functions for Debugging](#)" for details.

`-Nquickdbg=inf_detail`

In addition to the message and line number where the error occurred, information to help identify the cause is output, such as the variable name where the error occurred.

`-Nquickdbg=inf_simple`

The message indicating the error and the line number where the error occurred is output in diagnostic messages.

`-N{reordered_variable_stack|noreordered_variable_stack}`

The `-Nreordered_variable_stack` option directs the compiler the order in which to allocate the automatic variables to the stack area.

If the `-Nreordered_variable_stack` option is set, the allocation order of automatic variables is determined in the following order.

1. Descending order of their alignments
2. Descending order of data sizes if the alignments are equal
3. The order of appearance of the declaration statements in the source program if both the alignments and data sizes are equal

The stack area of the entire program can be reduced by allocating the automatic variables in descending order of their alignments.

When the `-Nnoreordered_variable_stack` option is specified, automatic variables are allocated in order of the appearance of the declaration statements in the source program. `-Nnoreordered_variable_stack` is set by default.

The order of allocation is not guaranteed when the `-Nnoline`, `-g0`, or `-Kocl` option is set.

For details, see Section "[D.4 Data Allocation to Stack Region](#)".

`-N{rt_tune|rt_notune}`

`-Nrt_tune` specifies to output the runtime information. The `-Nrt_notune` option invalidates the `-Nrt_tune` option. `-Nrt_notune` is set by default.

The runtime information can be used for the tuning of a user program. Refer to the "[Appendix G Runtime Information Output Function](#)" for details.

`-Nrt_tune_func`

In addition to the `-Nrt_tune` output, runtime information about user-defined functions is output.

If the `-Nrt_tune_func` option is set, the `-Nrt_tune` option also takes effect. When the `-Nrt_notune` option is specified after this option, runtime information is not output.

Refer to the "[Appendix G Runtime Information Output Function](#)" for details.

`-Nrt_tune_loop[={all|innermost}]`

In addition to the `-Nrt_tune` output, runtime information about loops is output. If the argument `={all|innermost}` is omitted, `-Nrt_tune_loop=all` is assumed.

If the `-Nrt_tune_loop` option is set, the `-Nrt_tune` option also takes effect. When the `-Nrt_notune` option is specified after this option, runtime information is not output.

Refer to the "[Appendix G Runtime Information Output Function](#)" for details.

`-Nrt_tune_loop=all`

Runtime information about all loops is output.

`-Nrt_tune_loop=innermost`

Runtime information about the innermost loops is output.

`-N{setvalue[=setarg]|nosetvalue}`

`-Nsetvalue` specifies to automatically fill up data allocated in heap or stack area with 0 value. `-Nnosetvalue` is set by default.

Note that execution time may increase by the `-Nsetvalue` option.

For *setarg*, specify one of the following:

{ {heap|noheap} | {stack|nostack} | {scalar|noscalar} | {array|noarray} | {struct|nostruct} }

If the argument `=setarg` is omitted, `-Nsetvalue=heap,setvalue=stack` is assumed.

The `-Nsetvalue` option is invalidated if the `-Nquickdbg` option is set.

`-Nsetvalue={heap|noheap}`

`-Nsetvalue=heap` specifies to automatically fill up data allocated in heap area with 0 value. `-Nsetvalue=noheap` is set by default.

`-Nsetvalue=heap` fills up following area with 0 value:

- areas allocated by malloc function
- areas allocated by new expression (depends on the new operator that is the default version defined by the C++ Standard library excluding placement forms.)

`-Nsetvalue={stack|nostack}`

`-Nsetvalue=stack` specifies to automatically fill up uninitialized variables allocated in stack area with 0 value. `-Nsetvalue=nostack` is set by default.

`-Nsetvalue=stack` fills up following variables with 0 value:

- automatic variables

However, variables specified by private clause or lastprivate clause of OpenMP specification are not filled up with 0 value, even if the `-Nsetvalue=stack` option is effective.

The `-Nsetvalue=stack` option is equivalent to the following:

`-Nsetvalue=stack,setvalue=scalar,setvalue=array,setvalue=struct`

`-Nsetvalue={scalar|noscalar}`

`-Nsetvalue=scalar` specifies to automatically fill up uninitialized scalar type variables allocated in stack area with 0 value. `-Nsetvalue=noscalar` is set by default.

`-Nsetvalue=scalar` option requires that the `-Nsetvalue=stack` option is set.

`-Nsetvalue={array|noarray}`

`-Nsetvalue=array` specifies to automatically fill up uninitialized array type variables allocated in stack area with 0 value. `-Nsetvalue=noarray` is set by default.

`-Nsetvalue=array` option requires that the `-Nsetvalue=stack` option is set.

`-Nsetvalue={struct|nostruct}`

`-Nsetvalue=struct` specifies to automatically fill up uninitialized class, struct, and union type variables allocated in stack area with 0 value. `-Nsetvalue=nostruct` is set by default.

`-Nsetvalue=struct` option requires that the `-Nsetvalue=stack` option is set.



## Example

How to specify `-Nsetvalue` option

To fill up only uninitialized scalar type variables with 0 value allocated in stack area, specify as follows:

```
-Nsetvalue=stack,setvalue=noarray,setvalue=nostruct
```

## -Nsrc

Specifies to output source list to the standard output.

The source list includes annotation of optimization which shows the situation of automatic parallelization, inline expansion, loop unrolling, etc.

Note that when the -Nlst or -Nlst\_out=*file* option is specified, the source list is output to the *file*.

For details on the output format on statements specified with the OpenMP C++ directives, see "[Chapter 4 Multiprocessing](#)".

## -Nsta

Specifies to output statistics list to the standard output.

Note that when the -Nlst or -Nlst\_out=*file* option is specified, the statistics list is output to the *file*.

## 2.2.3 Notes for Using Compiler Options

---

This section provides notes on compiler options.

### "-K" option and "-N" option

Multiple options can be specified after the "-K" option using commas to separate them, or they can be specified in two or more separate "-K" options. "-N" option works in the same way.

In the following example, although the options are specified differently, the two commands are equivalent.

Example:

```
$ FCCpx -o pgm.out -Kfast -Kparallel -Nsrc -Nsta pgm.cc
```

```
$ FCCpx -o pgm.out -Kfast,parallel -Nsrc,sta pgm.cc
```

### Ban of Blank Space

Options which are specified after the "-K" and "-N" options cannot have any blank spaces.

In the following example, the "parallel" option is invalid because there is a blank space in front of the option.

Example:

Invalid Specification:

```
$ FCCpx -Kfast, parallel pgm.cc
```

Correct Specification:

```
$ FCCpx -Kfast,parallel pgm.cc
```

## 2.3 Environment Variable for Compile Command

---

This section explains the environment variables that the compile command recognizes.

FCCpx\_ENV

FCC\_ENV

Environment variables to set compiler options. These are valid in Trad/Clang Mode.

The environment variable FCCpx\_ENV is for the cross compiler and the environment variable FCC\_ENV is for the native compiler.

The user can specify Trad/Clang Mode common compiler options for the value of these environment variables. Refer to "[9.1.2.3.1 Correspondences of Compile Options in Clang Mode and Trad Mode](#)" for a description of the common compiler options for Trad/Clang Mode.

Compiler options specified to these environment variables are valid in Trad/Clang Mode. It is useful to set compiler options that you always specify for these environment variables.

Refer to "[2.4 Compilation Profile File](#)" for the priority of compiler options.

## Example

The following two examples are equivalent:

Example 1: Specifying compiler options to environment variable `FCCpx_ENV`

```
$ FCCpx_ENV="-Kfast -I/usr/prv/include"
$ export FCCpx_ENV
$ FCCpx a.cc
```

Example 2: Specifying compiler options to operands of the compile command

```
$ FCCpx -Kfast -I/usr/prv/include a.cc
```

`FCCpx_trad_ENV`

`FCC_trad_ENV`

Environment variables to set compiler options. These are valid only in Trad Mode.

The environment variable `FCCpx_trad_ENV` is for the cross compiler and the environment variable `FCC_trad_ENV` is for the native compiler.

The user can specify Trad Mode specific and Trad/Clang Mode common compiler options for the value of these environment variables. Refer to "[9.1.2.3.1 Correspondences of Compile Options in Clang Mode and Trad Mode](#)" for a description of the common compiler options for Trad/Clang Mode.

Compiler options specified to these environment variables are valid only in Trad Mode.

Refer to "[2.4 Compilation Profile File](#)" for the priority of compiler options.

`FCOMP_LINK_FJOB`

In this system, original objects of this system are usually combined at linking, but they are not combined under the following conditions.

- The `-L` option which is passed to the linker by this system is directly specified on the compile command to pass the linker by user.

This can result in a link error (undefined reference to), which can be avoided by setting the environment variable `FCOMP_LINK_FJOB`.

These are valid in Trad/Clang Mode. The environment variable `FCOMP_LINK_FJOB` can have any value. When the environment variable `FCOMP_LINK_FJOB` is set, original objects of this system are combined at linking.

For the example of a link error, see "[B.2.12 Link Error \(undefined reference to\)](#)".

## Example

Example of using `FCOMP_LINK_FJOB` environment variable:

```
$ export FCOMP_LINK_FJOB=true
```

`FCOMP_UNRECOGNIZED_OPTION`

The behavior for unrecognized compiler options can be changed by specifying the environment variable `FCOMP_UNRECOGNIZED_OPTION`. This is valid only in Trad Mode.

Specify either `warning` or `error` to the environment variable `FCOMP_UNRECOGNIZED_OPTION`. If the environment variable `FCOMP_UNRECOGNIZED_OPTION` is not set or an invalid value is set to it, `warning` is set.

Value	Description
<code>warning</code>	A warning message is output for an unrecognized compiler option and the compilation is continued.
<code>error</code>	An error message is output for an unrecognized compiler option and the compilation is canceled.

## Example

Example 1: Specifying warning to the environment variable `FCOMP_UNRECOGNIZED_OPTION`

```
$ FCOMP_UNRECOGNIZED_OPTION=warning
$ export FCOMP_UNRECOGNIZED_OPTION
$ FCCpx -unrecognized_option a.cc
FCCpx: warning: -unrecognized_option is unrecognized option.
$ echo $?
0
$
```

Example 2: Specifying error to the environment variable `FCOMP_UNRECOGNIZED_OPTION`

```
$ FCOMP_UNRECOGNIZED_OPTION=error
$ export FCOMP_UNRECOGNIZED_OPTION
$ FCCpx -unrecognized_option a.cc
FCCpx: error: -unrecognized_option is unrecognized option.
$ echo $?
1
$
```

`CPATH`

`CPLUS_INCLUDE_PATH`

The directory that retrieves headers with names that do not start with '/' can be added by specifying a value for the environment variable `CPATH` or `CPLUS_INCLUDE_PATH`. These are valid in Trad/Clang Mode.

To specify directory and other directories, separate the arguments by a colon (:).

For details about the order of the search of the header, see the `-I` option in Section "2.2.2.1 General Options for Compiler".

## Example

Example of using `CPATH` and `CPLUS_INCLUDE_PATH` environment variables:

```
$ CPATH="/cpath_inc:/cpath_inc_2"
$ export CPATH
$ CPLUS_INCLUDE_PATH="/cplus_include_path:/cplus_include_path_2"
$ export CPLUS_INCLUDE_PATH
```

`LIBRARY_PATH`

The directory to be searched for libraries at linking can be added by specifying a value for the environment variable `LIBRARY_PATH`. This is valid in Trad/Clang Mode.

Multiple directories can be specified using colons to separate them.

Libraries are searched in the following order:

1. Directories specified as the argument of the `-L` option
2. The directory installed libraries provided by the compiler
3. The directory installed standard libraries
4. Directories specified to the environment variable `LIBRARY_PATH`

## Example

Example of using `LIBRARY_PATH` environment variable:

```
$ LIBRARY_PATH="/usr/local/lib64:/usr/local_2/lib64"
$ export LIBRARY_PATH
```

### TMPDIR

The temporary directory used by the compile command can be changed by specifying a value for the environment variable `TMPDIR`. This is valid in Trad/Clang Mode.

If the environment variable `TMPDIR` is not set, `/tmp` is used. To avoid output to a common directory, set the local directory to the environment variable `TMPDIR`.

## Example

Example of using `TMPDIR` variable:

```
$ TMPDIR=/usr/local/tmp
$ export TMPDIR
```

## 2.4 Compilation Profile File

The default values of compiler options can be changed by specifying the compilation profile file.

The following compilation profile files are available:

Table 2.4 Compilation Profile files (Trad Mode)

Kind	Mode	File Name
Cross compiler	Trad/Clang Mode common	/etc/opt/FJSVstclang/FCCpx_PROF
	Trad Mode specific	/etc/opt/FJSVstclang/FCCpx_trad_PROF
Native compiler	Trad/Clang Mode common	/etc/opt/FJSVstclang/FCC_PROF
	Trad Mode specific	/etc/opt/FJSVstclang/FCC_trad_PROF

Common compilation profile files for Trad/Clang Mode can specify common compiler options for Trad/Clang Mode. Refer to "[9.1.2.3.1 Correspondences of Compile Options in Clang Mode and Trad Mode](#)" for a description of the common compiler options for Trad/Clang Mode.

Trad Mode specific compilation profile files can specify Trad Mode specific and Trad/Clang common compiler options.

## Note

For contents about the settings in compilation profile files, contact the system administrator.

The format of a compilation profile file is as follows:

- Blank characters not enclosed in quotation marks (") or single quotation marks (') are not significant.
- Either single or double quotation marks can be used to begin and end character strings. A character string is terminated by the same delimiter used to begin the string. A character string is also terminated by the end of a line.
- A symbol (#) that is not part of a character string starts a comment. Any characters from the # to the end of the line are treated as part of the comment.



## Example

Example of a compilation profile file specification

```
#Default options
-w -O
```

Compiler options are determined by the following priorities:

1. Operand of the compile command
2. Environment variable to set compiler options (mode specific)
3. Environment variable to set compiler options (mode common)
4. Compilation profile file (mode specific)
5. Compilation profile file (mode common)
6. Default value

## 2.5 Procedure of Execution


This section explains the procedure of executing C++ language programs.

### 2.5.1 Environment Variable for Execution

The environment variables shown in "[Table 2.5 Environment Variables for Execution](#)" control the behavior of execution.

For details about OpenMP environment variables, see Section "[4.3.2.2 Environment Variable for OpenMP Specifications](#)".

Table 2.5 Environment Variables for Execution

Environment Variable Name	Operands	Description
FLIB_C_MESSAGE	NO_MESSAGE	<p>A message is output to standard error when an error occurs during execution of the user program. However, if this environment variable is set, an error message is not output.</p> <p>For details of the error message, refer to "Fortran/C/C++ Runtime Messages".</p> <p>When a program was compiled and linked with the -NRtrap option, this environment variable is invalidated.</p> <p> <b>Note</b></p> <p>.....</p> <p>This environment variable is invalidated for the error message of LLVM OpenMP Library. See "<a href="#">4.4 Runtime Messages</a>" for the message of LLVM OpenMP Library.</p> <p>.....</p>
FLIB_EXCEPT	u	If the -NRtrap option is set at compilation and linking, the floating-point underflow exception (jwe0012i-u) is caught.
FLIB_HEAPCHK_VALUE	<i>hex</i>	Changes the check values used in checks for buffer overrun. See Section " <a href="#">8.1.2.2 Buffer Overrun Check</a> " for details.
FLIB_HOOK_TIME	<i>time</i>	The user-defined function is called at interval of <i>time</i> millisecond(s). <i>time</i> can be a value between 0 and 2147483647. If 0 is specified for <i>time</i> , calling at regular interval is disabled. See Section " <a href="#">8.3 Hook Function</a> ", for information about the hook function.
FLIB_RTINFO	(none)	Outputs runtime information. For details, see " <a href="#">Appendix G Runtime Information Output Function</a> ".

Environment Variable Name	Operands	Description
FLIB_RTINFO_CSV	<i>file</i>	In the Runtime Information Output Function, outputs the fetched information to a CSV file.  <i>file</i> is optional. Any file name can be specified as <i>file</i> of the environment variable.  If <i>file</i> is omitted, <code>flib_rtinfo.csv</code> is assumed.  For details, see " <a href="#">Appendix G Runtime Information Output Function</a> ".
FLIB_RTINFO_NOFUNC	(none)	In the Runtime Information Output Function, even if the compiler option -Nrt_tune_func is effective, the output of information for each function is suppressed.  For details, see " <a href="#">Appendix G Runtime Information Output Function</a> ".
FLIB_RTINFO_NOLOOP	(none)	In the Runtime Information Output Function, even if the compiler option -Nrt_tune_loop is effective, the output of information for each loop is suppressed.  For details, see " <a href="#">Appendix G Runtime Information Output Function</a> ".
FLIB_TRACEBACK_MEM_SIZE	<i>size</i>	This environment variable changes the size of the heap memory used to debug information output to the trace back map information.  An integer value from 1 to 128 can be specified as <i>size</i> . The unit for specifying <i>size</i> is MiB.  If no environment variable is specified, or if the value of <i>size</i> is invalid, the default value of this system will be used as the value of <i>size</i> . The default value is 16 MiB.  See " <a href="#">5.2.2 Trace Back Information</a> " for the trace back map.

## 2.5.2 Notes for Execution

When executing programs created on this system, note the following.

### 2.5.2.1 Variable Allocation at Execution

Programs created by this system allocate the local and private variables in each function to the stack region. When much more space is required for local and private variables in the function, the stack area needs to be extended to the appropriate size.

The stack area of the process can be set by `ulimit` (bash built-in command), etc.

# Chapter 3 Optimization

This chapter describes the optimization functions and how to use them.

## 3.1 Overview of Optimization

Optimization aims to generate object modules (instructions and data areas) that allow the program to execute as quickly as possible. The compiler optimization functions perform the following tasks while generating object modules from C++ source code.

### Deletion

If there are statements or expressions that give the same result, any redundant portions are deleted. Statements or expressions that would be meaningless if executed are also deleted.

### Movement

Statements or expressions within loops that give the same result no matter how many times they are repeated are moved outside the loop. When the ultimate result is known without repeated execution, only the required portions are left inside the loop, with the remaining portion moved outside the loop.

### Modification

The operator of an arithmetic expression may be changed within a range in which no calculation errors occur. In addition, the data subject to operation may also be modified. For example, for a variable to which a constant value is assigned once and its value remains thereafter unchanged, the reference to that variable is changed to a reference to the constant.

### Expansion

When standard library functions or user-defined functions are called, whenever possible, the body of the function is expanded in the location where it was called. This obviates the need to pass arguments and function values or use instructions for calling and return. Moreover, the function's code is optimized as a single unit with the calling function, decreasing execution time.

### Execution

If the result of the execution of a statement or the calculated result of an expression is known at compilation time, the instructions that execute the statement or calculate the expression are omitted and only the result of the execution or calculation is used at compilation time.

### Other

The functions and characteristics of hardware (for example type of instructions, type and number of registers, addressing methods) are employed as much as possible in order to generate instructions and data areas for fast execution.

Optimization functions can shorten execution time. However, compilation time and work area increase. In addition, depending on the type of optimization, the object modules may be greatly enlarged, and some types of optimization may cause unexpected execution results. To perform optimization, select the appropriate compiler options.

Standard optimization specifies the optimization level and extended optimization supplements standard optimization. These optimization levels are applied by specifying compilation options.



### Point

By default, optimization does not cause unexpected results. This is because optimizations with the potential to affect the execution result must be specified by the user via the appropriate optimization option, and are not executed by default.

## 3.2 Standard Optimization

When an optimization level between level 1 (-O1 option) and level 3 (-O3 option) is specified, standard optimization is performed. This section describes standard optimization.

## 3.2.1 Elimination of Common Expressions

---

If two expressions that give equal calculation results (common expressions) are present, the result of the former expression can be used in the latter expression without calculation.

Elimination of common expressions applies to arithmetic operations, relational operations, logical operations, and some standard library functions.

Elimination of common expressions is performed in optimization level 1 (-O1 option) or higher.

An example of elimination of common expressions is given below.

Example: Elimination of Common Expressions

<pre>a = x * y + c; ... b = x * y + d;</pre>	->	<pre>t = x * y; a = t + c; ... b = t + d;</pre>
--	----	---

The  $x*y$  portion of the expression on the right hand side of the assignment is a common expression. The code is changed so that the second calculation is eliminated and the result of the first calculation  $t$  is used instead.  $t$  is a variable generated by the compiler.

## 3.2.2 Movement of Invariant Expressions

---

Expressions with values that do not change within a loop (invariant expressions) are moved outside the loop.

Movement of invariant expressions applies to arithmetic operations, relational operations, logical operations, and some standard library functions.

Movement of invariant expressions is performed in optimization level 1 (-O1 option) or higher.

An example of the movement of invariant expressions is given below.

Example: Movement of Invariant Expressions

<pre>for(i = 0; i &lt; n; i++) {     y = a[j] * 2;     x = x + y * z; }</pre>	->	<pre>y = a[j] * 2; t = y * z; for(i = 0; i &lt; n; i++) {     x = x + t; }</pre>
---	----	--

The entire statement  $y=a[j]*2$  and the  $y*z$  portion of the expression are moved outside the loop.  $t$  is a variable generated by the compiler.

The objects that are usually moved by this optimization are statements or parts of statements that are always executed within the loop. However, if the `-Kpreex` option is specified, invariant expressions in statements that are selectively executed within the loop depending on a decision statement are also moved. To differentiate this optimization from the normal movement of invariant expressions, it is called advance evaluation of invariant expressions. This optimization can further reduce execution time. However, side effects may occur due to the movement. For details of the side effects, see Section "[3.3.1.1 Advance Evaluation of Invariant Expressions](#)".

## 3.2.3 Reducing Strength of Operators

---

The "strength" of operators describes the relative amount of execution time required for operation. A "strong" operator requires a large amount of execution time. Typically, addition and subtraction have the same strength. Multiplication is stronger than addition or subtraction and division is even stronger than multiplication. In addition, type conversion between the integer and float types is stronger than addition or subtraction but weaker than multiplication. You can reduce execution time by changing strong operators to weaker ones; for example, by changing multiplication to addition or subtraction.

Reducing the strength of operators is performed in optimization level 1 (-O1 option) or higher.

## Information

### Induction Variable

This optimization applies to variables of type integer with values that change regularly and incrementally within a loop. This type of variable is called an induction variable.

An example of reducing strength is given below.

#### Example 1: Reducing the Strength of Operators (Reducing Multiplication to Addition)

```
for(i = 1; i < 10; i++) {  
    ...  
    j = k + i * 100;  
    ...  
}  
->  
t = 100;  
for(i = 1; i < 10; i++) {  
    ...  
    j = k + t;  
    ...  
    t = t + 100;  
}
```

Optimization is performed on the operation  $i * 100$  on the derivative variable  $i$ , so the operation  $i * 100$  within the loop is converted to the addition  $t = t + 100$ .  $t$  is a variable generated by the compiler.

#### Example 2: Reducing the Strength of Operators (Reducing Type Conversion to Addition)

```
for(i = 1; i < 10; i++) {  
    ...  
    x = (float)i;  
    ...  
}  
->  
t = 1.0;  
for(i = 1; i < 10; i++) {  
    ...  
    x = t;  
    ...  
    t = t + 1.0;  
}
```

Optimization is performed on the operation  $(float) i$  on the derivative variable  $i$ , so the operation of type conversion from integer to float is converted to the addition of float type numbers  $t = t + 1.0$ .  $t$  is a `float` type variable generated by the compiler.

## 3.2.4 Loop Unrolling

Loop unrolling is a type of optimization in which all of the statements executed within a loop are expanded  $N$  times within the loop (where  $N$  is called the "multiplicity"), and instead, the number of iterations of the loop is reduced to  $1/N$ . However, if the `Ksimd[={ 1|2|auto}]` option is set and the loop has been optimized using SIMD Extended instructions, the executable statement is unrolled "`SIMD Width`"\* $N$ , and the loop iteration count is reduced to "`SIMD Width`"\* $N$ .

This optimization is performed on loops which contain no jump from the inside to the outside and no jump from the outside to the inside.

The multiplicity  $N$  is an optimal value determined by the compiler depending on the number of iterations of the loop, the number and type of statements executed within the loop and the data types used. In addition, the multiplicity can be specified in the source code with the optimization control line (OCL). For details on this function, see Section "3.4.1 Using Optimization Control Line (OCL)".

Since the number of iterations is reduced and the executed statements that are expanded are optimized as a group within the loop, a faster object is obtained. However, since the statements executed within the loops are expanded, the size of the object module increases.

Loop unrolling works before and after SIMD.

- To apply SIMD or the software pipelining on the outer loop, the inner loop is fully unrolled before SIMD when the iteration count of the inner loop is small. This is called full unrolling before SIMD.
- The inner loop is unrolled or fully unrolled to promote the optimization such as common equations after SIMD.

Loop unrolling is performed in optimization level 2 (-O2 option) or higher.

Loop unrolling works before and after SIMD when `-Kunroll[=n]` option or the unroll specifier is the effective loop.

Loop unrolling can be suppressed by specifying `-Knounroll` option or the nounroll specifier.

You can control the behavior of full unrolling before SIMD by specifying the `fullunroll_pre_simd` or `nofullunroll_pre_simd` specifier. If the `fullunroll_pre_simd` or `nofullunroll_pre_simd` specifier is specified for the loop that has the `-Kunroll[=n]` option, `-Knounroll` option, `unroll` specifier, or `nounroll` specifier enabled, the `fullunroll_pre_simd` or `nofullunroll_pre_simd` specifier takes precedence in full unrolling before SIMD.

Moreover, the message to show the optimized loop and the multiplicity can be output by specifying the `-Koptmsg=2` option.

An example of loop unrolling is given below.

Example: Loop unrolling

```
int i, j, k;
float a[400][400], b[400][400], c[400][400];
for(i = 0; i < 400; i++) {
    for(j = 0; j < 400; j++) {
        c[i][j] = 0.0;
        for(k = 0; k < 400; k++) {
            c[i][j] += a[i][k] * b[k][j];
        }
    }
}
```

For example, if the multiplicity is 4, the innermost loop is converted to the following format in the process of generating the object:

```
for(k = 0; k < 400; k += 4) {
    c[i][j] += a[i][k ] * b[k ][j];
    c[i][j] += a[i][k+1] * b[k+1][j];
    c[i][j] += a[i][k+2] * b[k+2][j];
    c[i][j] += a[i][k+3] * b[k+3][j];
}
```

### 3.2.5 Loop Blocking

Loop blocking is the optimization to subdivide the access for array by multiplexing the loop in block size. As a result, the localization of data access is improved, and then efficient use of the cache is promoted.

This optimization is performed on loops which contain no jump from the inside to the outside and no jump from the outside to the inside.

Loop blocking is performed in optimization level 2 (`-O2` option). It can be suppressed by specifying the `-Kloop_noblocking` option.

The message to show the optimized loop can be output by specifying the `-Koptmsg=2` option.

An example of loop blocking is given below.

Example: Loop blocking

```
#define N 200
int a_array[N][N];
int b_array[N][N];

void sub() {
    int i,j;
    for(i = 0; i < N; i++) {
        for(j = 0; j < N; j++) {
            a_array[i][j] = b_array[j][i];
        }
    }
}
```

->

```
#define MIN(x,y) ((x<y)?x:y)
#define BLOCK 2
#define N 200
int a_array[N][N];
int b_array[N][N];

void sub() {
    int i,j,ii,jj;
    for(ii = 0; ii < N; ii += BLOCK) {
        for(jj = 0; jj < N; jj += BLOCK) {
            for(i = ii; i < MIN(N,(ii+BLOCK)); i++) {
                for(j = jj; j < MIN(N,(jj+BLOCK)); j++) {
                    a_array[i][j] = b_array[j][i];
                }
            }
        }
    }
}
```

## 3.2.6 Software Pipelining

Software pipelining schedules instructions in loops to execute as parallel as possible.

Software pipelining can be suppressed by specifying the compile-time option `-Knoswp`.

Software pipelining performs an instruction scheduling to arrange instructions of a particular iteration in a loop with instructions of the following iterations, and reshapes the loop. Thus software pipelining requires enough loop iteration count. Other optimizations affect the required iteration count because it depends on instructions in the loop.

If the original loop is short, software pipelining may be less effective because instruction scheduling is performed moderately to keep required iteration count small.

When software pipelining is applied to a loop whose iteration count is variable, as shown in the following example, the compiler inserts a branch instruction to choose the software-pipelined loop or the original loop in case the iteration count is short. It is decided at run time whether the software-pipelined loop is chosen. The size of the object program will increase.

Example:

```
void sub (int N) {
    (The original loop: the iteration count is N.)
}
```

[Optimized pseudo-code]

```
void sub (int N) {
    if (Is N enough?) {
        (The software-pipelined loop)
    } else {
        (The original loop: the iteration count is N.)
    }
}
```

If the compile-time option `-Koptmsg=2` is specified, optimization messages about results of software pipelining are output with the line numbers of loops. When software pipelining is applied to a loop, optimization messages `jwd8204o-i` and `jwd8205o-i` are output at once. When software pipelining is not applied to a loop, one of optimization messages from `jwd8662o-i` to `jwd8674o-i` is output. Optimization messages about software pipelining are not output when the `-Knoswp` option is specified, or when a loop disappears because of other optimizations before software pipelining, such as full unrolling.

The optimization message `jwd8205o-i` shows the minimum iteration count required to choose the software-pipelined loop. Make sure to consider following optimizations applied to the same loop because the target loop of software pipelining is the innermost loop; when loop collapse, loop interchange, thread parallelization, or inline expansion is applied, the iteration count of the target loop of software pipelining is changed as shown in the following table. If this is the case, the value shown in the optimization message `jwd8205o-i` should be compared with the iteration count in the table.

Optimization applied to the same loop	Message number	The iteration count of the target loop to which software pipelining is applied
Loop collapse	<code>jwd8330o-i</code>	Product of the iteration counts of the collapsed loops.
Loop interchange	<code>jwd8211o-i</code>	The iteration count of the innermost loop after loop interchange.
Thread parallelization	<code>jwd5001p-i</code> , etc.	Quotient of the iteration count divided by the number of threads.
Inline expansion	<code>jwd8101o-i</code>	The iteration count of the innermost loop after inline expansion.

As for the iteration count shown by optimization message `jwd8205o-i`, the value is output with taking account of the iteration counts executed at the same time by SIMD. Note that when the `-Ksimd_reg_size=agnostic` option is effective, the value is output as if the SVE vector register size is 128 bits is output, because the SVE vector register size is unknown at compilation.

## 3.2.7 SIMD

### 3.2.7.1 Normal SIMD

The compiler may apply optimizations using SIMD extensions. Using SIMD extensions, multiple operations of the same kind are executed at once.

This optimization is performed on loops which contain no jump from the inside to the outside and no jump from the outside to the inside.

This optimization is enabled by the -O2 option or higher. And this optimization is suppressed by the -Knosimd option. Moreover, the message to show the optimized loop can be output by specifying -Koptmsg=2 option.

### 3.2.7.2 SIMD Extension for Loop Containing "if" Statement

When -Ksimd={2|auto} option is specified, SIMD extension can be applied for the loop that contains "if" statement. The conditional execution instructions are only store instruction and move instruction in CPU architecture, and other instructions are executed speculatively (statements in the "if" statement are executed even if the condition of the "if" statement is false). Therefore the true rate of condition of the "if" statement is high, the performance may be improved. However, side effects, such as exceptions at execution, may occur in the execution results because instructions are executed speculatively.

### 3.2.7.3 List Vector Conversion

When the true rate of the "if" statement is low, and the "if" statement in loop contains many instructions, the performance may be improved by applying list vector conversion.

The list vector conversion generates the following two loops:

1. The loop to save the value of loop control variable into a new array when the condition of the original "if" statement is true.
2. The loop to execute the statements of the original "if" statement, and its loop iteration count is equal to the number of the new array element.

The loop of 2. becomes list access method, but the loop becomes the target of SIMD extension and software pipelining.

Example of list vector conversion is shown in the following. The list vector conversion is applied by specifying the optimization control line. For details about `simd_listv` specifier, see Section "3.4.1 Using Optimization Control Line (OCL)".

Example:

```
for (i=0; i<n; ++i) {  
#pragma statement simd_listv  
    if (m[i]) {  
        a[i] = b[i] + c[i];  
    }  
}
```

->


```
j = 0;  
for (i=0; i<n; ++i) {  
    if (m[i]) {  
        idx[j] = i;  
        j = j + 1;  
    }  
}  
for (k=0; k<j-1; ++k) {  
    i = idx[k];  
    a[i] = b[i] + c[i];  
}
```

### 3.2.7.4 SIMD with Redundant Executions for the SIMD Width

When the loop iteration count is not a multiple of the SIMD width, this function generates loops which use SIMD extensions over the whole iterations by using SIMD instructions with masks. (Refer to the example and figures below.) When SIMD is performed, this function is applied by default. The following figure is the image of this function.



Figure 3.1 SIMD with redundant executions for the SIMD width (4-wide SIMD)



[Explanation]

The whole iterations are executed by using SIMD instructions with masks.


The loops whose remainder dividing the loop iteration count by the SIMD width tend to be 1 or 2 may not obtain the effect of optimization. In order to suppress this function, specify the `simd_noredundant_vl` specifier in the optimization control line.

When the function is not applied to the source program in the example, SIMD instructions are used as shown in "Figure 3.2 The optimization control line is not effective (4-wide SIMD)". The remaining iterations, which correspond the remainder after dividing the loop iteration count by the SIMD width, do not use SIMD instructions.

Example: The optimization control line (OCL) is specified

```
#pragma loop simd_noredundant_vl
for(i=0; i<m; i++) {
    a[i] = b[i]+c[i];
}
```

Figure 3.2 The optimization control line is not effective (4-wide SIMD)



[Explanation]

When the loop control variable `i` is from 13 to 15, SIMD instructions are not used.

This function can be suppressed by specifying `-Kextract_stride_store` option.

### 3.2.7.5 Math functions that SIMD Extensions can be applied to

The following shows the math functions that SIMD extensions can be applied to when `-Klib` option is effective.

```
abs, acos, acosf, acosh, acoshf, asin, asinf, asinh, asinhf, atan, atan2, atan2f, atanf, atanh,
atanhf, cbrt, cbrtf, ceil, ceilf, cimag, cimagf, conj, conjf, copysign, copysignf, cos, cosf, cosh,
coshf, creal, crealf, erf, erfc, erfcf, erff, exp, exp2, exp2f, expf, fabs, fabsf, floor, floorf,
fmax, fmaxf, fmin, fminf, lgamma, lgammaf, log, log10, log10f, log2, log2f, logf, nearbyint,
nearbyintf, pow, powf, rint, rintf, round, roundf, sin, sinf, sinh, sinhf, sqrt, sqrtf, tan, tanf,
tanh, tanhf, tgamma, tgammaf, trunc, truncf
```

See also the description of the `-Kifunc` and `-Kmfunc` options in the section "[2.2 Compiler Options](#)" for information on math functions that SIMD extensions can be applied to.

### 3.2.8 Loop Unswitching

Loop unswitching is an optimization to modify loop by putting out the `"if"` statement out of the loop and creating a loop in each `"TRUE"` and `"FALSE"` blocks of the `"if"` statement in order to promote some optimizations such as instruction scheduling if the condition of the `"if"` statement is invariant in a loop.

This optimization is enabled by specifying the `-O3` or `-Kparallel` option.

This optimization is disabled by specifying the `-O2` option or less, and invalidating `-Kparallel` option.

```
void foo(double a[], int b, int c, int n) {
    int i;
    for(i = 0; i < n; i++) {
        if(b == c) {
            a[i] = 0;
        } else {
            a[i] = 1;
        }
    }
}

->

void foo(double a[], int b, int c, int n) {
    int i;
    if(b == c) {
        for(i = 0; i < n; i++) {
            a[i] = 0;
        }
    } else {
        for(i = 0; i < n; i++) {
            a[i] = 1;
        }
    }
}
```

### 3.2.9 Inline Expansion

Using inline expansion, when a function is referenced (called), the body of the function is expanded directly in the place where the function was called. Expansion eliminates the passing of arguments and function values and the overhead for calling and returning from the function (saving and restoring registers, branches). In addition, bringing the expanded portion together with the calling portion promotes other types of optimization. Therefore, execution time can be greatly reduced, but the size of the object module will increase.

Inline expansion is performed in optimization level 3 (`-O3` option). It can be suppressed by specifying the `-x0` option.

The message to show the name and its location of the expanded function can be output by specifying the `-Koptmsg=2` option.

## 3.3 Extended Optimization

When one of the extended function options is specified regardless of optimization, extended optimization is performed in addition to the standard optimization.

This section describes the functions comprising extended optimization.

Depending on the extended optimization function, the size of the object modules (instructions and data areas) generated by the compiler may be greatly enlarged or side effects may occur in the execution results.

### 3.3.1 Optimization by Modifying Evaluation Methods

This section describes optimization by advance evaluation of invariant expressions and arithmetic evaluation method modification.

#### 3.3.1.1 Advance Evaluation of Invariant Expressions

Invariant expressions in statements executed selectively within a loop (for example if they depend on an "if" statement or similar) are moved outside the loop. Advance evaluation of invariant expressions can reduce execution time, however may also cause side effects.

The advance evaluation of invariant expressions can be controlled with the `-Kpreex` and `-Knopreex` options.

In some cases, advance evaluation of invariant expressions causes the execution of instructions that should not be executed according to the logic of the program, resulting in errors. However, they will not affect the results of execution or their precision.

Moving the execution location may cause an unexpected standard library function call, invalid access to an element of an array, or a division-by-zero error. Examples of the effects that may occur due to advance evaluation of invariant expressions are given below.

##### Example 1: Referencing of Array Elements

<pre>int a[10], b[10]; ... for(i = 0; i &lt; 10; i++) {     if(j &lt; 10) {         a[i] = b[j] * f;     } }</pre>	->	<pre>int a[10], b[10]; ... t = b[j] * f; for(i = 0; i &lt; 10; i++) {     if(j &lt; 10) {         a[i] = t;     } }</pre>
--	----	---

In the program on the left, when the variable `j` is less than 10, array element `b[j]` is referenced, so the declared range of array `b` is not exceeded.

However, when array element `b[j]` is moved outside the loop due to advance evaluation of invariant expressions, it is referenced regardless of the value of `j`. As a result, if `j` is 10 or greater, the declared range of the array is exceeded when referenced. If the value of `j` is extremely large, an illegal area will be referenced and program execution will be aborted. `t` is a variable of type `int` generated by the compiler.

##### Example 2: Division

<pre>int a[100], b, f; ... for(i = 0; i &lt; n; i++) {     if(f != 0) {         a[i] = b / f;     } else {         a[i] = 0;     } }</pre>	->	<pre>int a[100], b, f; ... t = b / f; for(i = 0; i &lt; n; i++) {     if(f != 0) {         a[i] = t;     } else {         a[i] = 0;     } }</pre>
--	----	---

In the program on the left, the division is performed only when the variable `f` is not 0, so the error by division exception will not occur.

Since the division `b/f` is moved outside the loop due to the advance evaluation of invariant expressions, the division is performed regardless of the value of `f`. As a result, if the variable `f` is 0, program execution will be aborted. `t` is an `int` type variable generated by the compiler.

#### 3.3.1.2 Arithmetic Evaluation Method Modification

The method of evaluating arithmetic expressions is changed to shorten execution time. This modification may cause side effects.

##### Modification of the Operation Sequence

The operation sequence is modified in order to promote compile-time operations on constants, movement of invariant expressions, instruction scheduling, and other types of optimization.

There may be cases in which calculations that did not give exceptions (overflow or underflow) in the operation sequence prior to the modification will do so as a result of the modification. In addition, in floating-point operations, the precision of the operation results may differ depending on the number of valid digits in each operation term.

Modification of the operation sequence can be controlled with the `-Keval` and `-Knoeval` options.

An example of using this optimization is given below.

Example: Modification of the operation sequence

<pre>double a[100], b[100], c[100]; double x, y; ... for(i = 0; i &lt; 100; i++) {     a[i] = b[i] * x * c[i] * y; }</pre>	->	<pre>double a[100], b[100], c[100]; double x, y; ... t = x * y; for(i = 0; i &lt; 100; i++) {     a[i] = b[i] * c[i] * t; }</pre>
--	----	---

Once the operations among invariant terms are collected, `x*y` is moved outside the loop as an invariant expression.

Even though the calculation on the left operates as designed and does not generate exceptions, depending on the values of the arrays `b` and `c` and the variables `x` and `y`, the calculation on the right may generate exceptions or the intended result may not be obtained. `t` is a `double` type variable generated by the compiler.

Note that programs that generate exceptions or change the precision of operation results due to modification of the operation sequence are naturally programs that process values near the limit values or perform precision-sensitive calculations. Therefore, even if side effects occur due to the implementation of this optimization, this optimization may not be the basic reason.

### Conversion of Division to Multiplication

The division of invariant floating type data within a loop is changed to multiplication. This conversion increases execution speed. Note that this conversion may affect the precision of the execution results.

The conversion of division to multiplication can be controlled with the `-Keval` and `-Knoeval` options.

An example of this optimization and its effects is given below.

Example: Conversion of Division to Multiplication

<pre>float a[10], b[10], v; ... for(i = 0; i &lt; 10; i++) {     a[i] = b[i] / v; }</pre>	->	<pre>float a[10], b[10], v; ... t = 1.0 / v; for(i = 0; i &lt; 10; i++) {     a[i] = b[i] * t; }</pre>
---	----	--

The operation `b[i]/v` within the loop is converted to the multiplication `b[i]*t`. `t` is a variable of type `float` generated by the compiler.

In the program on the right, `1.0/v` is moved outside the loop and changed to multiplication of `t`, so there is potential for calculation errors.

## 3.3.2 Optimization of Pointers

When pointers are not being optimized, the compiler performs on the safe side assuming that it is not known what the pointer variable is pointing to.

In contrast, when the optimization of pointers is in effect, the compiler performs with the following prerequisites:

- Prerequisite 1

No other pointer variable points to the same area.

- Prerequisite 2

Situations in which an area is accessed with a pointer variable are not mixed with those in which the same area is accessed directly without a pointer.

For example, if the areas accessed by pointer variables are limited to those allocated dynamically with the `malloc` function or similar, prerequisite 2 is satisfied.

When these prerequisites are met, a high degree of optimization can be applied; for example, the elimination of common expressions and the movement of invariant expressions.

However, if optimization is performed on programs that do not satisfy the above prerequisites, the user may not obtain the intended results.

The optimization of pointers can be controlled with the `-Krestp` and `-Knorestp` options.

#### Example 1: Optimization of Pointers (Explanation of Prerequisite 1)

<pre>a1 = *p + 100; *q = 0; a2 = *p + 100;</pre>	->	<pre>a1 = *p + 100; *q = 0; a2 = a1; /*Changed by elimination of common expressions*/</pre>
--	----	---

Optimization is performed under the premise that the pointer variable `p` and the pointer variable `q` do not point to the same area. As a result, `*p+100` is considered to be a common expression.

If the pointer variable `p` and the pointer variable `q` do point to the same area, unexpected results may occur.

#### Example 2: Optimization of Pointers (Explanation of Prerequisite 2)

<pre>a1 = *p + b; x = 0; a2 = *p + b;</pre>	->	<pre>a1 = *p + b; x = 0; a2 = a1; /* Changed by elimination of common expressions */</pre>
---	----	--

Optimization is performed under the premise that the pointer variable `p` does not point to any of the variables `a1`, `a2`, `b`, or `x`. As a result, `*p+b` is considered to be a common expression.

If the pointer variable `p` points to the variable `a1` or the variable `x`, unexpected results may occur.

Note that these prerequisites are applied to headers, too.

### 3.3.3 Multi-Operation Function

---

A multi-operation function is a function achieving improved execution performance by performing multiple functions of the same type using one call.

When the `-Kmfunc` option is specified or the `mfunc` specifier is effective, the compiler analyzes and replaces a loop with a multi-operation function.

For details about the `-Kmfunc` option and the functions targeted by multi-operation functions, see Section "[2.2.2.6 -K Option](#)".

For details about the `mfunc` specifier, see Section "[3.4.1 Using Optimization Control Line \(OCL\)](#)".

#### 3.3.3.1 About Calling of Multi-Operation Functions

When the compiler cannot analyze complicated loops, the user program can call a multi-operation function directly.

Note the following:

- Include "`fjmfunc.h`" in the program.
- If a user program calls a multi-operation function directly, the argument used to return the result and the memory used for other arguments should be separate. If the memory addresses overlap, execution may be incorrect.
- Specify the `-Kmfunc` option when the program is linked.
- The argument check may be omitted in the multi-operation function for high-speed. Therefore, the user program may be terminated abnormally when the special value (NaN and Inf, etc.) defined by IEEE 754 is input.

"[Table 3.1 List of direct calling multi-operation function](#)" lists the multi-operation functions that can be called directly.

Table 3.1 List of direct calling multi-operation function

Function	Type Argument	Calling Format	Content of Calculation
acos	long n; float x[n],y[n];	void v_acos(x,&n,y);	for(i=0; i<n; i++) y[i] = acosf(x[i]);
	long n; double x[n],y[n];	void v_dacos(x,&n,y);	for(i=0; i<n; i++) y[i] = acos(x[i]);
asin	long n; float x[n],y[n];	void v_asin(x,&n,y);	for(i=0; i<n; i++) y[i] = asinf(x[i]);
	long n; double x[n],y[n];	void v_dasin(x,&n,y);	for(i=0; i<n; i++) y[i] = asin(x[i]);
atan	long n; float x[n],y[n];	void v_atan(x,&n,y);	for(i=0; i<n; i++) y[i] = atanf(x[i]);
	long n; double x[n],y[n];	void v_datan(x,&n,y);	for(i=0; i<n; i++) y[i] = atan(x[i]);
atan2	long n; float x1[n],x2[n],y[n];	void v_atan2(x1,x2,&n,y);	for(i=0; i<n; i++) y[i] = atan2f(x1[i],x2[i]);
	long n; double x1[n],x2[n],y[n];	void v_datan2(x1,x2,&n,y);	for(i=0; i<n; i++) y[i] = atan2(x1[i],x2[i]);
erf	long n; float x[n],y[n];	void v_erf(x,&n,y);	for(i=0; i<n; i++) y[i] = erff(x[i]);
	long n; double x[n],y[n];	void v_derf(x,&n,y);	for(i=0; i<n; i++) y[i] = erf(x[i]);
erfc	long n; float x[n],y[n];	void v_erfc(x,&n,y);	for(i=0; i<n; i++) y[i] = erfcf(x[i]);
	long n; double x[n],y[n];	void v_derfc(x,&n,y);	for(i=0; i<n; i++) y[i] = erfc(x[i]);
exp	long n; float x[n],y[n];	void v_exp(x,&n,y);	for(i=0; i<n; i++) y[i] = expf(x[i]);
	long n; double x[n],y[n];	void v_dexp(x,&n,y);	for(i=0; i<n; i++) y[i] = exp(x[i]);
exp10	long n; float x[n],y[n];	void v_exp10(x,&n,y);	for(i=0; i<n; i++) y[i] = exp10f(x[i]);
	long n; double x[n],y[n];	void v_dexp10(x,&n,y);	for(i=0; i<n; i++) y[i] = exp10(x[i]);
log	long n; float x[n],y[n];	void v_log(x,&n,y);	for(i=0; i<n; i++) y[i] = logf(x[i]);
	long n; double x[n],y[n];	void v_dlog(x,&n,y);	for(i=0; i<n; i++) y[i] = log(x[i]);
log10	long n; float x[n],y[n];	void v_log10(x,&n,y);	for(i=0; i<n; i++) y[i] = log10f(x[i]);
	long n; double x[n],y[n];	void v_dlog10(x,&n,y);	for(i=0; i<n; i++) y[i] = log10(x[i]);

Function	Type Argument	Calling Format	Content of Calculation
sin	long n; float x[n],y[n];	void v_sin(x,&n,y);	for(i=0; i<n; i++) y[i] = sinf(x[i]);
	long n; double x[n],y[n];	void v_dsin(x,&n,y);	for(i=0; i<n; i++) y[i] = sin(x[i]);
cos	long n; float x[n],y[n];	void v_cos(x,&n,y);	for(i=0; i<n; i++) y[i] = cosf(x[i]);
	long n; double x[n],y[n];	void v_dcos(x,&n,y);	for(i=0; i<n; i++) y[i] = cos(x[i]);
sincos	long n; float x[n],y1[n],y2[n];	void v_scn(x,&n,y1,y2);	for(i=0; i<n; i++) sincosf(x[i], &y1[i], &y2[i]);
	long n; double x[n],y1[n],y2[n];	void v_dscn(x,&n,y1,y2);	for(i=0; i<n; i++) sincos(x[i], &y1[i], &y2[i]);
pow	long n; float x1[n],x2[n],y[n];	void v_pow(x1,x2,&n,y);	for(i=0; i<n; i++) y[i] = powf(x1[i], x2[i]);
	long n; float x1[n],x2[n],y[n],a;	void v_pow1(x,&a,&n,y);	for(i=0; i<n; i++) y[i] = powf(x[i], a);
	long n; double x1[n],x2[n],y[n];	void v_dpow(x1,x2,&n,y);	for(i=0; i<n; i++) y[i] = pow(x1[i], x2[i]);
	long n; double x1[n],x2[n],y[n],a;	void v_dpow1(x,&a,&n,y);	for(i=0; i<n; i++) y[i] = pow(x[i], a);

The following is an example of calling a multi-operation function.

Example 1:

```
#include <math.h>
#define N 1000
...
double a;
double x[N], y[N], z[N];
int i;
...
for(i = 0; i < N; i++) {
    y[i] = exp(x[i]);
    z[i] = pow(x[i], a);
}
```

Direct calling a multi-operation function

```
#include <fjfunc.h>
#define N 1000
...
long n = N;
double a;
double x[N], y[N], z[N];
...
v_dexp(x, &n, y);
v_dpow1(x, &a, &n, z);
```

### Example 2:

When function calling exists in "if" statement and array in function can compress an array only value with true, it can use a multi-operation function.

```
#include <math.h>
#define N 1000
...
double a, x, y;
int i;
...
for(i = 0; i < N; i++) {
    x = ...
    if(x > a) {
        y = y + sin(x);
    }
}
```

### Direct calling multi-operation function

```
#include <fjfunc.h>
#define N 1000
...
double a, x, y;
double wx[N], wy[N];
int i;
long j;
...
j = 0;
for(i = 0; i < N; i++) {
    x = ...
    if(x > a) {
        wx[j] = x;
        j++;
    }
}
v_dsine(wx, &j, wy);
for(i = 0; i < j; i++) {
    y = y + wy[i];
}
```

### 3.3.3.2 Effects of Compiler Option -Kmfunc=3

When the -Kmfunc=3 option is specified, a multi-operation function is used in the loop which contains an "if" statement. This may cause an abend as shown in the following example.

The following example also illustrates the effects of the -Kmfunc=3 option.

### Example:

```
double x[1000], y[1000];
int i;
...
for(i = 0; i < 2000; i++) {
    if(i < 1000) {
        y[i] = exp(x[i]);
    }
}
```

### Explanation:

In the above example, when the condition  $i < 1000$  is true, the `exp` function is executed. Since the value of the argument in the `exp` function is less than 1000, the value of subscripts of array `x` and `y` never exceed the declared range.



However, if a multi-operation function is used in the loop, array `x` exceeds the declared range because the array elements for the iteration count of the "for" loop are used as the argument of the `exp` function regardless of the condition of the "if" statement. In this case, a storage protection exception, and an abnormal end may occur.

In this case, do not specify the compiler option `-Kmfunc=3`.


### 3.3.4 Loop Striping

Using loop striping all of the statements executed within a loop are expanded up to  $N$  times within the loop (where  $N$  is called the "striping size"). The overhead for loop iteration is reduced, and software pipelining may be promoted.

This optimization is performed on loops which contain no jump from the inside to the outside and no jump from the outside to the inside.

Example: Loop Striping

When the striping size is 4 by specifying `-Kstriping=4`, the loop is expanded as follows:



In loop striping optimization, the statements executed within the loops are expanded. Therefore, the size of the object program increases as loop unrolling. Further, the compile time may increase. Note that the execution performance may decrease because the number of registers used are increased.

Loop striping can be suppressed by specifying the `-Knostriping` option.

Moreover, the message to show the optimized loop and the multiplicity can be output by specifying the `-Koptmsg=2` option.

### 3.3.5 zfill

The `zfill` optimization speeds up write operations for array data by using an instruction that allocates space on the cache for writing (DC ZVA) without loading data from the memory. The `zfill` is applied to array data to be written in a loop.

However, the `zfill` is not applied to array data which is referred in the same loop, accessed non-sequentially, or stored in an "if" statement. Further, if the `zfill` optimization is applied, prefetch instructions to the second level cache are not generated.

The `zfill` optimization works on the data  $N$  blocks ahead of the address pointed to by the target store instruction where one block is 256 byte-long and  $N$  is an integer value between 1 and 100. If a value is not specified for  $N$ , the compiler will automatically determine a value.

The following optimizations are suppressed because the loop is modified that data is always stored to the cache lines allocated by the `zfill` optimization:

- Loop unrolling
- Loop striping

Further, performance may also be reduced under the following conditions:

- The program is not affected by memory bandwidth bottleneck.

- When the loop iteration count is too few.
- When the block size is explicitly specified with the `-Kzfill=N` option and the memory size where the data is written in the loop is smaller than `N` blocks.

Do not specify the `-Kzfill` option if the execution performance is likely to decrease as a result.

This optimization is not expected to be effective for all the loops in the program, so it is recommended to use the optimization control specifier `zfill` for each loop rather than the `-Kzfill` option for the entire program.

### 3.3.6 Loop Versioning

The loop versioning generates two loops to which optimization is applied and is not applied. The object program created by the compiler judges the data dependency of array at execution time and selects either loop.

The optimization, such as SIMD, software pipelining or automatic parallelization, is promoted by loop versioning.

The size of the object program and compile time may increase because the loop versioning generates two loops.

The execution performance may decrease because the processing for judgement is overhead.

The loop versioning is applied only to the innermost loop which contains a single array whose data dependency is unknown. The loop versioning may decrease the overhead because of the unknown dependency.

Moreover, the message to show where optimized by the loop versioning can be output by specifying the `-Koptmsg=2` option.

The following example shows the loop versioning.

Example:

```
void f(double *a, double *b, int n) {
    for (int i = 0; i < n; i++) {
        a[i] += b[i];
    }
}
```

Optimized pseudo-code

```
void f(double *a, double *b, int n) {
    if ((&a[n-1] < &b[0]) || (&b[n-1] < &a[0])) {
#pragma loop norecurrence
        for (int i = 0; i < n; i++) { /* optimized loop */
            a[i] += b[i];
        }
    } else {
        for (int i = 0; i < n; i++) { /* no-optimized loop */
            a[i] += b[i];
        }
    }
}
```

It is unknown at compile time that the arrays "a" and "b" are overlapped. When the `-Kloop_versioning` option is specified, the compiler generates two loops with "if" statement for judging the overlap of the arrays at execution time. If the array "a" and "b" are found to have no data dependency, SIMD may be promoted for the loop.

### 3.3.7 Clone Optimization

Clone optimization generates conditional branches of variables for loops in order to promote other optimizations, such as full unrolling. Clone optimization is applied by specifying the optimization control line. The calculation result is not guaranteed if the optimization control line clone is used incorrectly. See Section "[3.4.1.2 Optimization Control Specifier](#)" and "[3.4.1.3 Notes for Optimization Control Specifiers](#)" for more information about clone.

Example:

```
#pragma loop clone n==10
for (i=0;i<n; ++i) {
    a[i] = i;
}
```

->

```
if (n==10) {
    for (i=0;i<10;++i) {
        a[i] = i;
    }
} else {
    for (i=0;i<n;++i) {
        a[i] = i;
    }
}
```

### 3.3.8 Unroll-and-Jam

Unroll-and-jam is the optimization to unroll an outer loop of a nested loop  $N$  times and jam the unrolled statements into the inner loop as loop fusion. Unroll-and-jam promotes removing common expression and improves the execution performance. The increase of data stream and change of the order of data access may decrease the cache efficiency and the execution performance. This optimization is not applied to innermost loop.

This optimization is not expected to be effective for all the loop in the program, so it is recommended to use the optimization control specifier `unroll_and_jam` or `unroll_and_jam_force` for each loop rather than the `-Kunroll_and_jam` option for the entire program.

For details about `unroll_and_jam` and `unroll_and_jam_force` specifier, see Section "3.4.1 Using Optimization Control Line (OCL)".

To output line numbers and expansion counts of optimized loops, specify `-Koptmsg=2` option.

Example: Unroll-and-jam

```
#pragma loop unroll_and_jam_force 2
for (i=0; i<128; i++) {
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        ...
    }
}
```

->

```
for (i=0; i<128; i+=2) {
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        a[i+1][j] = b[i+1][j] + b[i+2][j];
        ...
    }
}
```

### 3.3.9 Tree-Height-Reduction Optimization

The tree-height-reduction optimization increases the instruction-level parallelism by changing the order of calculations in a loop to keep the calculation tree as short as possible.

The calculation tree is an expression of arithmetic operations in a form of tree structure in the order of operations. Leaf nodes of the tree structure hold values and other nodes hold operators. Nodes whose operator has higher priority are placed lower in the following example.

This optimization is applied to floating-point operations when the `-Keval_concurrent` option is set. In this case, side effects (calculation errors) may occur in the execution results. See "3.6.1 Side Effect of Optimizations for Floating-Point Operation" for the side effects of optimizations.

An example of optimization is shown below.

Example:

```
for(i = 0; i < n; i++) {
    x[i] = a[i] * b[i] + c[i] * d[i] + e[i] * f[i] + g[i] * h[i];
}
```

Figure 3.3 The calculation tree when the -Keval\_noconcurrent option or optimization control specifier eval\_noconcurrent is set





Figure 3.4 The calculation tree when the -Keval\_concurrent option or optimization control specifier eval\_concurrent is set



### 3.3.10 Loop Fission

Loop fission is a function to fission a loop into two or more small loops for aiming at the following objectives:

- Promotion of the software pipelining
- Improvement of the cache memory efficiency
- Resolution of register shortage

This optimization is performed on loops which contain no jump from the inside to the outside and no jump from the outside to the inside.

Loop fission is performed on loops to which the loop\_fission\_target specifier is specified when the -Kloop\_fission and -Kocl options are set. It can be suppressed by specifying the -Kloop\_nofission option.

The message to show the optimized loop can be output by specifying the -Koptmsg=2 option.

An example of automatic loop fission is given below.

Example: Automatic loop fission

```

#pragma loop loop_fission_target
for(i = 0; i < N; i++) {
    e[i] = a1[i]*b1[i] + a2[i]*b2[i] +
          a3[i]*b3[i] + a4[i]*b4[i] +
          a5[i]*b5[i];
    f[i] = c1[i]*d1[i] + c2[i]*d2[i] +
          c3[i]*d3[i] + c4[i]*d4[i] +
          c5[i]*d5[i];
}

```

->

```

for(i = 0; i < N; i++) {
    e[i] = a1[i]*b1[i] + a2[i]*b2[i] +
          a3[i]*b3[i] + a4[i]*b4[i] +
          a5[i]*b5[i];
}
for(i = 0; i < N; i++) {
    f[i] = c1[i]*d1[i] + c2[i]*d2[i] +
          c3[i]*d3[i] + c4[i]*d4[i] +
          c5[i]*d5[i];
}

```

### 3.3.10.1 Strip-Mining

Strip-mining is an optimization to fragment the loop by small iteration.

Applied by this optimization for loops after automatic loop fission, improvement of the cache memory efficiency is expected for the data accessed between fissioned loops.

Strip-mining is performed when the `-Kloop_fission_stripmining` option is specified. The message to show the optimized loop can be output by specifying the `-Koptmsg=2` option.

Examples of automatic loop fission and strip-mining with specifying the `-Kloop_fission_stripmining=256` are shown as follows.

Example: Automatic loop fission and strip-mining

Source code

```

#pragma loop loop_fission_target
float a[N], b[N], p[N], q;
for(i = 0; i < N; i++) {
    q = a[i] + b[i];
    ...
    p[i] = p[i] + q;
    ...
}

```

Automatic loop fission pseudo-code

```

float a[N], b[N], p[N], q;
float temparray_q[N];
for(i = 0; i < N; i++) {
    temparray_q[i] = a[i] + b[i];
    ...
}
for(i = 0; i < N; i++) {
    p[i] = p[i] + temparray_q[i];
    ...
}

```

The temporary array `temparray_q` is generated by the compiler with loop fission. The array is required to hand over data between the fissioned loops in the source code. The number of array elements is the same value as the number of loop iteration.

Automatic loop fission and strip-mining pseudo-code

```

#define MIN(x,y) ((x<y)?x:y)
float a[N], b[N], p[N], q;
float temparray_q[256];
for(ii = 0; ii < N; ii = ii+256) {
    for(i = ii; i < MIN(N, ii+256); i++) {
        temparray_q[i-ii] = a[i] + b[i];
        ...
    }
    for(i = ii; i < min(N, ii+256); i++) {

```

```

        p[i] = p[i] + tmparray_q[i-ii];
        ...
    }
}

```

A loop is generated outside of the fissioned loop. The original loop iteration is fragmented by the strip length 256. The number of elements of temporary array tmparray\_q is 256.

### 3.3.11 Strict Aliasing

Strict Aliasing is optimization with considering overlaps of memory regions according to the strict aliasing rules defined by language standard. Some optimizations such as the movement of invariant expressions and elimination of common expressions are promoted.

Strict Aliasing can be controlled with the `-Kstrict_aliasing` and `-Knostrict_aliasing` options.

Examples of alias are shown below.

#### Example 1:

If the `-Kstrict_aliasing` option is specified, the int type and the float type have no alias.

```

void sub(int *a, float *b) {
    long i;
    for (i=0; i<256; i++) {
        a[i] = *b;
    }
}

```

#### Example 2:

Even if the `-Kstrict_aliasing` option is specified, the int type and the unsigned int type have alias.

```

void sub(int *a, unsigned int *b) {
    long i;
    for (i=0; i<256; i++) {
        a[i] = *b;
    }
}

```

#### Example 3:

Even if the `-Kstrict_aliasing` option is specified, the char type has alias with all types.

```

void sub(int *a, char *b) {
    long i;
    for (i=0; i<256; i++) {
        a[i] = *b;
    }
}

```

#### Example 4:

If the `-Kstrict_aliasing` option is specified to the source code that does not conform to the rules, the compiler may incorrectly determined that there is no alias. The execution result may be incorrect.

```

void sub(int *i_pointer, long *l_pointer) {
    *i_pointer = 10;
    *l_pointer = 15;
}
int main() {
    long double l_double = 0.0;
    int *i_pointer = &l_double;
    long *l_pointer = &l_double;
    sub(i_pointer, l_pointer);
    if (*i_pointer == *l_pointer) {
        // OK
    }
}

```

```

} else {
    // NG
}
return 0;
}

```

## 3.3.12 Fujitsu Matrix Template Library

Fujitsu Matrix Template Library (hereinafter called FMTL) is the high speed linear algebra operation library for C++. There are the following four types:

1. Serial execution version
2. High speed serial execution version using math library (BLAS)
3. Thread parallel execution version
4. High speed thread parallel execution version using math library (BLAS)

When using FMTL, specify `-Nfmtl=fmtl_mode` at both compile time and link time. For more detail, see "[2.2.2.7 -N Option](#)".



### Note

Note the following points regarding thread parallel execution version of FMTL.

- Can not call the FMTL from the following sections:
  - In the parallel region by OpenMP
  - In the parallel region by automatic parallelization
- Thread parallel execution version of FMTL use function of OpenMP C++ language. Therefore, program using thread parallel execution version of FMTL operates as an OpenMP program. For details on how to execute the OpenMP program, refer to "[4.3.2 Execution Process](#)".

FMTL consists of the following six classes:

Class name	Description
matrix	General matrix class (" <a href="#">3.3.12.2 matrix</a> ")
banded_matrix	Band matrix class (" <a href="#">3.3.12.3 banded_matrix</a> ")
hermitian_matrix	Hermitian matrix class (" <a href="#">3.3.12.4 hermitian_matrix</a> ")
symmetric_matrix	Symmetric matrix class (" <a href="#">3.3.12.5 symmetric_matrix</a> ")
triangular_matrix	Triangular matrix class (" <a href="#">3.3.12.6 triangular_matrix</a> ")
sparse_matrix	Sparse matrix class (" <a href="#">3.3.12.7 sparse_matrix</a> ")

### 3.3.12.1 Common Definition

#### 3.3.12.1.1 Class Constant

The following class constants can be used in FMTL.

Table 3.2 Kind of matrix

Constant	Description
<code>fj::general</code>	General matrix
<code>fj::hermitian</code>	Hermitian matrix
<code>fj::symmetric</code>	Symmetric matrix

Constant	Description
<code>fj::triangular</code>	Triangular matrix

Table 3.3 Access order

Constant	Description
<code>fj::row_major</code>	Access rows with priority.
<code>fj::col_major</code>	Access column with priority.

Table 3.4 Used part of matrix

Constant	Description
<code>fj::upper</code>	Use the upper triangle of the matrix.
<code>fj::lower</code>	Use the lower triangle of the matrix.

### 3.3.12.1.2 Member Function

From all matrix classes of FMTL, the following member functions exist.

Table 3.5 Iterator functions

Function prototype	Description
<code>iterator begin();</code>	Returns an iterator pointing to the first element in the matrix.
<code>const_iterator begin() const;</code>	
<code>iterator end();</code>	Returns an iterator pointing to the next of the last element in the matrix.
<code>const_iterator end() const;</code>	
<code>const_iterator cbegin() const;</code>	Returns a read-only iterator pointing to the first element in the matrix.
<code>const_iterator cend() const;</code>	Returns a read-only iterator pointing to the next of the last element in the matrix.
<code>reverse_iterator rbegin();</code>	Returns a reverse iterator pointing to the last element in the matrix.
<code>const_reverse_iterator rbegin() const;</code>	
<code>reverse_iterator rend();</code>	Returns a reverse iterator pointing to the first element in the matrix.
<code>const_reverse_iterator rend() const;</code>	
<code>const_reverse_iterator crbegin() const;</code>	Returns a read-only reverse iterator pointing to the last element in the matrix.
<code>const_reverse_iterator crend() const;</code>	Returns a read-only reverse iterator pointing to the first element in the matrix.

Table 3.6 Size functions

Function prototype	Description
<code>size_type row_size();</code>	Returns the size of the row.
<code>size_type col_size();</code>	Returns the size of the col.
<code>size_type max_size();</code>	Returns the maximum size ( <code>row * col</code> ).



Table 3.7 Element access functions/ Container handling functions

Function prototype	Description
<code>T &amp; operator[](size_type i);</code>	Returns "i"th element. If row or col is larger than the size, operation is not guaranteed.
<code>T const &amp; operator[](size_type i) const;</code>	
<code>T &amp; operator()(size_type row, size_type col);</code>	Returns "row * col"th element. If row or col is larger than the size, operation is not guaranteed.
<code>T const &amp; operator()(size_type row, size_type col) const;</code>	
<code>T data();</code>	Returns a pointer to the head of the matrix.
<code>T const data() const;</code>	
<code>T &amp; insert_element(size_type row, size_type col, T const &amp;value);</code>	Rewrites the "row * col"th number element to value.
<code>T &amp; erase_element(size_type row, size_type col);</code>	Initialize the "row * col"th number element with 0.
<code>void resize(size_type size1, size_type size2);</code>	Resize the matrix to "size1 * size2" size. Values before construction are not guaranteed.
<code>matrix&lt;T, L&gt; transpose();</code>	Transposes a matrix.
<code>void swap(matrix&lt;T, L&gt; &amp;x);</code>	Exchanges all element of the matrix for all element of x.
<code>void clear();</code>	Removes all elements from the matrix (which are destroyed), leaving the container with a size of 0.

### 3.3.12.2 matrix

General matrix class

Standard header

```
fmtl/matrix
```

Format

```
fj::matrix<T[,L]>
```

Template argument

Argument name	Description	Possible attributes	Default value
T	Type of general matrix	float or double or complex<float> or complex<double>	-

Argument name	Description	Possible attributes	Default value
L	Access order	fj::row_major or fj::col_major	fj::row_major

-: No settings

### Note

If an attribute is specified with a combination other than the above, the calculation result is guaranteed, but performance improvement can not be expected.

### Example

Declaration of general matrix mat of  $n * n$  (T: double type, L: Access rows with priority)

```
fj::matrix<double, fj::row_major> mat(n, n);
```

### Member function

This class provides the following functions and "3.3.12.1.2 Member Function".

Table 3.8 Constructors and destructors of general matrix

Function prototype	Description
matrix();	Constructs an empty general matrix.
matrix( size_type row, size_type col);	Constructs an empty general matrix in "row * col" size.
matrix( size_type row, size_type col, T const &value);	Constructs a general matrix in "row * col" size. Each element is initialized to "value".
matrix( matrix const &x);	Constructs a general matrix with a copy of each of the elements in general matrix x, in the same order.
matrix( matrix &&x);	Constructs a general matrix that acquires the elements of general matrix x. (Their ownership is directly transferred)
matrix &operator=( matrix const &x);	Deletes all elements of the general matrix and then duplicates the elements of general matrix x.
matrix &operator=( matrix &&x);	Deletes all elements of the general matrix and then moves general matrix x.
~matrix();	Destroys the general matrix.

### Non-member function

This class provides the following functions.

Table 3.9 Non-member functions of general matrix

Function prototype	Description
template <class T, class L> vector<T> row_value( 	Returns the elements (Array) of "i"th row.

Function prototype	Description
<code>matrix&lt;T, L&gt; const &amp;x, size_type i);</code>	
<code>template &lt;class T, class L&gt; vector&lt;T&gt; col_value( matrix&lt;T, L&gt; const &amp;x, size_type j);</code>	Returns the elements (Array) of "j"th column.
<code>template &lt;class T, class L&gt; vector&lt;T&gt; operator*( matrix&lt;T, L&gt; const &amp;x, vector&lt;T&gt; const &amp;y);</code>	Returns the operation result of matrix-vector product of general matrix x and vector y.
<code>template &lt;class T, class L&gt; matrix&lt;T, L&gt; operator*( matrix&lt;T, L&gt; const &amp;x, matrix&lt;T, L&gt; const &amp;y);</code>	Returns the operation result of matrix-matrix product of general matrix x and general matrix y.
<code>template &lt;class T, class L&gt; void swap( matrix&lt;T, L&gt; &amp;x, matrix&lt;T, L&gt; &amp;y);</code>	The contents of general matrix x are exchanged with those of general matrix y.

### 3.3.12.3 banded\_matrix

Band matrix class

Standard header

```
fmtl/banded
```

Format

```
fj::banded_matrix<T[,C][,L][,UL]>
```

Template argument

Argument name	Description	Possible attributes		Default value
T	Type of band matrix	float or double	complex<float> or complex<double>	-
C	Kind of band matrix	fj::general or fj::symmetric or fj::triangular	fj::general or fj::hermitian	fj::general
L	Access order	fj::row_major or fj::col_major		fj::row_major
UL	Used part of band matrix	fj::upper or fj::lower		fj::upper

-: No settings

#### Note

If an attribute is specified with a combination other than the above, the calculation result is guaranteed, but performance improvement can not be expected.



## Example

Declaration of band matrix `mat` of  $n * n$  (T: double type, C: General matrix, L: Access rows with priority, UL: Use the upper triangle of the matrix)

```
fj::banded_matrix<double, fj::general, fj::row_major, fj::upper> mat(n, n);
```

### Member function

This class provides the following functions and "[3.3.12.1.2 Member Function](#)".

Table 3.10 Constructors and destructors of band matrix

Function prototype	Description
<code>banded_matrix();</code>	Constructs an empty band matrix.
<code>banded_matrix( size_type row, size_type col);</code>	Constructs an empty band matrix in "row * col" size.
<code>banded_matrix( size_type row, size_type col, size_type lower, size_type upper);</code>	Constructs an empty band matrix in "row * col" size, lower bandwidth "lower", upper bandwidth "upper".
<code>banded_matrix( size_type row, size_type col, T const &amp;value);</code>	Constructs a band matrix in "row * col" size. Each diagonal element is initialized to "value".
<code>banded_matrix( size_type row, size_type col, size_type lower, size_type upper, T const &amp;value);</code>	Constructs a band matrix in "row * col" size, lower bandwidth "lower", upper bandwidth "upper". Each element between upper and lower is initialized to value.
<code>banded_matrix( banded_matrix const &amp;x);</code>	Constructs a band matrix with a copy of each of the elements in band matrix x, in the same order.
<code>banded_matrix( banded_matrix &amp;&amp;x);</code>	Constructs a band matrix that acquires the elements of band matrix x. (Their ownership is directly transferred).
<code>banded_matrix &amp;operator=( banded_matrix const &amp;x);</code>	Deletes all elements of the band matrix and then duplicates the elements of band matrix x.
<code>banded_matrix &amp;operator=( banded_matrix &amp;&amp;x);</code>	Deletes all elements of the band matrix and then moves band matrix x.
<code>~banded_matrix();</code>	Destroys the band matrix.

Table 3.11 Size functions of band matrix

Function prototype	Description
<code>size_type diag_low_size();</code>	Returns the bandwidth of the lower triangular matrix.
<code>size_type diag_up_size();</code>	Returns the bandwidth of the upper triangular matrix.
<code>set_diag_low_size(size_type lower);</code>	Set the bandwidth of the lower triangular matrix.

Function prototype	Description
<code>set_diag_up_size(size_type lower);</code>	Set the bandwidth of the upper triangular matrix.
<code>set_diag_low_up_size( size_type lower, size_type upper);</code>	Sets the bandwidth of the lower triangular matrix and upper triangular matrix.

#### Non-member function

This class provides the following functions.

Table 3.12 Non-member functions of band matrix

Function prototype	Description
<code>template &lt;class T, class C, class L, class UL&gt; vector&lt;T&gt; row_value( banded_matrix&lt;T, C, L, UL&gt; const &amp;x, size_type i);</code>	Returns the elements (Array) of "i"th row.
<code>template &lt;class T, class C, class L, class UL&gt; vector&lt;T&gt; col_value( banded_matrix&lt;T, C, L, UL&gt; const &amp;x, size_type j);</code>	Returns the elements (Array) of "j"th column.
<code>template &lt;class T, class C, class L, class UL&gt; vector&lt;T&gt; operator*( banded_matrix&lt;T, C, L, UL&gt; const &amp;x, vector&lt;T&gt; &amp;y);</code>	Returns the operation result of matrix-vector product of band matrix x and vector y.
<code>template &lt;class T, class C, class L, class UL&gt; void swap( banded_matrix&lt;T, C, L, UL&gt; &amp;x, banded_matrix&lt;T, C, L, UL&gt; &amp;y);</code>	The contents of band matrix x are exchanged with those of band matrix y.

### 3.3.12.4 hermitian\_matrix

Hermitian matrix class

Standard header

```
fmt1/hermitian
```

Format

```
fj::hermitian_matrix<T[,L][,UL]>
```

Template argument

Argument name	Description	Possible attributes	Default value
T	Type of hermitian matrix	<code>complex&lt;float&gt;</code> or <code>complex&lt;double&gt;</code>	-
L	Access order	<code>fj::row_major</code> or <code>fj::col_major</code>	<code>fj::row_major</code>
UL	Used part of hermitian matrix	<code>fj::upper</code> or <code>fj::lower</code>	<code>fj::upper</code>

--: No settings

## Note

If an attribute is specified with a combination other than the above, the calculation result is guaranteed, but performance improvement can not be expected.

## Example

Declaration of hermitian matrix mat of  $n * n$  (T: double type, L: Access rows with priority, UL: Use the upper triangle of the matrix)

```
fj::hermitian_matrix<double, fj::row_major, fj::upper> mat(n, n);
```

### Member function

This class provides the following functions and "[3.3.12.1.2 Member Function](#)".

Table 3.13 Constructors and destructors of hermitian matrix

Function prototype	Description
<code>hermitian_matrix();</code>	Constructs an empty hermitian matrix.
<code>hermitian_matrix( size_type row, size_type col);</code>	Constructs an empty hermitian matrix in "row * col" size.
<code>hermitian_matrix( size_type row, size_type col, T const &amp;value);</code>	Constructs a hermitian matrix in "row * col" size. Each element is initialized to "value".
<code>hermitian_matrix( hermitian_matrix const &amp;x);</code>	Constructs a hermitian matrix with a copy of each of the elements in hermitian matrix x, in the same order.
<code>hermitian_matrix( hermitian_matrix &amp;&amp;x);</code>	Constructs a hermitian matrix that acquires the elements of hermitian matrix x. (Their ownership is directly transferred).
<code>hermitian_matrix &amp;operator=( hermitian_matrix const &amp;x);</code>	Deletes all elements of the hermitian matrix and then duplicates the elements of hermitian matrix x.
<code>hermitian_matrix &amp;operator=( hermitian_matrix &amp;&amp;x);</code>	Deletes all elements of the hermitian matrix and then moves hermitian matrix x.
<code>~hermitian_matrix();</code>	Destroys the hermitian matrix.

### Non-member function

This class provides the following functions.

Table 3.14 Non-member functions of hermitian matrix

Function prototype	Description
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; row_value( hermitian_matrix&lt;T, L, UL&gt; const &amp;x, size_type i);</code>	Returns the elements (Array) of "i"th row.
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; col_value( hermitian_matrix&lt;T, L, UL&gt; const &amp;x, size_type j);</code>	Returns the elements (Array) of "j"th column.

Function prototype	Description
<pre>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; operator*(     hermitian_matrix&lt;T, L, UL&gt; const &amp;x,     vector&lt;T&gt; const &amp;y);</pre>	Returns the operation result of matrix-vector product of hermitian matrix x and vector y.
<pre>template &lt;class T, class L, class UL&gt; matrix&lt;T, L&gt; operator*(     hermitian_matrix&lt;T, L, UL&gt; const &amp;x,     matrix&lt;T, L&gt; const &amp;y);</pre>	Returns the operation result of matrix-matrix product of hermitian matrix x and general matrix y.
<pre>template &lt;class T, class L, class UL&gt; void swap(     hermitian_matrix&lt;T, L, UL&gt; &amp;x,     hermitian_matrix&lt;T, L, UL&gt; &amp;y);</pre>	The contents of hermitian matrix x are exchanged with those of hermitian matrix y.

### 3.3.12.5 symmetric\_matrix

Symmetric matrix class

Standard header

```
fmtl/symmetric
```

Format

```
fj::symmetric_matrix<T[,L][,UL]>
```

Template argument

Argument name	Description	Possible attributes	Default value
T	Type of symmetric matrix	float or double or complex<float> or complex<double>	-
L	Access order	fj::row_major or fj::col_major	fj::row_major
UL	Used part of symmetric matrix	fj::upper or fj::lower	fj::upper

-: No settings



#### Note

If an attribute is specified with a combination other than the above, the calculation result is guaranteed, but performance improvement can not be expected.



#### Example

Declaration of symmetric matrix mat of  $n * n$  (T: double type, L: Access rows with priority, UL: Use the upper triangle of the matrix)

```
fj::symmetric_matrix<double, fj::row_major, fj::upper> mat(n, n);
```

Member function

This class provides the following functions and "[3.3.12.1.2 Member Function](#)".

Table 3.15 Constructors and destructors of symmetric matrix

Function prototype	Description
<code>symmetric_matrix();</code>	Constructs an empty symmetric matrix.
<code>symmetric_matrix( size_type row, size_type col);</code>	Constructs an empty symmetric matrix in "row * col" size.
<code>symmetric_matrix( size_type row, size_type col, T const &amp;value);</code>	Constructs a symmetric matrix in "row * col" size. Each element is initialized to "value".
<code>symmetric_matrix( symmetric_matrix const &amp;x);</code>	Constructs a symmetric matrix with a copy of each of the elements in symmetric matrix x, in the same order.
<code>symmetric_matrix( symmetric_matrix &amp;&amp;x);</code>	Constructs a symmetric matrix that acquires the elements of symmetric matrix x. (Their ownership is directly transferred).
<code>symmetric_matrix &amp;operator=( symmetric_matrix const &amp;x);</code>	Deletes all elements of the symmetric matrix and then duplicates the elements of symmetric matrix x.
<code>symmetric_matrix &amp;operator=( symmetric_matrix &amp;&amp;x);</code>	Deletes all elements of the symmetric matrix and then moves symmetric matrix x.
<code>~symmetric_matrix();</code>	Destroys the symmetric matrix.

#### Non-member function

This class provides the following functions.

Table 3.16 Table 1.15 Non-member functions of symmetric matrix

Function prototype	Description
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; row_value( symmetric_matrix&lt;T, L, UL&gt; const &amp;x, size_type i);</code>	Returns the elements (Array) of "i"th row.
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; col_value( symmetric_matrix&lt;T, L, UL&gt; const &amp;x, size_type j);</code>	Returns the elements (Array) of "j"th column.
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; operator*( symmetric_matrix&lt;T, L, UL&gt; const &amp;x, vector&lt;T&gt; const &amp;y);</code>	Returns the operation result of matrix-vector product of symmetric matrix x and vector y.
<code>template &lt;class T, class L, class UL&gt; matrix&lt;T, L&gt; operator*( symmetric_matrix&lt;T, L, UL&gt; const &amp;x, matrix&lt;T, L&gt; const &amp;y);</code>	Returns the operation result of matrix-matrix product of symmetric matrix x and general matrix y.
<code>template &lt;class T, class L, class UL&gt; void swap( symmetric_matrix&lt;T, L, UL&gt; &amp;x, symmetric_matrix&lt;T, L, UL&gt; &amp;y);</code>	The contents of symmetric matrix x are exchanged with those of symmetric matrix y.

### 3.3.12.6 triangular\_matrix

Triangular matrix class



## Standard header

```
fmt1/triangular
```

## Format

```
fj::triangular_matrix<T[,L][,UL]>
```

## Template argument

Argument name	Description	Possible attributes	Default value
T	Type of triangular matrix	float or double or complex<float> or complex<double>	-
L	Access order	fj::row_major or fj::col_major	fj::row_major
UL	Used part of triangular matrix	fj::upper or fj::lower	fj::upper

-: No settings



### Note

If an attribute is specified with a combination other than the above, the calculation result is guaranteed, but performance improvement can not be expected.



### Example

Declaration of triangular matrix mat of  $n * n$  (T: double type, L: Access rows with priority, UL: Use the upper triangle of the matrix)

```
fj::triangular_matrix<double, fj::row_major, fj::upper> mat(n, n);
```

## Member function

This class provides the following functions and "[3.3.12.1.2 Member Function](#)".

Table 3.17 Constructors and destructors of triangular matrix

Function prototype	Description
<code>triangular_matrix();</code>	Constructs an empty triangular matrix.
<code>triangular_matrix( size_type row, size_type col);</code>	Constructs an empty triangular matrix in "row * col" size.
<code>triangular_matrix( size_type row, size_type col, T const &amp;value);</code>	Constructs a triangular matrix in "row * col" size. Each element is initialized to "value".
<code>triangular_matrix( triangular_matrix const &amp;x);</code>	Constructs a triangular matrix with a copy of each of the elements in triangular matrix x, in the same order.
<code>triangular_matrix( triangular_matrix &amp;&amp;x);</code>	Constructs a triangular matrix that acquires the elements of triangular matrix x. (Their ownership is directly transferred).

Function prototype	Description
<code>triangular_matrix &amp;operator=( triangular_matrix const &amp;x);</code>	Deletes all elements of the triangular matrix and then duplicates the elements of triangular matrix x.
<code>triangular_matrix &amp;operator=( triangular_matrix &amp;&amp;x);</code>	Deletes all elements of the triangular matrix and then moves triangular matrix x.
<code>~triangular_matrix();</code>	Destroys the triangular matrix.

#### Non-member function

This class provides the following functions.

Table 3.18 Non-member functions of triangular matrix

Function prototype	Description
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; row_value( triangular_matrix&lt;T, L, UL&gt; const &amp;x, size_type i);</code>	Returns the elements (Array) of "i"th row.
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; col_value( triangular_matrix&lt;T, L, UL&gt; const &amp;x, size_type j);</code>	Returns the elements (Array) of "j"th column.
<code>template &lt;class T, class L, class UL&gt; vector&lt;T&gt; operator*( triangular_matrix&lt;T, L, UL&gt; const &amp;x, vector&lt;T&gt; &amp;y);</code>	Returns the operation result of matrix-vector product of triangular matrix x and vector y.
<code>template &lt;class T, class L, class UL&gt; matrix&lt;T, L&gt; operator*( triangular_matrix&lt;T, L, UL&gt; const &amp;x, matrix&lt;T, L&gt; &amp;y);</code>	Returns the operation result of matrix-matrix product of triangular matrix x and general matrix y.
<code>template &lt;class T, class L, class UL&gt; void swap( triangular_matrix&lt;T, L, UL&gt; &amp;x, triangular_matrix&lt;T, L, UL&gt; &amp;y);</code>	The contents of triangular matrix x are exchanged with those of triangular matrix y.

### 3.3.12.7 sparse\_matrix

Sparse matrix class

Standard header

```
fmtl/sparse
```

Format

```
fj::sparse_matrix<T[,C][,L][,UL]>
```

Template argument

Argument name	Description	Possible attributes		Default value
T	Type of sparse matrix	float or double	complex<float> or complex<double>	-
C	Kind of sparse matrix	fj::symmetric or fj::triangular	fj::general or fj::triangular	fj::symmetric
L	Access order	fj::row_major or fj::col_major		fj::row_major

Argument name	Description	Possible attributes	Default value
UL	Used part of sparse matrix	fj::upper or fj::lower	fj::upper

-: No settings

### Note

If an attribute is specified with a combination other than the above, the calculation result is guaranteed, but performance improvement can not be expected.

### Example

Declaration of sparse matrix mat of  $n * n$  (T: double type, C: Symmetric matrix, L: Access rows with priority, UL: Use the upper triangle of the matrix)

```
fj::matrix<double, fj::row_major> mat(n, n);
```

### Member function

This class provides the following functions and "3.3.12.1.2 Member Function".

Table 3.19 Constructors and destructors of sparse matrix

Function prototype	Description
<code>sparse_matrix();</code>	Constructs an empty sparse matrix.
<code>sparse_matrix( size_type row, size_type col);</code>	Constructs an empty sparse matrix in "row * col" size.
<code>sparse_matrix( size_type row, size_type col, T const &amp;value);</code>	Constructs a sparse matrix in "row * col" size. Each element is initialized to "value".
<code>sparse_matrix( sparse_matrix const &amp;x);</code>	Constructs a sparse matrix with a copy of each of the elements in sparse matrix x, in the same order.
<code>sparse_matrix( sparse_matrix &amp;&amp;x);</code>	Constructs a sparse matrix that acquires the elements of sparse matrix x. (Their ownership is directly transferred).
<code>sparse_matrix &amp;operator=( sparse_matrix const &amp;x);</code>	Deletes all elements of the sparse matrix and then duplicates the elements of sparse matrix x.
<code>sparse_matrix &amp;operator=( sparse_matrix &amp;&amp;x);</code>	Deletes all elements of the sparse matrix and then moves sparse matrix x.
<code>~sparse_matrix();</code>	Destroys the sparse matrix.

### Non-member function

This class provides the following functions.

Table 3.20 Non-member functions of sparse matrix

Function prototype	Description
<code>template &lt;class T, class C, class L, class UL&gt; vector&lt;T&gt; row_value( </code>	Returns the elements (Array) of "i"th row.

Function prototype	Description
<pre>sparse_matrix&lt;T, C, L, UL&gt; const &amp;x, size_type i);</pre>	
<pre>template &lt;class T, class C, class L, class UL&gt; vector&lt;T&gt; col_value(     sparse_matrix&lt;T, C, L, UL&gt; const &amp;x,     size_type j);</pre>	Returns the elements (Array) of "j"th column.
<pre>template &lt;class T, class C, class L, class UL&gt; vector&lt;T&gt; operator*(     sparse_matrix&lt;T, C, L, UL&gt; const &amp;x,     vector&lt;T&gt; &amp;y);</pre>	Returns the operation result of matrix-vector product of sparse matrix x and vector y.
<pre>template &lt;class T, class C, class L, class UL&gt; void swap(     sparse_matrix&lt;T, C, L, UL&gt; &amp;x,     sparse_matrix&lt;T, C, L, UL&gt; &amp;y);</pre>	The contents of sparse matrix x are exchanged with those of sparse matrix y.

### 3.3.12.8 Usage Example

An example of using FMTL (General matrix) is shown below.

Before applying FMTL (General matrix):

```
#include <iostream>
#include <vector>
using namespace std;

int main(void)
{
    int n = 10;
    vector<vector<double>> mat(n, vector<double>(n));
    vector<double> vec(n),ret(n);

    for (int i = 0; i < n; i++) {
        vec[i] = i;
        for (int j = 0; j < n; j++) {
            mat[i][j] = i * j;
        }
    }

    for(int i = 0; i < n; i++){
        ret[i] = 0;
        for(int j = 0; j < n; j++){
            ret[i] += mat[i][j] * vec[j];
        }
    }

    cout << "ret[0] =" << ret[0] << endl;
    cout << "ret[9] =" << ret[9] << endl;

    return 0;
}
```

After applying FMTL (General matrix):

```
#include <iostream>
#include <vector>
#include <fmtl/matrix>
using namespace std;

int main(void)
```

```

{
    int n = 10;
    fj::matrix<double> mat(n, n);
    vector<double> vec(n),ret(n);

    for (int i = 0; i < n; i++) {
        vec[i] = i;
        for (int j = 0; j < n; j++) {
            mat(i,j) = i * j;
        }
    }

    ret = mat * vec;

    cout << "ret[0] =" << ret[0] << endl;
    cout << "ret[9] =" << ret[9] << endl;

    return 0;
}

```

## 3.4 Using Optimization Functions

This section describes techniques for effectively using the optimization functions.

### 3.4.1 Using Optimization Control Line (OCL)

It may be possible to increase the optimization level using certain `#pragma` directives in the source code. The following `#pragma` directive falls into this category. The `-Kocl` option enables the `#pragma` directive optimization, and the `-Knoocl` option suppresses the effect of the `#pragma` directive optimization.

#### 3.4.1.1 Types of Optimization Control Lines

The types of optimization control lines and their scope are shown in "[Table 3.21 Scope of Optimization Control Lines](#)".

Table 3.21 Scope of Optimization Control Lines

Preprocessing Literal Line	Type of optimization control line	Scope
global	global line	Compilation unit
procedure	procedure line	Function
loop	loop line	Subsequent loop
statement	statement line	Subsequent statement

#### global line

The global line is used when important information in the compilation unit is given to the compiler.

#### Description Format

```
#pragma[ fj] global Optimization-control-specifier
```

#### Insert Position

This is positioned at the top of the corresponding compilation unit.

#### procedure line

The procedure is used when important information in a function is given to the compiler.

#### Description Format

```
#pragma[ fj] procedure Optimization-control-specifier
```

### Insert Position

This is positioned between a declaration statement of the corresponding function (including the statement in which both the declaration and initialization value are specified) and an initial execution statement.

### loop line

The loop line is used when valuable information is given to the compiler.

### Description Format

```
#pragma[ fj] loop Optimization-control-specifier
```

### Insert Position

This is positioned just before the corresponding loop.

### statement line

The statement line is used when the valuable information is instructed to the compiler.

### Description Format

```
#pragma[ fj] statement Optimization-control-specifier
```

### Insert Position

This is positioned just before the corresponding statement.



### Note

Only one *optimization-control-specifier* can be specified for one optimization control line.

Two or more *optimization-control-specifiers* can be specified with consecutive lines.

## 3.4.1.2 Optimization Control Specifier

The optimization control specifiers that can be used in the optimization control line are shown in "[Table 3.22 Optimization Control Specifiers that can be specified to optimization control line](#)".

Table 3.22 Optimization Control Specifiers that can be specified to optimization control line

Optimization Control Specifiers	Explanation	Optimization control line that can be Specified <sup>[a]</sup>			
		global line	procedure line	loop line	statement line
<code>array_declaration_opt</code>	Specifies to perform optimization assuming that the array subscript does not exceed the range of the array declaration.	*	*	*	-
<code>noarray_declaration_opt</code>	Specifies to perform the optimization without assuming that the array subscript does not exceed the range of the array declaration.	*	*	*	-
<code>assume {shortloop noshortloop memory_bandwidth nomemory_bandwidth time_saving_compilation notime_saving_compilation}</code>	Directs to control optimization considering features of the program.	*	*	-	-
<code>clone var==n1[,n2]...<sup>[b]</sup></code>	Directs to generate conditional branches along to the arguments and generate loop copies and	-	-	*	-

Optimization Control Specifiers	Explanation	Optimization control line that can be Specified <sup>[a]</sup>			
		global line	procedure line	loop line	statement line
	<p>put them for the conditional blocks assuming that the <i>var</i> is invariant in the loops.</p> <p>The conditional expressions are equals-to expressions which operands are <i>var</i> and <i>n1[,n2]...</i> of specified arguments.</p> <p>Variable <i>var</i> is a name of integer variable.</p> <p>An integer value from -9223372036854775808 to 9223372036854775807 can be specified as constant values(<i>n1[ , n2]...</i>).</p>				
<code>eval</code>	Performs the optimization that changes the expression evaluation sequence.	*	*	*	-
<code>noeval</code>	Suppresses the optimization that changes the expression evaluation sequence.	*	*	*	-
<code>eval_concurrent</code>	Gives the priority to the instruction-level parallelism in the tree-height-reduction optimization.	*	*	*	-
<code>eval_noconcurrent</code>	Suppresses the instruction-level parallelism and give priority to utilizing FMA instructions in the tree-height-reduction optimization.	*	*	*	-
<code>extract_stride_store</code>	Directs to use scalar instructions for stride access store in the target loop.	*	*	*	-
<code>noextract_stride_store</code>	Directs to use SIMD instructions for stride access store in the target loop.	*	*	*	-
<code>fission_point [n]</code>	<p>Performs loop fission optimization.</p> <p><i>n</i> indicates the count of nested loops to be distributed from innermost loop, which should be an integer value from 1 to 6.</p>	-	-	-	*
<code>fp_contract</code>	Performs optimizations using the Floating-Point Multiply-Add/Subtract instructions.	*	*	*	-
<code>nofp_contract</code>	Suppresses optimizations using the Floating-Point Multiply-Add/Subtract instructions.	*	*	*	-
<code>fp_relaxed</code>	Converts floating point divisions and <code>sqrt</code> functions into reciprocal approximation instructions.	*	*	*	-
<code>nofp_relaxed</code>	Converts floating point divisions and <code>sqrt</code> functions into normal instructions and <code>sqrt</code> instruction.	*	*	*	-
<code>fullunroll_pre_simd [n]</code>	Performs the prioritize promotion of the full unrolling before SIMD. <i>n</i> is the upper limit of	-	-	*	-

Optimization Control Specifiers	Explanation	Optimization control line that can be Specified <sup>[a]</sup>			
		global line	procedure line	loop line	statement line
	the iteration count for a target loop. <i>n</i> is a number from 2 to 100.				
<code>nofullunroll_pre_simd</code>	Suppresses the full unrolling before SIMD.	-	-	*	-
<code>iterations max=n1</code> <code>iterations avg=n2</code> <code>iterations min=n3</code>	Performs optimization assuming the maximum iteration count is <code>max=n1</code> , the average of the iteration count is <code>avg=n2</code> , and the minimum iteration count is <code>min=n3</code> .  An integer value from 0 to 2147483647 can be specified as <i>n1</i> , <i>n2</i> and <i>n3</i> .  One or more <code>max</code> , <code>avg</code> , and <code>min</code> specifiers can be specified in random order. When specifying <code>max</code> , <code>avg</code> , and <code>min</code> at the same time, use space to separate them like " <code>iterations max=n1 avg=n2 min=n3</code> ".	*	*	*	-
<code>loop_blocking n</code>	Performs the loop-blocking optimization.  <i>n</i> indicates the block size which should be an integer value from 2 to 10000.	*	*	*	-
<code>loop_noblocking</code>	Suppresses the loop-blocking optimization.	*	*	*	-
<code>loop_fission_stripmining [n "c-level"]</code>	Performs the strip-mining optimization for automatic fissioned loops. <i>n</i> is a number from 2 to 100000000. <i>c-level</i> is L1 or L2.	*	*	*	-
<code>loop_nofission_stripmining</code>	Suppresses the strip-mining optimization for automatic fissioned loops.	*	*	*	-
<code>loop_fission_target [c1 ls]</code>	Performs automatic loop fission optimization.	-	-	*	-
<code>loop_fission_threshold n</code>	Specifies the threshold to decide the granularity of loop after automatic loop fission. <i>n</i> is a number from 1 to 100.	*	*	*	-
<code>loop_interchange var1,var2[,var3]...</code> <sup>[b]</sup>	Designates that the sequence of a nested loop is changed according to the specified order.	-	-	*	-
<code>loop_nointerchange</code>	Suppresses the change of the sequence of a nested loop.	-	-	*	-
<code>loop_nofusion</code>	Suppresses the loop fusion.	*	*	*	-
<code>loop_part_simd</code>	Performs optimization which divides loops and partially uses SIMD extensions.	*	*	*	-
<code>loop_nopart_simd</code>	Suppresses optimization which divides loops and partially uses SIMD extensions.	*	*	*	-
<code>loop_perfect_nest</code>	Performs optimization which fissions the imperfectly nested loop into the perfectly nested loops.	*	*	*	-



Optimization Control Specifiers	Explanation	Optimization control line that can be Specified <sup>[a]</sup>			
		global line	procedure line	loop line	statement line
loop_noperfect_nest	Suppress optimization which fissions the imperfectly nested loop into the perfectly nested loops.	*	*	*	-
loop_versioning	Designates to perform optimization by the loop versioning.	*	*	*	-
loop_noversioning	Cancels the loop_versioning.	*	*	*	-
mfunc [ <i>level</i> ]	Designates to change the function to a multi-operation function. Decimal 1, 2 or 3 can be specified as <i>level</i> .	*	*	*	-
nomfunc	Designates canceling to change the function to a multi-operation function.	*	*	*	-
noalias	Designates that pointer variables and other variable do not share memory area.	*	*	*	-
norecurrence [ <i>array1[,array2]...</i> ] <sup>[b]</sup>	Declares array for which the loop slice optimization can be performed.	*	*	*	-
novrec [ <i>array1[,array2]...</i> ] <sup>[b]</sup>	Instructs to the system that there is no recurrent operation with the specified array in the specified loop.	*	*	*	-
preex	Optimizes by evaluating an invariant first.	*	*	*	-
nopreex	Suppresses the optimization of evaluating an invariant first.	*	*	*	-
prefetch	Performs the automatic prefetch function of the compiler.	*	*	*	-
noprefetch	Suppresses the automatic prefetch function of the compiler.	*	*	*	-
prefetch_cache_level <i>c-level</i>	Designates cache-level to prefetch of data. 1, 2 or all can be specified as <i>c-level</i> .	*	*	*	-
prefetch_conditional	Directs the system to generate the prefetch instruction to array data used in blocks in "if" statements or "switch" statements.	*	*	*	-
prefetch_noconditional	Suppresses the prefetch_conditional	*	*	*	-
prefetch_indirect	Directs the system to generate the prefetch instruction to array data accessed indirectly in loops.	*	*	*	-
prefetch_noindirect	Suppresses the prefetch_indirect	*	*	*	-
prefetch_infer	Directs the system to generate the prefetch instruction for sequential access when the distance of prefetch is unknown.	*	*	*	-
prefetch_noinfer	Suppresses the prefetch_infer.	*	*	*	-

Optimization Control Specifiers	Explanation	Optimization control line that can be Specified <sup>[a]</sup>			
		global line	procedure line	loop line	statement line
<code>prefetch_iteration n</code>	Directs the system to be target of a prefetch instruction would be data that is referred to <i>n</i> iterations of a loop. This targets prefetch instructions only for the first level cache. Decimal from 1 to 10000 can be specified as <i>n</i> .  If the optimization using SIMD extensions is applied to a loop, specify the loop iteration count after using SIMD extensions for <i>n</i> .	*	*	*	-
<code>prefetch_iteration_L2 n</code>	Directs the system to be target of a prefetch instruction would be data that is defined or is referred to <i>n</i> iterations of a loop. This targets prefetch instructions only for the second level cache.  Decimal from 1 to 10000 can be specified as <i>n</i> .  If the optimization using SIMD extensions is applied to a loop, specify the loop iteration count after using SIMD extensions for <i>n</i> .	*	*	*	-
<code>prefetch_line n</code>	Directs the system to generate prefetch instructions to prefetch data after <i>n</i> cache line(s). This targets prefetch instructions only for first level cache.  Decimal from 1 to 100 can be specified as <i>n</i> .	*	*	*	-
<code>prefetch_line_L2 n</code>	Directs the system to generate prefetch instructions to prefetch data after <i>n</i> cache line(s). This targets prefetch instructions only for second level cache.  Decimal from 1 to 100 can be specified as <i>n</i> .	*	*	*	-
<code>prefetch_sequential [auto soft]</code>	Directs that prefetch instructions are created for array data that is accessed sequentially. If neither <code>auto</code> nor <code>soft</code> is specified, <code>auto</code> is set by default.	*	*	*	-
<code>prefetch_nosequential</code>	Directs that prefetch instructions are not created for array data that is accessed sequentially.	*	*	*	-
<code>prefetch_stride [soft hard_auto hard_always]</code>	Directs that prefetch instructions are generated for array data that is accessed with a stride larger than the cache line size used in a loop.	*	*	*	-
<code>prefetch_nostride</code>	Directs that prefetch instructions are not generated for array data that is accessed with a stride larger than the cache line size used in a loop.	*	*	*	-

Optimization Control Specifiers	Explanation	Optimization control line that can be Specified <sup>[a]</sup>			
		global line	procedure line	loop line	statement line
<code>prefetch_strong</code>	Directs that the prefetch instructions for the first level cache are to be the strong prefetch.	*	*	*	-
<code>prefetch_nostrong</code>	Directs that the prefetch instructions generated for the first level cache will not be strong prefetch.	*	*	*	-
<code>prefetch_strong_L2</code>	Directs that the prefetch instructions generated for the second level cache are to be strong prefetch.  The prefetch instructions for the second level cache generated for the target loop will be strong prefetch.	*	*	*	-
<code>prefetch_nostrong_L2</code>	Directs that the prefetch instructions generated for the second level cache are not to be strong prefetch.	*	*	*	-
<code>preload</code>	Directs to perform the speculative execution of load instructions	*	*	*	-
<code>nopreload</code>	Directs to suppress the speculative execution of load instructions	*	*	*	-
<code>scache_isolate_way L2=n1 [L1=n2]</code> and <code>end_scache_isolate_way</code>	Directs the maximum number of ways in sector 1 of cache.	-	*	-	*
<code>scache_isolate_assign array1[,array2]...</code> <sup>[b]</sup> and <code>end_scache_isolate_assign</code>	Directs the array to store in sector 1 of cache.	-	*	-	*
<code>simd [aligned unaligned]</code>	Performs the optimization that uses SIMD Extensions.	*	*	*	-
<code>nosimd</code>	Suppresses the optimization that uses SIMD Extensions.	*	*	*	-
<code>simd_listv [all then else]</code>	Performs list vector conversion to the statements in the block of "if" statement.	-	-	-	*
<code>simd_noredundant_vl</code>	This directs that SIMD with redundant executions for the SIMD width is not applied.	*	*	*	-
<code>simd_use_multiple_structures</code>	Directs to use the SVE Load Multiple Structures instructions and SVE Store Multiple Structures instructions.	*	*	*	-
<code>simd_nouse_multiple_structures</code>	Directs not to use the SVE Load Multiple Structures instructions and SVE Store Multiple Structures instructions.	*	*	*	-
<code>striping [n]</code>	Performs loop striping.	*	*	*	-

Optimization Control Specifiers	Explanation	Optimization control line that can be Specified <sup>[a]</sup>			
		global line	procedure line	loop line	statement line
	<i>n</i> indicates the stripe-length (number of expansions) which should an integer value be from 2 to 100.				
<code>nostriping</code>	Suppresses loop striping.	*	*	*	-
<code>swp</code>	Performs software pipelining.	*	*	*	-
<code>noswp</code>	Suppresses software pipelining.	*	*	*	-
<code>swp_freq_rate n</code> <code>swp_ireg_rate n</code> <code>swp_preg_rate n</code>	Changes the condition of the register number of software pipelining.  <i>n</i> indicates the rate (percentage) of the register number that can be used by software pipelining. <i>n</i> is a number from 1 to 1000.	*	*	*	-
<code>swp_policy {auto small large}</code>	Specifies a policy to select an instruction scheduling algorithm used in software pipelining.	*	*	*	-
<code>swp_weak</code>	Adjusts software pipelining to make overlapping of the instructions less.	*	*	*	-
<code>unroll [n "full"]</code>	Performs unrolling for the corresponding loop.  <i>n</i> is an iteration count for unrolling and is an integer value from 2 to 100.	-	-	*	-
<code>nounroll</code>	Suppresses unrolling for the corresponding loop.	-	-	*	-
<code>unroll_and_jam [n]</code>	Performs unroll-and-jam optimization to the loop when the optimization is expected to be effective.  <i>n</i> is an iteration count for unrolling and is an integer value from 2 to 100.	*	*	*	-
<code>unroll_and_jam_force [n]</code>	Performs unroll-and-jam optimization.  <i>n</i> is an iteration count for unrolling and is an integer value from 2 to 100.	-	-	*	-
<code>nounroll_and_jam</code>	Suppresses unroll-and-jam optimization.	*	*	*	-
<code>unswitching</code>	Instructs loop unswitching to "if" statement.	-	-	-	*
<code>zfill [N]</code>	Directs that the zfill optimization to be performed.  The optional parameter <i>N</i> is an integer between 1 and 100 that specifies the number of blocks the DC ZVA instruction writes.	-	-	*	-
<code>nozfill</code>	Directs that the zfill optimization not to be performed.	-	-	*	-

[a]

\* : An optimization control line accepts the optimization control specifier.

- : An optimization control line does not accept the optimization control specifier.

[b] *var*, *var1*, *var2*, *var3*, *array1*, and *array2* must be declared before they are used.

#### array\_declaration\_opt specifier

The `array_declaration_opt` specifier instructs to perform the optimization assuming that the array subscript does not exceed the range of the array declaration.

An example is shown in the following.

Example:

```
double a[8];
#pragma loop array_declaration_opt
for(int i = 0; i < n; i++) {
    a[i] = 0;
}
```

The optimization is performed assuming that the operation range of the subscript is equal to or less than the SIMD length.

#### noarray\_declaration\_opt specifier

The `noarray_declaration_opt` specifier instructs to perform the optimization without assuming that the array subscript does not exceed the range of the array declaration.

An example is shown in the following.

Example:

```
double a[8];
#pragma loop noarray_declaration_opt
for(int i = 0; i < n; i++) {
    a[i] = 0;
}
```

Performs optimization with the operating range of the subscript unknown.

#### assume specifier

The `assume` specifier instructs to control optimization considering features of the program. Multiple `assume` specifiers can be specified at the same time.

Either of specifiers shown below is instructed after `assume` specifier.

##### shortloop

The `assume shortloop` specifier instructs to assume that iteration counts are small when the iteration count of the innermost loop in the program is unknown at compilation. Optimizations such as automatic parallelization, loop unrolling, and software pipelining may be controlled or invalidated.

##### noshortloop

The `assume noshortloop` specifier instructs to assume that iteration counts are not small when the iteration count of the innermost loop in the program is unknown at compilation.

##### memory\_bandwidth

The `assume memory_bandwidth` specifier instructs to assume that innermost loops in the program has a memory bandwidth bottleneck. Optimizations such as `zfill` optimization promotion and `s` software pipelining may be controlled or invalidated.

##### nomemory\_bandwidth

The `assume nomemory_bandwidth` specifier instructs to assume that innermost loops in the program does not have a memory bandwidth bottleneck.

## time\_saving\_compilation

The `assume time_saving_compilation` specifier instructs to control optimization for decreasing compilation time of the program.

Optimization is controlled to decrease compilation time of the program. General optimizations and inline expansions may be controlled or invalidated.

## notime\_saving\_compilation

The `assume notime_saving_compilation` specifier instructs to optimize with giving priority to the speed of the executable program rather than decrease of the compilation time.

An example is shown in the following.

### Example:

```
#pragma global assume shortloop
#pragma global assume memory_bandwidth
#pragma global assume time_saving_compilation
void sub() {
    for(int i = 0; i < N; i++) {
        a[i] = b[i];
    }
}
```

The optimization is controlled considering the specified features in the target loop.

## clone specifier

The `clone` specifier instructs to generate conditional branches along to the arguments and generate loop copies and put them for the conditional blocks assuming that the *var* is invariant in the loops. The order of the generated branches is the order of the corresponding arguments. This specifier promotes other optimizations, such as full unrolling. The execution result is not guaranteed in case that the value of variable *var* changes in the loop. For more details, refer to "[3.4.1.3 Notes for Optimization Control Specifiers](#)". The size of the object program and compile time may increase because the clone makes copies of loops. This specifier is effective only when the `-O3` option is set.

Variable *var* is a name of integer variable.

There are some cases that cannot be specified like follows.

- Structure member variable
- Union member variable
- Class member variable
- Reference argument
- Array elements
- Threadprivate variable
- Variable with scope resolution operator

An integer value from -9223372036854775808 to 9223372036854775807 can be specified as constant values ( $n1[,n2]...$ ).

You can use comma to enumerate constant values and colon to specify a range. The following two specifications are equivalent.

```
#pragma loop clone var==1,3:5,7
```

```
#pragma loop clone var==1,3,4,5,7
```

You can specify up to 20 values at once. When more than 20 values are specified, excess values are ignored. The following specification is limited in the number of values because it specifies 30 values.

```
#pragma loop clone var==11:40
```

The above specification is treated as the same as the following specification because excess values are ignored.

```
#pragma loop clone var==11:30
```

Examples of correct usage are following.

#### Example 1:

Loop is copied with "if" statements in the order of the specification.

```
#pragma loop clone N==10,20
for (i=0;i<N; ++i) {
    a[i] = i;
}
```

->

```
if (N==10) {
    for (i=0;i<10;++i) {
        a[i] = i;
    }
} else if (N==20) {
    for (i=0;i<20;++i) {
        a[i] = i;
    }
} else {
    for (i=0;i<N;++i) {
        a[i] = i;
    }
}
```

#### Example 2:

Loop is copied with "if" statements in the order of the multiple specifications.

```
#pragma loop clone N==10
#pragma loop clone M==20
for (i=0;i<N;++i) {
    a[i] = M;
}
```

->

```
if (N==10) {
    for (i=0;i<10;++i) {
        a[i] = M;
    }
} else if (M==20) {
    for (i=0;i<N;++i) {
        a[i] = 20;
    }
} else {
    for (i=0;i<N;++i) {
        a[i] = M;
    }
}
```

#### Example 3:

clone specifiers for nested loops generates nested "if" statements.

```
#pragma loop clone M==10
for (i=0;i<M;++i) {
    #pragma loop clone N==20
    for (j=0;j<N;++j) {
        a[i][j] = 0;
    }
}
```

->

```
if (M==10) {
    for (i=0;i<10;++i) {
        if (N==20) {
            for (j=0;j<20;++j) {
                a[i][j] = 0;
            }
        } else {
            for (j=0;j<N;++j) {
                a[i][j] = 0;
            }
        }
    }
} else {
    for (i=0;i<M;++i) {
        if (N==20) {
            for (j=0;j<20;++j) {
                a[i][j] = 0;
            }
        }
    }
}
```

```
    } else {  
        for (j=0;j<N; j++) {  
            a[i][j] = 0;  
        }  
    }  
}
```

### eval specifier

The `eval` specifier instructs to perform the optimization that changes the expression evaluation sequence.

An example is shown in the following.

Example:

```
#pragma loop eval  
for(i = 0; i < n; i++) {  
    a[i] = a[i] + b[i] + c[i] + d[i];  
}
```

->

```
for(i = 0; i < n; i++) {  
    a[i] = (a[i] + b[i]) + (c[i] + d[i]);  
}
```

The optimization that changes the expression evaluation sequence is performed in the target loop.

### noeval specifier

The `noeval` specifier instructs to suppress the optimization that changes the expression evaluation sequence.

An example is shown in the following.

Example:

```
#pragma loop noeval  
for(i = 0; i < n; i++) {  
    a[i] = a[i] + b[i] + c[i] + d[i];  
}
```

->

```
for(i = 0; i < n; i++) {  
    a[i] = (a[i] + b[i]) + (c[i] + d[i]);  
}
```

The optimization that changes the expression evaluation sequence is suppressed in the target loop.

### eval\_concurrent specifier

The `eval_concurrent` specifier instructs to give priority to the instruction-level parallelism in the tree-height-reduction optimization.

An example is shown in the following.

Example:

```
#pragma loop eval_concurrent  
for(i = 0; i < n; i++) {  
    x[i] = a[i] * b[i] + c[i] * d[i] + e[i] * f[i] + g[i] * h[i];  
}
```

See section "[3.3.9 Tree-Height-Reduction Optimization](#)" for the tree-height-reduction optimization.

### eval\_noconcurrent specifier

The `eval_noconcurrent` specifier instructs to suppress the instruction-level parallelism and specifies to give priority to utilizing FMA instructions in the tree-height-reduction optimization.

An example is shown in the following.

Example:

```
#pragma loop eval_noconcurrent  
for(i = 0; i < n; i++) {  
    x[i] = a[i] * b[i] + c[i] * d[i] + e[i] * f[i] + g[i] * h[i];  
}
```



See section "3.3.9 Tree-Height-Reduction Optimization" for the tree-height-reduction optimization.

#### `extract_stride_store` specifier

The `extract_stride_store` specifier instructs to use scalar instructions for stride access store in the target loop.

An example is shown in the following.

Example:

```
#pragma loop extract_stride_store
for(int i = 0; i < n; i+=2) {
    a[i] = b[i] + c[i];
}
```

Performs optimization using scalar instructions for store of array a that can be converted to SIMD instruction.

#### `noextract_stride_store` specifier

The `noextract_stride_store` specifier instructs to use SIMD instructions for stride access store in the target loop.

An example is shown in the following.

Example:

```
#pragma loop extract_stride_store
for(int i = 0; i < n; i+=2) {
    a[i] = b[i] + c[i];
}
```

Performs optimization using SIMD instructions for array a.

#### `fission_point` specifier

The `fission_point` specifier instructs to perform loop fission optimization.

The number from 1 to 6 following the specifier instructs the count of nested loops to be distributed from innermost loop. If the count is omitted, the compiler distributes innermost loop.

The `fission_point` specifier is effective only when it is placed in the innermost loop.

An example is shown in the following.

Example:

```
for(j = 1; j < n; j++) {
    for(i = 1; i < n; i++) {
        a[i][j] = a[i-1][j-1];
        #pragma statement fission_point 1
        a[i][j] = a[i][j] + a[i-1][j];
    }
}
```

->

```
for(j = 1; j < n; j++) {
    for(i = 1; i < n; i++) {
        a[i][j] = a[i-1][j-1];
    }
    for(i = 1; i < n; i++) {
        a[i][j] = a[i][j] + a[i-1][j];
    }
}
```

### Note

If loop fission is performed on a loop that cannot be executed because of the program logic, Some side effects, such as exceptions on runtime, can happen.

For instance, the compiler generates some temporary arrays for passing data between the split loops. In this case, the necessary instructions may be executed speculatively.

In the following example, the load instructions of the array element `IDX[M][M]` is speculatively executed. If `N` is less than or equal to 0 and `M` used in array element `IDX[M][M]` is greater than the declared size, an exception may be thrown at runtime.

Example:

```

for(j = 0; j < N; j++) {
    for(i = 0; i < IDX[M][M]; i++) {
        x = a[i] + b[i];
#pragma statement fission_point 1
        c[i] = x;
    }
}

```

→

```

double *tmp;
tmp =(double *)malloc(IDX[M]
[M]*sizeof(double)); // speculative execution
for(j = 0; j < N; j++) {
    for(i = 0; i < IDX[M][M]; i++) {
        tmp[i] = a[i] + b[i];
    }
    for(i = 0; i < IDX[M][M]; i++) {
        c[i] = tmp[i];
    }
}
free(tmp);

```

### fp\_contract specifier

The `fp_contract` specifier instructs to apply optimizations using the Floating-Point Multiply-Add/Subtract instructions. However, using these instructions may cause calculation errors in the result due to rounding errors.

An example is shown in the following.

Example:

```

#pragma loop fp_contract
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i] * c[i];
}

```

The optimization that uses the Floating-Point Multiply-Add/Subtract instructions is performed in the target loop.

### nofp\_contract specifier

The `nofp_contract` specifier instructs to suppress optimizations using the Floating-Point Multiply-Add/Subtract instructions.

An example is shown in the following.

Example:

```

#pragma loop nofp_contract
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i] * c[i];
}

```

The optimization that uses the Floating-Point Multiply-Add/Subtract instructions is not performed in the target loop.

### fp\_relaxed specifier

The `fp_relaxed` specifier instructs to convert floating point division operations and `sqrt` functions into the reciprocal approximation operation with the reciprocal approximation instructions and the Floating-Point Multiply-Add/Subtract instructions. This conversion is applied to the single-precision and double-precision real type calculation. At the result of this conversion, the program is executed faster. However it may cause calculation errors in the result due to rounding errors. Also, this optimization may result in generating a floating-point exception whether the `-NRtrap` option is valid or not.

When the `-NRtrap` option and either of the `-Knosimd` option or the `-KNOSVE` option are effective, the optimization to convert a `sqrt` function into reciprocal approximation instructions is suppressed. Therefore, the execution performance may decrease as compared with when the `-NRnotrap` option is effective. Note that the reciprocal approximation instructions generated by this specifier may be evaluated in advance even if the `-Knopreex` option is enabled. And a floating-point exception may occur when the `-Knopreex` option, `-NRtrap` option, and this specifier are valid.

An example is shown in the following.

Example:

```

#pragma loop fp_relaxed
for(i = 0; i < n; i++) {

```

```

    a[i] = sqrt(b[i] / c[i]);
}

```

The optimization of using the reciprocal approximation instructions is performed in the target loop.

#### `nofp_relaxed` specifier

The `nofp_relaxed` specifier instructs to convert floating point division operations and `sqrt` functions into the normal division and `sqrt` instructions without using reciprocal approximation instructions.

An example is shown in the following.

Example:

```

#pragma loop nofp_relaxed
for(i = 0; i < n; i++) {
    a[i] = sqrt(b[i] / c[i]);
}

```

The normal division and `sqrt` instructions are used for the target loop.

#### `fullunroll_pre_simd` specifier

The `fullunroll_pre_simd` specifier instructs to perform the prioritize promotion of the full unrolling before SIMD.

The number from 2 to 100 following the specifier instructs the upper limit of the iteration count for a target loop. If the upper limit is not specified, the compiler will automatically determine the value.

Note that the `fullunroll_pre_simd` specifier targets only the loop immediately after the optimization control line.

If the number of the iteration is unknown at compilation, no optimization of the full unrolling before SIMD is performed.

An example is shown in the following.

Example:

```

for(i = 0; i < n; i++) {
#pragma loop fullunroll_pre_simd
    for(j = 0; j < 16; j++) {
        a[i][j] = b[j][i] + c[j][i];
    }
}

```

Full unrolling before SIMD is applied for the inner loop.

#### `nofullunroll_pre_simd` specifier

The `nofullunroll_pre_simd` specifier instructs to suppress the full unrolling before SIMD.

Note that the `nofullunroll_pre_simd` specifier targets only the loop immediately after the optimization control line.

An example is shown in the following.

Example:

```

for(i = 0; i < n; i++) {
#pragma loop nofullunroll_pre_simd
    for(j = 0; j < 5; j++) {
        a[i][j] = b[j][i] + c[j][i];
    }
}

```

Full unrolling before SIMD is not applied to the inner loop.

#### `iterations max=n1` specifier

#### `iterations avg=n2` specifier

#### `iterations min=n3` specifier

The `iterations` specifier instructs to optimize assuming the specified value is the loop iteration count.

max=n1, avg=n2, and min=n3 are specified following the specifier. An integer value from 0 to 2147483647 can be specified as *n1*, *n2*, and *n3*. One or more max, avg, and min specifiers can be specified in random order. When specifying max, avg, and min at the same time, use space to separate them like "iterations max=n1 avg=n2 min=n3". Note that values *n1*, *n2* and *n3* specified as max, avg, and min are required to be  $n1 \geq n2 \geq n3$ .

When max=n1 is specified following the specifier, the optimization assuming the maximum loop iteration count is *n1* is performed.

When avg=n2 is specified following the specifier, the optimization assuming the average loop iteration count is *n2* is performed.

When min=n3 is specified following the specifier, the optimization assuming the minimum loop iteration count is *n3* is performed.

This specifier is effective when the loop iteration count is unknown at compilation.

The execution result is guaranteed even if the actual average iteration loop count is different from *n2* because the specified average loop iteration count *n2* is used as an optimization hint.



The execution result is not guaranteed in the following cases:

- When the actual loop iteration count is greater than *n1*
- When the actual loop iteration count is less than *n3*

For details, see the iterations max=*n1* specifier and iterations min=*n3* specifier in Section "[3.4.1.3 Notes for Optimization Control Specifiers](#)".

This specifier can be specified at global line, procedure line, and loop line, but it is recommended to specify at loop line because the iteration counts of loops in the program may not be the same.

An example is shown in the following.

Example 1: Single iteration count avg

```
#pragma loop iterations avg=32
for(i = 0; i < m; i++) {
    a[i] = b[i] + c[i];
}
```

Optimized assuming average loop iteration count is 32.

Example 2: Multiple iteration count min and avg

```
#pragma loop iterations min=1 avg=8
for(i = 0; i < m; i++) {
    a[i] = b[i] + c[i];
}
```

Optimized assuming minimum loop iteration count is 1, and average loop iteration count is 8.

Example 3: Multiple iteration count max, avg, and min

```
#pragma loop iterations max=128 avg=16 min=16
for(i = 0; i < m; i++) {
    a[i] = b[i] + c[i];
}
```

Optimized assuming maximize loop iteration count is 128, average loop iteration count is 16, and minimum iteration count is 16.

#### loop\_blocking specifier

The loop\_blocking specifier instructs to apply the loop-blocking optimization. *n* is a decimal value ranging from 2 to 10000, and specifies the block size. If *n* is omitted, the compiler automatically determines a suitable value for *n*.

An example is shown in the following.

Example:

```

#define MIN(x,y) ((x<y)?x:y)
#pragma loop loop_blocking 80
for(i = 0; i < n; i++) {
    for(j = 0; j < n; j++) {
        a[i][j] = a[i][j] + b[i][j];
    }
}

```

->

```

for(ii = 0; ii < n; ii += 80) {
    for(jj = 0; jj < n; jj += 80) {
        for(i = ii; i < MIN(n,(ii+80)); i++) {
            for(j = jj; j < MIN(n,(jj+80)); j++) {
                a[i][j] = a[i][j] + b[i][j];
            }
        }
    }
}

```

The optimization of loop blocking by block size 80 is performed in the target loop.

#### loop\_noblocking specifier

The `loop_noblocking` specifier instructs to suppress the loop-blocking optimization.

An example is shown in the following.

Example:

```

#pragma loop loop_noblocking
for(i = 0; i < n; i++) {
    for(j = 0; j < n; j++) {
        a[i][j] = a[i][j] + b[i][j];
    }
}

```

The optimization of loop blocking by block size 80 is suppressed in the target loop.

#### loop\_fission\_stripmining specifier

The `loop_fission_stripmining` specifier instructs to perform the strip-mining optimization for automatic fissioned loops.

The strip length is specified by  $n$ , "L1", or "L2" following the specifier. If the strip length is omitted, the compiler chooses the strip length automatically.

When  $n$  is specified, the strip length is as  $n$ .  $n$  is a number from 2 to 100000000.

When "L1" is specified, the strip length is adjusted to the size of the first level cache for the cache memory efficiency.

When "L2" is specified, the strip length is adjusted to the size of the second level cache as well.

This specifier is effective when the `loop_fission_target` specifier is specified in the optimization control line, and the `-Kloop_fission` option and the `-O2` option or higher are set.

For details about strip-mining, see Section "3.3.10.1 Strip-Mining".

An example is shown in the following.

Example:

```

float a[N],b[N],p[N],q;
#pragma loop loop_fission_target
#pragma loop loop_fission_stripmining 256
for(i = 0; i < N; i++) {
    q = a[i] + b[i];
    ...
    p[i] = p[i] + q;
    ...
}

```

->

```

#define MIN(x,y) ((x<y)?x:y)
float a[N],b[N],p[N];
float temparray_q[256];
for(ii = 0; ii < N; ii=ii+256) {
    for(i = ii; i < MIN(N, ii+256); i++) {
        temparray_q[i-ii] = a[i] + b[i];
        ...
    }
    for(i = ii; i < min(N, ii+256); i++) {
        p[i] = p[i] + temparray_q[i-ii];
        ...
    }
}

```

### loop\_nofission\_stripmining specifier

The `loop_nofission_stripmining` specifier instructs to suppress the stripmining optimization for automatic fissioned loops.

An example is shown in the following.

Example:

```
float a[N], b[N], p[N], q;  
#pragma loop loop_fission_target  
#pragma loop loop_nofission_stripmining  
for(i = 0; i < N; i++) {  
    q = a[i] + b[i];  
    ...  
    p[i] = p[i] + q;  
    ...  
}
```

### loop\_fission\_target specifier

The `loop_fission_target` specifier instructs to perform automatic loop fission for target loop.

The `loop_fission_target c1` specifier instructs to perform loop fission by the clustering algorithm for target loop. This optimization performs the loop fission to prioritize reduction of work array for temporary data transfer by the loop fission. This optimization may increase the translation time.

The `loop_fission_target ls` specifier instructs to perform loop fission by the local search algorithm for target loop. This optimization performs the loop fission to prioritize promotion of the software pipelining. This optimization may increase the translation time longer than the clustering algorithm.

If neither `c1` nor `ls` is specified, `c1` is set by default.

For more details, refer to "[3.3.10 Loop Fission](#)".

Examples are shown in the following.

Example 1: `c1` specified

```
#pragma loop loop_fission_target c1  
for(i = 0; i < n; i++) {  
    s1 = a[i] + b[i];  
    s2 = c[i] + d[i];  
    ...  
    p[i] = s1 + q[i];  
    x[i] = s2 + y[i];  
    ...  
}
```

The loop is fissioned to prioritize reduction of work array for temporary data transfer by the loop fission.

Example 2: `ls` specified

```
#pragma loop loop_fission_target ls  
for(i = 0; i < n; i++) {  
    s1 = a[i] + b[i];  
    s2 = c[i] + d[i];  
    ...  
    p[i] = s1 + q[i];  
    x[i] = s2 + y[i];  
    ...  
}
```

The loop is fissioned to prioritize promotion of the software pipelining.

### loop\_fission\_threshold specifier

The `loop_fission_threshold` specifier specifies the threshold to decide the granularity of loop after automatic loop fission. *n* is a number from 1 to 100, and specifies the threshold to decide the granularity of loop after fission.

This specifier is effective when the `loop_fission_target` specifier is specified in the optimization control line, and the `-Kloop_fission` option and the `-O2` option or higher are set.

Examples are shown in the following. In the example 2 which specifying a small value for the `loop_fission_threshold` specifier, the compiler fissions the loop into small granularity compared with example 1 and the number of fissioned loops increase.

Example 1:

<pre>#pragma loop loop_fission_target #pragma loop loop_fission_threshold 50 for(i = 0; i &lt; N; i++) {     a1[i] = a1[i] + b1[i];     ...     a2[i] = a2[i] + b2[i];     ...     a3[i] = a3[i] + b3[i];     ...     a4[i] = a4[i] + b4[i];     ... }</pre>	->	<pre>for(i = 0; i &lt; N; i++) {     a1[i] = a1[i] + b1[i];     ...     a2[i] = a2[i] + b2[i];     ... } for(i = 0; i &lt; N; i++) {     a3[i] = a3[i] + b3[i];     ...     a4[i] = a4[i] + b4[i];     ... }</pre>
--	----	--

Example 2:

<pre>#pragma loop loop_fission_target #pragma loop loop_fission_threshold 20 for(i = 0; i &lt; N; i++) {     a1[i] = a1[i] + b1[i];     ...     a2[i] = a2[i] + b2[i];     ...     a3[i] = a3[i] + b3[i];     ...     a4[i] = a4[i] + b4[i];     ... }</pre>	->	<pre>for(i = 0; i &lt; N; i++) {     a1[i] = a1[i] + b1[i];     ... } for(i = 0; i &lt; N; i++) {     a2[i] = a2[i] + b2[i];     ... } for(i = 0; i &lt; N; i++) {     a3[i] = a3[i] + b3[i];     ... } for(i = 0; i &lt; N; i++) {     a4[i] = a4[i] + b4[i];     ... }</pre>
--	----	--

### loop\_interchange specifier

The `loop_interchange` specifier instructs the compiler to interchange the nested loops from the inner loops in order of the specified loop control variable names. As a result of this interchange, the execution speed may improve. However, when the interchange is impossible, this optimization is not applied.

An example is shown in the following.

Example:

<pre>#pragma loop loop_interchange j, i for(j = 0; j &lt; n; j++) {     for(i = 0; i &lt; m; i++) {         a[i][j] = b[i][j] * c[j][i];     } }</pre>	->	<pre>for(i = 0; i &lt; m; i++) {     for(j = 0; j &lt; n; j++) {         a[i][j] = b[i][j] * c[j][i];     } }</pre>
--	----	---

In left example, the nested loops are interchanged from the inner loop in order of the loop control variables `i` and `j`. The result of this interchange is equivalent to the loops shown in right.

### loop\_nointerchange specifier

The `loop_nointerchange` specifier instructs the compiler not to interchange the nested loops.

An example is shown in the following.

Example:

```
#pragma loop loop_nointerchange
for(j = 0; j < n; j++) {
    for(i = 0; i < m; i++) {
        a[i][j] = b[i][j] * c[j][i];
    }
}
```

The nested loops are not interchanged.

#### loop\_nofusion specifier

The `loop_nofusion` specifier instructs the compiler not to fuse the neighboring loop.

An example is shown in the following.

Example:

```
#pragma loop loop_nofusion
for(i = 0; i < n; i++) {
    i++;
}
for(j = 0; j < n; j++) {
    j++;
}
for(k = 0; k < n; k++) {
    k++;
}
```

Fusing of the first and second loops is suppressed. However fusing of the second and third loops are not suppressed.

#### loop\_part\_simd specifier

The `loop_part_simd` specifier instructs to perform optimization which divides loops and partially uses SIMD extensions.

An example is shown in the following.

Example:

```
#pragma loop loop_part_simd
for(i = 1; i < n; i++) {
    a[i] = a[i] - b[i] + log(c[i]); /* SIMD extensions can be applied to this statement */
    d[i] = d[i-1] + a[i];          /* SIMD extensions cannot be applied to this statement */
}
```

↓

```
#pragma loop loop_part_simd
for(i = 1; i < n; i++) {
    a[i] = a[i] - b[i] + log(c[i]); /* SIMD execution */
}
for(i = 1; i < n; i++) {
    d[i] = d[i-1] + a[i];          /* no SIMD execution */
}
```

The loop is divided and SIMD extensions are applied to the part of the divided loop.

#### loop\_nopart\_simd specifier

The `loop_nopart_simd` specifier instructs to suppress optimization which divides loops and partially uses SIMD extensions.

An example is shown in the following.

Example:



```

#pragma loop loop_nopart_simd
for(i = 1; i < n; i++) {
    a[i] = a[i] - b[i] + log(c[i]);
    d[i] = d[i-1] + a[i];
}

```

Neither loop dividing nor using SIMD extensions is applied to the loop.

#### loop\_perfect\_nest specifier

The `loop_perfect_nest` specifier instructs to perform optimization which fissions the imperfectly nested loop into the perfectly nested loops.

An example is shown in the following.

Example:

```

#pragma loop loop_perfect_nest
for(i = 0; i < n; i++) {
    /* imperfectly nested loop */
    a[i] = b[i] + 1;
    for(j = 0; j < n; j++) {
        c[j][i] = d[j][i] + a[i];
    }
}

```

->

```

for(i = 0; i < n; i++) {
    a[i] = b[i]+1;
}
for(i = 0; i < n; i++) {
    /* perfectly nested loop */
    for(j = 0; j < n; j++) {
        c[j][i] = d[j][i]+a[i];
    }
}

```

The imperfectly nested loops `i` and `j` are fissioned into the perfectly nested loops.

#### loop\_noperfect\_nest specifier

The `loop_noperfect_nest` specifier instructs to suppress optimization which fissions the imperfectly nested loop into the perfectly nested loops.

An example is shown in the following.

Example:

```

#pragma loop loop_noperfect_nest
for(i = 0; i < n; i++) { /* imperfectly nested loop*/
    a[i] = b[i] + 1;
    for(j = 0; j < n; j++) {
        c[j][i] = d[j][i] + a[i];
    }
}

```

The imperfectly nested loops `i` and `j` are not optimized into the perfectly nested loops.

#### loop\_versioning specifier

The `loop_versioning` specifier instructs to perform optimization by the loop versioning.

The loop versioning is applied to only an innermost loop, and also the loop contains only one array with unknown data dependency.

The following shows an example:

Example:

```

    for (i = 0; i < n; i++) {
#pragma loop loop_versioning
        for (j = 0; j < n; j++) {
            a[j] = a[j+m] + b[i][j];
        }
    }

```

The data dependency of array "a" is unknown at compile time because values of variable "n" and "m" are unknown at compile time.

When `loop_oversioning` is specified, the compiler generates two loops with "if" statement for judging the values of variable "n" and "m" so that the data dependency of the array "a" can be analyzed at execution time.

If the array "a" judges no data dependency, SIMD may be promoted for the loop.

#### `loop_oversioning` specifier

The `loop_oversioning` specifier instructs that loop versioning is suppressed.

The `loop_oversioning` specifier is effective only for the innermost loop.

The following shows an example:

Example:

```
for (i = 0; i < n; i++) {
#pragma loop loop_oversioning
    for (j = 0; j < n; j++) {
        a[j] = a[j+m] + b[i][j];
    }
}
```

#### `mfunc` specifier

The `mfunc` specifier instructs to convert the function in a loop to the multi-operation function.

The number 1, 2, or 3 following the specifier instructs the level of conversion to multi-operation function. If the level is omitted, 1 is applied. For more details about the levels of conversion, see the `-Kmfunc` option in Section "2.2.2.6 -K Option".

An example is shown in the following.

Example:

```
#pragma loop mfunc 1
for(i = 0; i < n; i++) {
    a[i] = log(b[i]);
}
```

The function `log(b[i])` is converted into multi-operation functions in the target loop.

#### `nomfunc` specifier

The `nomfunc` specifier instructs not to convert the intrinsic function in a loop to the multi-operation function.

An example is shown in the following.

Example:

```
#pragma loop nomfunc
for(i = 0; i < n; i++) {
    a[i] = log(b[i]);
}
```

The function `log(b[i])` is not converted into multi-operation functions in the target loop.

#### `noalias` specifier

The `noalias` specifier instructs that two or more pointer variables never point to the same memory area.

The compiler can assume from the `noalias` specifier that the pointer variables never point to the same memory area, however the memory area pointed to by the pointer variables is determined at execution.

An example is shown in the following. This promotes optimizations on the pointer variable. Note that this specifier may not promote the optimization when the value of the pointer variable is changed in the loop.

Example:

```
float *a;
float *b;
#pragma loop noalias
for (i=0; i<10; i++) {
```

```
b[i] = a[i] + 1.0;
}
```

 **Note**

If there are two or more pointer variables referring the same memory area in a loop with the `noalias` specifier, the execution result may be incorrect.

**norecurrence specifier**

The `norecurrence` specifier instructs that each element of the arrays which are specified by this optimization control specifier is defined and referenced within the limits of iteration.

This enables some optimizations to be performed for the loop because the order of definitions and references become certain.

For example, the following optimizations are applied:

- Loop Slicing (Automatic Parallelization)
- SIMD
- Software Pipelining

In the code example of "Example code without a `norecurrence` specifier", the system cannot determine whether array `a` can be safely sliced, because the subscript expression of array `a` is an array element `l[i]`, which is unknown at compile-time. Therefore, the system will not slice the outer loop.

**Example 1: Example code without a `norecurrence` specifier**

```
for(j = 0; j < 1000; j++) {
  for(i = 0; i < 1000; i++) {
    a[j][l[i]] = a[j][l[i]] + b[j][i];
  }
}
```


If the programmer knows that array `a` can safely be sliced, the outer loop will be sliced if the `norecurrence` specifier is used as shown in "Usage of the `norecurrence` specifier".

**Example 2: Usage of the `norecurrence` specifier**

---

```
#pragma loop norecurrence a

for(j = 0; j < 1000; j++) {
  for(i = 0; i < 1000; i++) {
    a[j][l[i]] = a[j][l[i]] + b[j][i];
  }
}
```

 parallel execution

---

 **Note**

Use of this specifier is the programmer's responsibility. If an array is specified in `norecurrence` specifier, but cannot safely be sliced, the results will be unpredictable.

**novrec specifier**

The `novrec` specifier instructs the compiler that there is no recurrent operation with array in the specified loop.

The SIMD instructions may be generated for the array looping using this specifier. However, the SIMD instructions may not be generated, depending on the type of operation and the loop structure.

When an array name is specified as an operand of the `novrec` specifier, it indicates that the array is excluded from the target of recurrent operation. When no operand is specified as the `novrec` specifier, it indicates that all arrays in the target loop are excluded from the target of recurrent operation.

An example is shown in the following.

Example:

```
double a[20], b[20];
...
#pragma loop novrec a
for(i = 1; i < 10; i++) {
    a[i] = a[i+n] + 1;
    b[i] = b[i] + 2;
}
```

The `novrec` specifier instructs that array `a` whose data dependency is unknown is not a target for recurrence operation. The operation of arrays `a` and `b` uses SIMD instructions.

#### `preex` specifier

The `preex` specifier instructs to optimize of evaluating the invariant first.

An example is shown in the following.

Example:

```
#pragma loop preex
for(i = 0; i < n; i++) {
    if(m[i] != 0) {
        a[i] = a[i] / b[k];
    }
}
```

->

```
t = 1 / b[k];
for(i = 0; i < n; i++) {
    if(m[i] != 0) {
        a[i] = a[i] * t;
    }
}
```

The optimization of evaluating the invariant first is performed in the target loop.

#### `nopreex` specifier

The `nopreex` specifier instructs to suppress the optimization of evaluating the invariant first. Note that a floating-point exception may occur when the `-Kfp_relaxed` option, `-NRtrap` option, and this specifier are valid.

An example is shown in the following.

Example:

```
#pragma loop nopreex
for(i = 0; i < n; i++) {
    if(m[i] != 0) {
        a[i] = a[i] / b[k];
    }
}
```

The optimization of evaluating the invariant first is suppressed in the target loop.

#### `prefetch` specifier

The `prefetch` specifier instructs to perform the automatic prefetch function. This function inserts the prefetch instruction in the optimal position to execute, as determined by the compiler.

An example is shown in the following.

Example:

```
#pragma loop prefetch
for(i = 0; i < 10; i++) {
```

```
a[i] = b[i];  
}
```

The automatic prefetch function is performed in the loop.

#### **noprefetch specifier**

The `noprefetch` specifier suppresses the automatic prefetch function.

An example is shown in the following.

Example:

```
#pragma loop noprefetch  
for(i = 0; i < 10; i++) {  
    a[i] = b[i];  
}
```

The automatic prefetch function is suppressed in the loop.

#### **prefetch\_cache\_level specifier**

The `prefetch_cache_level` specifier instructs the data cache-level of prefetch instructions.

When `prefetch_cache_level 1` specifier is used, prefetch instructions are generated to prefetch the data to the first level cache.

When `prefetch_cache_level 2` specifier is used, prefetch instructions are generated to prefetch the data to the second level cache.

When `prefetch_cache_level all` specifier is used, prefetch instructions are generated to prefetch the data to the first level and second level caches.

Examples are shown in the following.

Example 1:

```
#pragma loop prefetch_cache_level 1  
for(i = 0; i < 10; i++) {  
    a[i] = b[i];  
}
```

Prefetch instructions are generated to prefetch the data to the first level cache in the target loop.

Example 2:

```
#pragma loop prefetch_cache_level 2  
for(i = 0; i < 10; i++) {  
    a[i] = b[i];  
}
```

Prefetch instructions are generated to prefetch the data to the second level cache in the target loop.

Example 3:

```
#pragma loop prefetch_cache_level all  
for(i = 0; i < 10; i++) {  
    a[i] = b[i];  
}
```

Prefetch instructions are generated to prefetch the data to the first level and second level caches in the target loop.

#### **prefetch\_conditional specifier**

The `prefetch_conditional` specifier instructs to generate the prefetch instructions to array data used in blocks in "if" statements or "switch" statements.

An example is shown in the following.

Example:

```
#pragma loop prefetch_conditional
for(i = 0; i < n; i++) {
    if (x[i] = y[i]) {
        a[i] = b[i];
    }
}
```

The `prefetch_conditional` specifier generates the prefetch instructions to array data used in blocks in "if" statements.

#### `prefetch_noconditional` specifier

The `prefetch_noconditional` specifier instructs not to generate the prefetch instructions to array data used in blocks in "if" statements or "switch" statements.

An example is shown in the following.

Example:

```
#pragma loop prefetch_noconditional
for(i = 0; i < n; i++) {
    if (x[i] = y[i]) {
        a[i] = b[i];
    }
}
```

The `prefetch_noconditional` specifier does not generate the prefetch instructions to array data used in blocks in "if" statements.

#### `prefetch_indirect` specifier

The `prefetch_indirect` specifier instructs to generate the prefetch instruction to array data accessed indirectly (list access) in loops.

An example is shown in the following.

Example:

```
#pragma loop prefetch_indirect
for(i = 0; i < n; i++) {
    a[c[i]] = b[c[i]];
}
```

The `prefetch_indirect` specifier generates the prefetch instruction to array data accessed indirectly (list access) in loops.

#### `prefetch_noindirect` specifier

The `prefetch_noindirect` specifier instructs not to generate the prefetch instruction to array data accessed indirectly (list access) in loops.

An example is shown in the following.

Example:

```
#pragma loop prefetch_noindirect
for(i = 0; i < n; i++) {
    a[c[i]] = b[c[i]];
}
```

The `prefetch_noindirect` specifier does not generate the prefetch instruction to array data indirectly accessed (list access) in loops.

#### `prefetch_infer` specifier

The `prefetch_infer` specifier directs the system to generate the prefetch instruction for sequential access when the distance of prefetch is unknown.

An example is shown in the following.

Example:

```
#pragma loop prefetch_infer
for(i = 0; i < n; i++) {
    a[c[i]] = b[c[i]];
}
```

Prefetch instructions are generated regarding data access sequential in the target loop.

#### `prefetch_noinfer` specifier

The `prefetch_noinfer` specifier directs the system to generate the prefetch instruction without assuming sequential access when the distance of prefetch is unknown.

An example is shown in the following.

Example:

```
#pragma loop prefetch_noinfer
for(i = 0; i < n; i++) {
    a[c[i]] = b[c[i]];
}
```

Prefetch instructions are generated without regarding data access sequential in the target loop.

#### `prefetch_iteration` specifier

The `prefetch_iteration` specifier directs the system to be target of a prefetch instruction would be data that is referred to  $n$  iterations of a loop. This targets prefetch instructions only for the first level cache. Decimal from 1 to 10000 can be specified as  $n$ .

If the optimization using SIMD extensions is applied to a loop, specify the loop iteration count after using SIMD extensions for  $n$ .

An example is shown in the following.

Example:

```
#pragma loop prefetch_iteration 50
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Prefetch instructions of the first level cache generated in the target loop prefetch the data of 50 iterations.

#### `prefetch_iteration_L2` specifier

The `prefetch_iteration_L2` specifier directs the system to be target of a prefetch instruction would be data that is defined or is referred to  $n$  iterations of a loop. This targets prefetch instructions only for the second level cache. Decimal from 1 to 10000 can be specified as  $n$ .

If the optimization using SIMD extensions is applied to a loop, specify the loop iteration count after using SIMD extensions for  $n$ .

An example is shown in the following.

Example:

```
#pragma loop prefetch_iteration_L2 50
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Prefetch instructions of the second level cache generated in the target loop prefetch the data of 50 iterations.

#### `prefetch_line` specifier

The `prefetch_line` specifier instructs to target data after  $n$  cache line(s) for generating prefetch instruction to the first level cache. Decimal from 1 to 100 can be specified as  $n$ .

An example is shown in the following.

Example:

```
#pragma loop prefetch_line 5
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

The `prefetch_line` specifier targets data after 5 cache lines for generating prefetch instructions to first level cache.

#### `prefetch_line_L2` specifier

The `prefetch_line_L2` specifier instructs to target data after  $n$  cache line(s) for generating prefetch instruction to the second level cache. Decimal from 1 to 100 can be specified as  $n$ .

An example is shown in the following.

Example:

```
#pragma loop prefetch_line_L2 20
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

The `prefetch_line_L2` specifier targets data after 20 cache lines for generating prefetch instructions to second level cache.

#### `prefetch_sequential` specifier

The `prefetch_sequential` specifier directs that prefetch instructions are created for array data that is accessed sequentially.

When `prefetch_sequential auto` specifier is used, the compiler automatically selects whether to use hardware-prefetch or to create prefetch instructions for array data that is accessed sequentially in the loop.

When `prefetch_sequential soft` specifier is used, the compiler does not use hardware-prefetch, but rather creates prefetch instructions for array data that is accessed sequentially in the loop.

If neither `auto` nor `soft` is specified, `auto` is set by default.

Examples are shown in the following.

Example 1: `auto` specified

```
#pragma loop prefetch_sequential auto
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Hardware prefetch is used for the target loop rather than generating prefetch instructions.

Example 2: `soft` specified

```
#pragma loop prefetch_sequential soft
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Prefetch instructions are generated for the target loop rather than using hardware-prefetch.

#### `prefetch_nosequential` specifier

The `prefetch_nosequential` specifier instructs that prefetch instructions are not generated for array data that is accessed sequentially.

An example is shown in the following.

Example:

```
#pragma loop prefetch_nosequential
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```



Prefetch instructions are not generated for the target loop.

#### prefetch\_stride specifier

The `prefetch_stride` specifier instructs that prefetch instructions are generated for array data that is accessed with a stride larger than the cache line size used in a loop.

When `prefetch_stride soft` specifier is used, the compiler generates prefetch instructions and perform the prefetch.

When `prefetch_stride hard_auto` specifier is used, the compiler does not generate prefetch instructions, but rather use hardware prefetcher. This specifier sets the stride prefetcher to prefetch only data that is not on cache.

When `prefetch_stride hard_always` specifier is used, the compiler does not generate prefetch instructions, but rather use hardware prefetcher. In contrast with the `prefetch_stride hard_auto` specifier, this specifier sets the stride prefetcher to always prefetch.

If neither `soft`, `hard_auto` nor `hard_always` is specified, `soft` is set by default.

See Section "[3.6.3 Optimizations Using the Hardware Stride Prefetcher](#)" for the prefetch using the hardware stride prefetcher.

Examples are shown in the following.

##### Example 1: `soft` specified

```
#pragma loop prefetch_stride soft
for(i = 0; i < n; i=i+k) {
    a[i] = b[i];
}
```

Prefetch instructions are generated for the target loop.

##### Example 2: `hard_auto` specified

```
#pragma loop prefetch_stride hard_auto
for(i = 0; i < n; i=i+k) {
    a[i] = b[i];
}
```

Hardware prefetch is used for the target loop rather than generating prefetch instructions.

##### Example 3: `hard_always` specified

```
#pragma loop prefetch_stride hard_always
for(i = 0; i < n; i=i+k) {
    a[i] = b[i];
}
```

Hardware prefetch is used for the target loop rather than using prefetch instructions. Hardware prefetcher is set to always prefetch.

#### prefetch\_nostride specifier

The `prefetch_nostride` specifier directs that prefetch instructions are not generated for array data that is accessed with a stride larger than the cache line size used in a loop.

An example is shown in the following.

##### Example:

```
#pragma loop prefetch_nostride
for(i = 0; i < n; i=i+k) {
    a[i] = b[i];
}
```

Prefetch instructions are not generated for the target loop.

#### prefetch\_strong specifier

The `prefetch_strong` specifier directs that prefetch instructions for the first level cache are to be the strong prefetch.

An example is shown in the following.

Example:

```
#pragma loop prefetch_strong
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Prefetch instructions for the first level cache generated for the target loop will be strong prefetch.

**prefetch\_nostrong specifier**

The `prefetch_nostrong` specifier directs that prefetch instructions generated for the first level cache will not be strong prefetch.

An example is shown in the following.

Example:

```
#pragma loop prefetch_nostrong
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Prefetch instructions for the first level cache generated for the target loop will not be strong prefetch.

**prefetch\_strong\_L2 specifier**

The `prefetch_strong_L2` specifier directs that prefetch instructions for the second level cache are to be the strong prefetch.

An example is shown in the following.

Example:

```
#pragma loop prefetch_strong_L2
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Prefetch instructions for the second level cache generated for the target loop will be strong prefetch.

**prefetch\_nostrong\_L2 specifier**

The `prefetch_nostrong_L2` specifier directs that prefetch instructions generated for the first level cache will not be strong prefetch.

An example is shown in the following.

Example:

```
#pragma loop prefetch_nostrong_L2
for(i = 0; i < n; i++) {
    a[i] = b[i];
}
```

Prefetch instructions for the second level cache generated for the target loop will not be strong prefetch.

**preload specifier**

The `preload` specifier instructs to perform the speculative execution of load instructions.

An example is shown in the following.

Example:

```
#pragma loop preload
for(i = 0; i < n; i++) {
    if(m[i] > 0) {
        c[i] = a[i] + b[i];
    }
}
```

The speculative execution of load instructions is performed in the target loop.

#### nopreload specifier

The `nopreload` specifier instructs to suppress the speculative execution of load instructions.

An example is shown in the following.

Example:

```
#pragma loop nopreload
for(i = 0; i < n; i++) {
    if(m[i] > 0) {
        c[i] = a[i] + b[i];
    }
}
```

The speculative execution of load instructions is not performed in the target loop.

#### `scache_isolate_way` specifier

#### `end_scache_isolate_way` specifier

#### `scache_isolate_assign` specifier

#### `end_scache_isolate_assign` specifier

See Section "3.5 Software Control of Sector Cache".

#### `simd` specifier

The `simd` specifier instructs to perform the optimization that uses SIMD Extensions. However, SIMD Extensions may not be used depending on the type of operation and the loop structure.

"`simd`", "`simd aligned`", or "`simd unaligned`" can be specified as the `simd` specifier, but they all have the same effect. The "`aligned`" and "`unaligned`" parameters are deprecated for this product and left for backward compatibility with previous products only.

The `simd` specifier cannot be specified for the "`if-goto`" loop.

Examples are shown in the following.

Example:

```
double a[10], b[10];
...
#pragma loop simd
for(i = 0; i < 10; i++) {
    a[i] = a[i] + b[i];
}
```

The `simd` specifier is in effect for the data in the loop.

The loop structure is modified to adjust boundaries for using SIMD Extensions.

#### `nosimd` specifier

The `nosimd` specifier instructs to suppress the optimization that uses SIMD Extensions.

An example is shown in the following.

Example:

```
#pragma loop nosimd
for(i = 0; i < 10; i++) {
    a[i] = a[i] + b[i];
}
```

The data operation is not the subject of SIMD Extensions in the loop.

#### `simd_listv` specifier

The `simd_listv` specifier instructs to perform list vector conversion to the statements in the block of "`if`" statement.

When `then` is specified with `simd_listv`, list vector conversion is performed to the first substatement of the specified "if" statement. When `else` is specified with `simd_listv`, list vector conversion is performed to the second substatement of the specified "if" statement.

When `all` or `nothing` is specified with `simd_listv`, list vector conversion is performed to the first and second substatements of the specified "if" statement.

When the `-O2` option or higher is set, and the `-Ksimd=2` option or the same effect by the `simd` specifier is effective, this specifier is effective.

For details about list vector conversion, see Section "[3.2.7.3 List Vector Conversion](#)".

Examples are shown in the followings.

Example 1:

```
#pragma loop simd
for(i = 0; i < n; i++ ) {
#pragma statement simd_listv then
  if(a[i] == 0.0) {
    a[i] = a[i] + b[i];
  } else {
    a[i] = a[i] + c[i];
  }
}
```

The list vector conversion is performed to the first substatement of the specified "if" statement.

Example 2:

```
#pragma loop simd
for(i = 0; i < n; i++ ) {
#pragma statement simd_listv else
  if(a[i] == 0.0) {
    a[i] = a[i] + b[i];
  } else {
    a[i] = a[i] + c[i];
  }
}
```

The list vector conversion is performed to the second substatement of the specified "if" statement.

Example 3:

```
#pragma loop simd
for(i = 0; i < n; i++ ) {
#pragma statement simd_listv all
  if(a[i] == 0.0) {
    a[i] = a[i] + b[i];
  } else {
    a[i] = a[i] + c[i];
  }
}
```

The list vector conversion is performed to the both substatements of the specified "if" statement.

`simd_noredundant_vl` specifier

The `simd_noredundant_vl` specifier instructs to invalidate SIMD with redundant executions for the SIMD width. For details about SIMD with redundant executions for the SIMD width, see section "[3.2.7.4 SIMD with Redundant Executions for the SIMD Width](#)".

Examples are shown in the followings. This specifier directs only specific loop to suppress SIMD with redundant executions for the SIMD width as follows.

Example:

```
void foo(double *a, double *b, double *c) {
    ...
    #pragma loop simd_noredundant_vl
        for(i=0; i<m; i++) {
            a[i] = b[i]+c[i];
        }
    ...
}
```

`simd_use_multiple_structures` specifier

The `simd_use_multiple_structures` specifier instructs to use the SVE Load Multiple Structures instructions and the Store Multiple Structures instructions when using SIMD extensions.

This specifier is effective when the `-Ksimd[={1|2|auto}]` option or the `simd` specifier is effective, and the `-KSVE` option is set.

An example is shown in the following.

Example:

```
#include <complex.h>
#define N 10000
float _Complex a[N];
float _Complex b[N];
float _Complex c[N];
double y[N];
double x[N][4];
...
#pragma loop simd_use_multiple_structures
    for(int i=0; i<N; ++i) {
        a[i] = b[i] * c[i];
    }
...
#pragma loop simd_use_multiple_structures
    for(int i=0; i<N; ++i) {
        y[i] = x[i][0] + x[i][1] + x[i][2] + x[i][3];
    }
}
```

The SVE Load Multiple Structures instruction and the Store Multiple Structures instruction are used in this loop.

`simd_nouse_multiple_structures` specifier

The `simd_nouse_multiple_structures` specifier instructs not to use the SVE Load Multiple Structures instructions and the Store Multiple Structures instructions when using SIMD extensions.

An example is shown in the following.

Example:

```
#include <complex.h>
#define N 10000
float _Complex a[N];
float _Complex b[N];
float _Complex c[N];
double y[N];
double x[N][4];
...
#pragma loop simd_nouse_multiple_structures
    for(int i=0; i<N; ++i) {
        a[i] = b[i] * c[i];
    }
...
#pragma loop simd_nouse_multiple_structures
    for(int i=0; i<N; ++i) {
```

```
y[i] = x[i][0] + x[i][1] + x[i][2] + x[i][3];  
}
```

The SVE Load Multiple Structures instruction and the Store Multiple Structures instruction are not used in this loop.

#### striping specifier

The `striping` specifier directs that the loop striping optimization is performed. The number from 2 to 100 following the specifier specifies stripe length (number of expansions).

If the stripe length is omitted, 2 is applied.

When the iteration count of a loop in the source program is known, the expansion number determined by the compiler will be used if a number that exceeds the iteration count is specified for stripe length.

Examples are shown in the following.

#### Example 1:

```
#pragma loop striping  
for(i = 0; i < n; i++) {  
    statement;  
}
```

Loop striping is to be performed on the *statement* using the striping length determined by the compiler.

#### Example 2:

```
#pragma loop striping 4  
for(i = 0; i < n; i++) {  
    statement;  
}
```

Loop striping is to be performed on the *statement* using the striping length 4.

#### Example 3:

```
#pragma loop striping 8  
for(i = 0; i < 4; i++) {  
    statement;  
}
```

The specified stripe length (8) exceeds the loop iteration count (4), so the stripe length automatically determined by the compiler will be used.

#### Example 4:

```
void func() {  
#pragma procedure striping 8  
...  
    for(i = 0; i < n; i++) {  
        statement;  
    }  
...  
    for(j = 0; j < m; j++) {  
        statement;  
    }  
}
```

All the "for" loops in the program are targeted of loop striping.

#### nostriping specifier

The `nostriping` specifier directs that the loop striping optimization is suppressed.

An example is shown in the following.

Example:

```
#pragma loop nostriping
for(i = 0; i < n; i++) {
    statement;
}
```

Loop striping is suppressed.

#### swp specifier

The `swp` specifier instructs the compiler to perform software pipelining.

Note that the `swp` specifier is effective only if the `-O2` option or higher is set.

An example is shown in the following.

Example:

```
#pragma loop swp
for(i = 0; i < n; i++) {
    ...
}
```

Software pipelining is performed to the specified loop.

#### noswp specifier

The `noswp` specifier instructs the compiler to suppress software pipelining.

An example is shown in the following.

Example:

```
#pragma loop noswp
for(i = 0; i < n; i++) {
    ...
}
```

Software pipelining is suppressed in the target loop.

#### swp\_freg\_rate specifier

#### swp\_ireg\_rate specifier

#### swp\_preg\_rate specifier

Specifies the rate (percentage) about the following registers that can be used by software pipelining.

- Floating-point register and SVE vector register
- Integer register
- SVE predicate register

When specifying to `swp_freg_rate` specifier, the rate about floating-point register and SVE vector register is adjusted.

When specifying to `swp_ireg_rate` specifier, the rate about integer register is adjusted.

When specifying to `swp_preg_rate` specifier, the rate about SVE predicate register is adjusted.

The number from 1 to 1000 following specifiers instructs the rate (percentage) of the number of registers which is assumed to be usable by the software pipelining. This specifier is effective only if the `-O2` option or higher is set.

An example is shown in the following.

Example:

```
#pragma loop swp_freg_rate 120
#pragma loop swp_ireg_rate 150
#pragma loop swp_preg_rate 80
for(i = 0; i < n; i++) {
```

```
} ...
```

Perform the optimization of the software pipelining which adjusts the condition about the number of each kind of register for the target loop.

#### swp\_policy specifier

Specifies a policy to select an instruction scheduling algorithm used in software pipelining.

Either of specifiers shown below is instructed after `swp_policy` specifier.

`auto`

An algorithm fit for a small loop, such as a loop with low register pressure, is used.

`small`

An algorithm fit for a small loop, such as a loop with low register pressure, is used.

`large`

An algorithm fit for a large loop, such as a loop with high register pressure, is used.

#### swp\_weak specifier

The `swp_weak` specifier instructs the compiler to adjust software pipelining to make smaller overlapping of the instructions. This specifier is effective only if the `-O2` option or higher is set.

An example is shown in the following.

Example:

```
#pragma loop swp_weak
for(i = 0; i < n; i++) {
    ...
}
```

Perform the optimization of the software pipelining which adjust software pipelining for the target loop and make smaller overlapping of the instructions.

#### unroll specifier

The `unroll` specifier instructs the compiler to perform the optimization of loop unrolling for the corresponding loop.

The number from 2 to 100 following the specifier instructs the upper bound on the number of loops to be unrolled. If the upper bound of the number is omitted, the compiler automatically determines a suitable value.

When the `unroll "full"` specifier is used, the statement is unrolled up to the number of the iteration in the specified "for" statement. If the number of the iteration is unknown at compilation, no optimization of loop unrolling is performed.

Note that the `unroll` specifier targets only the loop specified immediately after.

An example is shown in the following.

Example:

```
#pragma loop unroll 8
for(i = 0; i < n; i++) {
    statement;
}
```

The "`statement`" is unrolled 8 times in the target loop.

#### nounroll specifier

The `nounroll` specifier instructs the compiler to suppress unrolling for the corresponding loop.

Note that the `nounroll` specifier targets only the loop specified immediately after.

An example is shown in the following.



Example:

```
#pragma loop nounroll
for(i = 0; i < n; i++) {
    ...
}
```

Loop unrolling is suppressed in the target loop.

#### `unroll_and_jam` specifier

The `unroll_and_jam` specifier instructs the compiler to apply unroll-and-jam. However, the optimization is suppressed in the following case:

- Assuming that the optimization is not effective.
- Assuming that there is a data dependency over iterations of the loops.

The number from 2 to 100 following the specifier instructs the upper limit of the unrolling expansion number. If the upper limit is omitted, the compiler automatically determines a value.

This optimization is not applied to innermost loop.

This specifier is effective only if the `-O2` option or higher is set.

The following shows examples:

Example 1:

```
#pragma loop unroll_and_jam 2
for (i=0; i<128; i++) {
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        ...
    }
}
```

->

```
for (i=0; i<128; i+=2) {
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        a[i+1][j] = b[i+1][j] + b[i+2][j];
        ...
    }
}
```

Unroll-and-jam is applied to the loop whose control variable is `i`.

Example 2:

```
for (i=0; i<128; i++) {
    #pragma loop unroll_and_jam 2
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        ...
    }
}
```

Unroll-and-jam is not applied because the specified loop is the innermost.

Example 3:

```
#pragma loop unroll_and_jam 2
for (i=0; i<128; i++) {
    for (j=0; j<127; j++) {
        a[i][j] = a[i-1][j+1] + b[i][j];
        ...
    }
}
```

Unroll-and-jam is not applied because an array "a" has a data dependency over iterations of the loops.

#### `unroll_and_jam_force` specifier

The `unroll_and_jam_force` specifier instructs the compiler to perform the unroll-and-jam optimization for the corresponding loop assuming that there is no data dependency over iterations of the loops.

The result of executions is not guaranteed if the specifier is used incorrectly. For details, see the `unroll_and_jam_force` specifier in Section "3.4.1.3 Notes for Optimization Control Specifiers".

The number from 2 to 100 following the specifier instructs the upper limit of the unrolling expansion number. If the upper limit is omitted, the compiler automatically determines a value.

Note that the `unroll_and_jam_force` specifier targets only the loop immediately after the optimization control line. If the loop is the innermost, the unroll-and-jam is not applied.

This specifier is effective only if the `-O2` option or higher is set.

The following shows examples:

Example 1:

```
#pragma loop unroll_and_jam_force 2
for (i=0; i<128; i++) {
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        ...
    }
}

->

for (i=0; i<128; i+=2) {
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        a[i+1][j] = b[i+1][j] + b[i+2][j];
        ...
    }
}
```

Unroll-and-jam is applied to the loop whose control variable is `i`.

Example 2:

```
for (i=0; i<128; i++) {
    #pragma loop unroll_and_jam_force 2
    for (j=0; j<128; j++) {
        a[i][j] = b[i][j] + b[i+1][j];
        ...
    }
}
```

Unroll-and-jam is not applied because the specified loop is the innermost.

This specifier instructs that unroll-and-jam is performed assuming that there is no data dependency over iterations of the loops. The execution result is not guaranteed if the specifier is used incorrectly. See Section "3.4.1.3 Notes for Optimization Control Specifiers".

#### unswitching specifier

The `unswitching` specifier instructs loop unswitching to "if" statement.

This specifier must be described immediately before "if" statement which is invariant in a loop. This specifier is not effective in the case of specified other than the above.

The following shows an example:

Example 1:

```
void foo(double *a, double *b, double *c, int x, int n) {
    int i;
    for (i=0; i<n; i++) {
        #pragma statement unswitching
        if (x == 0) {
            a[i] = b[i];
        } else {
            a[i] = c[i];
        }
    }
}
```

Loop unswitching to "if" statement is performed.

Note that "if" statement without this specifier is not the target of loop unswitching.

The following shows an example:

### Example 2:

```
void foo(double *a, double *b, double *c, double *d, int x, int n) {
    int i;
    for (i=0; i<n; i++) {
        if (x == 0) {
            a[i] = b[i];
#pragma statement unswitching
        } else if (x == 1) {
            a[i] = c[i];
        } else {
            a[i] = d[i];
        }
    }
}
```

Only "if" statement with this specifier is the target of loop unswitching.

The memory consumption and the compilation time may increase drastically when loops to which loop unswitching is applied include a lot of execution statements.

### zfill specifier

The `zfill` specifier directs to apply the `zfill` optimization. The `zfill` optimization speeds up write operations for array data that is only written in a loop, by using an instruction that allocates space on the cache for writing without loading data from the memory. The `zfill` optimization works on the data  $N$  blocks ahead of the address pointed to by the target store instruction where one block is 256 byte-long and  $N$  is an integer value between 1 and 100. If a value is not specified for  $N$ , the compiler will automatically determine a value.

For details about `zfill`, see "3.3.5 `zfill`".

Examples are shown in the following.

### Example 1:

```
#pragma loop zfill
for(i = 0; i < n; i++) {
    ...
}
```

This specifies to improve the cache utilization for the data on certain blocks ahead, and the number is determined by the compiler.

### Example 2:

```
#pragma loop zfill 1
for(i = 0; i < n; i++) {
    ...
}
```

This specifies to improve the cache utilization for the data on 1 block ahead.

### nozfill specifier

The `nozfill` specifier directs not to apply the `zfill` optimization.

An example is shown in the following.

### Example:

```
#pragma loop nozfill
for(i = 0; i < n; i++) {
    ...
}
```

This specifies that the `zfill` optimization will not be performed.

### 3.4.1.3 Notes for Optimization Control Specifiers

#### clone specifier

In the copied loops under the generated branches, the specified variables are treated as invariants in the loops. Thus, the execution result is not guaranteed in case that the value of the variables changes in the loop.

##### Example 1:

Do not specify a variable which is changed in the loop because the execution result is not guaranteed.

```
int *M = &N;
#pragma loop clone N==10
for(i=0;i<32;++i) {
    *M = 5;
    a[i] = N;
}
```

##### Example 2:

Do not specify a variable which may be changed in the loop because the execution result is not guaranteed.

```
#pragma loop clone N==10
for(i=0;i<32;++i) {
    sub(&N);
}
```

#### novrec specifier

If the novrec specifier is specified, the compiler generates the SIMD instruction given that there is no array for recurrent operation in the loop.

However, if there are some arrays for recurrent operation in the loop, the compiler uses SIMD instructions, but the execution result will be unpredictable.

An example is shown in the following.

##### Example:

```
#pragma loop novrec
for(i = 0; i < 100; i++) {
    a[i] = a[i + m] + ...;    <- When array 'a' is recurrence data and novrec specifier is
                             effective, the compiler uses SIMD instructions.
}
```

This is illegal usage of novrec specifier.

#### unroll\_and\_jam\_force specifier

The unroll\_and\_jam\_force specifier instructs that unroll-and-jam is performed assuming that there is no data dependency over iterations of the loops. The result of executions is not guaranteed if the specifier is used incorrectly.

##### Example:

```
#pragma loop unroll_and_jam_force 2
for (i=1; i<128; i++) {
    for (j=0; j<127; j++) {
        a[i][j] = a[i-1][j+1] + b[i][j];
        ...
    }
}
```

Do not use the unroll\_and\_jam\_force specifier because there exists data dependency of array a over iterations of the loops.

#### iterations max=n1 specifier

#### iterations min=n3 specifier

The iterations max=n1 and iterations min=n3 specifiers are effective when the loop iteration count is unknown at compilation, and optimize assuming the maximum loop iteration count is n1 and the minimum loop iteration count is n3.

The result of executions is not guaranteed if the actual loop iteration count is greater than iteration count  $n1$ , or less than iteration count  $n3$ .

Example:

```
#pragma iterations max=100 min=1
for (i=1; i<m; i++) { // The actual maximum count is 1000 or minimum count is 0.
    a[i] = a[i] + b[i];
}
```

When the actual maximum iteration loop count is greater than 100, or minimum iteration loop count is 0, the execution result is not guaranteed.

## 3.5 Software Control of Sector Cache

---

### 3.5.1 Use of Sector Cache

---

Sector cache is a mechanism to prevent reusable data on the cache from being driven out by non-reusable data. In particular, Sector cache is effective for parallel execution when the shared L2 cache is accessed by multiple cores. By using Sector cache, reusable data is separated from non-reusable data on the cache. There are two methods to control Sector caches with software; environment variables or optimization control lines (OCL). The environment variables can specify the maximum number of ways for each sector, while the OCL can specify that of sector 1.

Performance will not be improved if the maximum number of ways is specified without consideration for the size of the array that is to be protected from being driven out. The cache utilization will decline, and may cause a reduction in speed. Further, even if software control of the Sector cache is not used, cache control with LRU (Least Recently Used) still works, so speed may not improve even if this is specified.

One effective way to improve the performance of Sector caches is to determine the reusable array size and specify the number of ways for sector 1 to store the array, after ensuring L2 cache miss has occurred by confirming the Hardware Monitor Information. Note that in order to use Sector caches, the compile-time option `-Khpctag` must be set.



#### Information

The maximum number of cache way for each sector

In the A64FX processor, the maximum number of the first level cache is 4, and that of the second level cache is 16.

### 3.5.2 Controlling Sector Cache via Software

---

The following two methods are provided to control Sector cache via software.

- Optimization Control Line (OCL).
- Environment variables and Optimization Control Line.

#### 3.5.2.1 Software Control with Optimization Control Lines

The following OCLs exist to control Sector cache.

```
scache_isolate_way L2= $n1$  [L1= $n2$ ]  
and  
end_scache_isolate_way
```

The `scache_isolate_way` specifier directs the maximum number of ways in sector 1 of the caches.

The argument of the `scache_isolate_way` specifier directs the maximum number of ways of sector 1 of the cache. The `L2= $n1$`  directs the maximum number of ways of sector 1 in the second level cache and the `L1= $n2$`  directs the maximum number of ways of sector 1 in the first level cache. Specifying the `L1= $n2$`  is optional. In this case, this specifier controls the maximum number of ways of sector 1 of the second level cache.

The arguments  $n1$  and  $n2$  must be as follows;

- $0 \leq n1 \leq$  "the maximum number of ways in the second level cache"
- $0 \leq n2 \leq$  "the maximum number of ways in the first level cache"

To direct in function, the `scache_isolate_way` specifier must be used in procedure line. In this case, the `scache_isolate_way` specifier is effective in the scope of the function.

To direct a part of function, the range must be indicated by the `scache_isolate_way` and `end_scache_isolate_way` specifiers in statement lines.

Note that the range must not be nested. However, combined use of procedure line and statement line is accepted.

```
scache_isolate_assign array1[,array2]...
and
end_scache_isolate_assign
```

The `scache_isolate_assign` specifier directs the array and pointer to the array to store in sector 1 of cache. However, it becomes effective only when the array of arithmetic type or pointer is specified.

To direct in function, the `scache_isolate_assign` specifier must be used in procedure line. In this case, the `scache_isolate_assign` specifier is effective in the scope of the function.

To direct a part of function, the range must be indicated by the `scache_isolate_assign` and `end_scache_isolate_assign` specifiers in statement lines.

Note that the range must not be nested. However, combined use of procedure line and statement line is accepted.

### Example

Example of software method to control Sector cache

```
/* Reuse the array a that can fit in 10 ways in the L2 cache */
#pragma statement scache_isolate_way L2=10
#pragma statement scache_isolate_assign a
for(int j = 0; j < n; j++) {
#pragma omp parallel for
    for(int i = 0; i < m; i++) {
        a[i] = a[i] + b[j][i];
    }
}
#pragma statement end_scache_isolate_assign
#pragma statement end_scache_isolate_way
```

### 3.5.2.2 Software Control with Environment Variables and Optimization Control Line

The initial value for the maximum number of ways for each Sector can also be specified with the following environment variables. The maximum number of ways for the object program at startup can be specified with environment variables.

`FLIB_SCCR_CNTL`

This environment variable specifies to use Sector cache. Allowed values for the variable are shown below.

Value	Explanation
TRUE	Use Sector cache. The default is TRUE.
FALSE	Do not use Sector cache.

`FLIB_L1_SCCR_CNTL`

The Sector cache for the second level cache is not available when multiple processes run on a NUMA node. This environment variable specifies whether to use the Sector cache for the first level cache when the Sector cache for the second level cache is unavailable. `FLIB_L1_SCCR_CNRTL` is effective only when `FLIB_SCCR_CNTL` is assigned to the value TRUE.

The values that can be set in the environment variable are shown below, along with their meanings.

Value	Explanation
TRUE	If the Sector cache for the second level cache is not available, the Sector cache for the first level cache is used. The value is set by default.
FALSE	If the Sector cache for the second level cache is not available, the Sector cache for the first level cache is not used. And the following error message is output, the Sector cache control function is disabled, and program execution continues. <pre>jwe1047i-w A sector cache couldn't be used.</pre>

#### FLIB\_L2\_SECTOR\_NWAYS\_INIT

This environment variable specifies the initial maximum number of cache ways for each sector of the second level cache. FLIB\_L2\_SECTOR\_NWAYS\_INIT is effective only when the value of FLIB\_SCCR\_CNTL is TRUE.

The values for the first sector (sector 0) *n0* and for the second sector (sector 1) *n1* should be specified as follows:

```
n0, n1
```

The possible values are as follows:

- 0 <= *n0* <= "the maximum number of ways in the second level cache"
- 0 <= *n1* <= "the maximum number of ways in the second level cache"

It is recommended to satisfy the following condition to avoid conflicts.

- $n0 + n1 =$  "the maximum number of ways in the second level cache"



### Example

Example of software method to control Sector cache with environment variables and optimization control line

As shown in the following example of software control of Sector cache, by setting the value "2,10" to the environment variable FLIB\_L2\_SECTOR\_NWAYS\_INIT before execution starts, the maximum number of ways is set to 10 and the array data is cached on the sector 1 cache.

1. Set environment variables (bash)

```
FLIB_SCCR_CNTL=TRUE
export FLIB_SCCR_CNTL
FLIB_L2_SECTOR_NWAYS_INIT=2,10
export FLIB_L2_SECTOR_NWAYS_INIT
```

2. Specify an array data stored in sector 1 by optimization control lines

```
/* Assign sector 1 to the array a */
#pragma statement scache_isolate_assign a
for(int j = 0; j < n; j++) {
#pragma omp parralel for
    for(int i = 0; i < m; i++) {
        a[i] = a[i] + b[j][i];
    }
}
#pragma statement end_scache_isolate_assign
```

### 3.5.2.3 Behavior when an Exceptional Value Is Specified

If a value specified for the `scache_isolate_way` specifier or the environment variable `FLIB_L2_SECTOR_NWAYS_INIT` exceeds the upper limit, it is assumed that the upper limit is specified. If a value less than 0 is specified, it has no effect. If a value outside the range of two-byte signed integer is specified, the behavior is undefined.

## 3.6 Notes



### 3.6.1 Side Effect of Optimizations for Floating-Point Operation

Optimizations for floating-point operation might cause the side effect. This section explains the side effect (mainly, computation error).

This system basically creates objects which comply with IEEE 754 arithmetic. Note that the numerical operations may not comply with IEEE 754 arithmetic due to the optimizations with calculation errors in "Table 3.23 Side effect of optimization for floating-point operation".

See Section "2.2.2.6 -K Option" for compiler options. See section "3.4.1.2 Optimization Control Specifier" and "4.2.6.3 Optimization Control Specifiers for Automatic Parallelization" for optimization control specifiers.

Table 3.23 Side effect of optimization for floating-point operation

Compiler Option	Optimization Control Specifier	Side Effect
-Keval	eval	<p>Calculation errors may occur when optimization that changes the method of operator evaluation.</p> <p> <b>Example</b></p> <hr/> <pre>x*y + x*z -&gt; x*(y+z)</pre> <hr/> <p>When the -Keval option is set, the following options are valid. See the description about the side effect of each option.</p> <ul style="list-style-type: none"> <li>- -Kfsimple option</li> <li>- -Kreduction option (When -Kparallel option is valid)</li> <li>- -Ksimd_reduction_product option (When -Ksimd[={1 2 auto}] option is valid)</li> </ul>
-Kfsimple	-	<p>Simplification of floating point operation on source programs is performed. Therefore, the numerical operations do not comply with IEEE 754 arithmetic, as in the example below.</p> <p> <b>Example</b></p> <hr/> <pre>x*0.0 -&gt; 0.0</pre> <hr/> <p>Operations such as "x*0.0" will be simplified to "0.0".</p>
-Kfp_contract	fp_contract	<p>Rounding errors may occur when optimization using Floating-Point Multiply-Add/Subtract instructions is performed on source programs.</p>
-Kfp_relaxed	fp_relaxed	<p>Side effects may occur because reciprocal approximation operation instructions are used on single-precision or double-precision floating point division or <code>sqrt</code> functions.</p> <p>The side effects that may occur are:</p> <ul style="list-style-type: none"> <li>- Rounding errors.</li> </ul>



Compiler Option	Optimization Control Specifier	Side Effect
		<ul style="list-style-type: none"> <li>- Replacing denormalized numbers found in the arguments or the return value with zero regardless of set of -Kfz option.</li> <li>- Replacing negative zeroes found in the arguments or the return value with positive zeroes.</li> <li>- Behaviors not conforming to IEEE 754 when NaN, positive or negative Inf, numbers which are close to maximal normalized number or numbers which are close to minimal normalized number are found in the arguments or the return value.</li> </ul>
-Kfz	-	Flush-to-zero mode is used. If a result or a source operand is a denormalized number, flush-to-zero mode replaces it with zero with the same sign.
-Kfast_matmul	-	Calculation errors may occur when using high speed library call for the loop of matrix multiplication.
-Kilfunc	-	Side effects similar to the ones caused by -Kfp_relaxed may occur when using inline expansion for the math functions, because reciprocal approximation instructions and trigonometric instructions, etc. are used. Plus, use of reciprocal approximation instructions or Floating-Point Multiply-Add/Subtract instructions regardless of set of -Knofp_contract and/or -Knofp_relaxed options or their setting order.
-Kmfunc[={1 2 3}]	mfunc [ <i>level</i> ]	Side effects similar to the ones caused by -Kfp_relaxed and/or -Kilfunc may occur because different algorithms, reciprocal approximation instructions, trigonometric instructions, etc. are used internally when the function is converted to the multi-operation function. See Section " <a href="#">3.3.3 Multi-Operation Function</a> ".
-Kreduction	reduction	The automatic parallelization loop reduction is performed. Therefore, calculation errors may occur when the optimization that changes the method of operator evaluation.
-Ksimd_reduction_product	-	SIMD extensions are used to the reduction operation of product. Therefore, calculation errors may occur when the optimization that changes the method of operator evaluation.
-Kfast	-	-Kfast option induces the -Keval option, -Kilfunc option, and so on.
-Kvisimpact	-	-Kvisimpact option induces the -Kfast option.

-: None

If the -O1 option or higher is set, this system performs optimizations that execute floating-point arithmetic instructions that may not be executed based on the logic of the program. This may result in generating extra floating-point exceptions at execution time. In this case, the following phenomena may occur.

- If a program uses the fetestexcept function included in standard library functions and accesses a floating-point status flag by it, the flag which should not be set based on the logic of the program is set.
- If a program uses the feenableexcept function included in the GNU C Library and enables a trap of a floating-point exception by it, the SIGFPE signal which should not be raised based on the logic of the program is raised.

This phenomenon may be avoided by setting the option -NRtrap.



## Information

To avoid side effect for floating-point operations error at parallelization (automatic parallelization/OpenMP)

Specify the following options to avoid calculation error of a real type or a complex type operation that is caused by the difference of the parallel number of threads. When these options are set, the execution performance may decrease.

- `-Kparallel_fp_precision` option (When the `-Kparallel` option or `-Kopenmp` option is valid.)
- `-Kopenmp_ordered_reduction` option (When the `-Kopenmp` option is valid.)

See Section "[2.2.2.6 -K Option](#)" for compiler options.

---

## 3.6.2 Notes on Specified SVE Vector Register Size

---

`-Ksimd_reg_size={ 128|256|512 }` specifies the bitwise SVE vector register size. When this option is specified, the optimization is performed considering the vector register size as a fixed value at compilation. Therefore, the generated executable program works normally on the CPU architecture which has the same size of the SVE vector register as the specified size.

When the program is executed on the CPU whose implemented vector register size is different from the vector register size specified by `-Ksimd_reg_size={ 128|256|512 }` option, an abnormal end occurs. And, the result of executions is not guaranteed.

It is required to change the effective size of the vector register by the system call `prctl ( 2 )` or in other ways when the vector register size specified by `-Ksimd_reg_size={ 128|256|512 }` option is smaller than that of the CPU.

When `-Ksimd_reg_size=agnostic` is effective, the executable program does not depend on the SVE vector register size.

See also Section "[B.2.9 Changing SVE Vector Register Size](#)" for notes on changing the SVE vector register size.

## 3.6.3 Optimizations Using the Hardware Stride Prefetcher

---

The `-Kprefetch_stride=hard_auto` option, `-Kprefetch_stride=hard_always` option, `prefetch_stride hard_auto` specifier, and `prefetch_stride hard_always` specifier use the hardware stride prefetcher at runtime.

To use it, configuration of the hardware prefetcher should be permitted on the system. If it is not permitted, these optimizations do not take effect.

Contact the system administrator to know permission setting.

## 3.6.4 Notes of Wrong Erroneous Program

---

When optimization is performed on the following erroneous programs, execution results may differ with the optimization level. If results differ, check the following and correct the programs.

- Variables containing undefined values are referenced
- Variables containing values with unacceptable formats are referenced
- Array elements are referenced using index values that exceed array declarations
- There are violations of C++ language syntax

# Chapter 4 Multiprocessing

This chapter describes the method of processing a C++ program in parallel or "multiprocessing" using LLVM OpenMP Library.

See "[Appendix J Fujitsu OpenMP Library](#)" when you use Fujitsu OpenMP Library.

## Information

The table below shows a connectable combination of library for multiprocessing (LLVM OpenMP Library and Fujitsu OpenMP Library) and connectable object file.

	Fortran object file	C/C++ object file	
		Trad Mode (-Nnoclang option)	Clang Mode (-Nclang option)
LLVM OpenMP Library (-Nlibomp option)	Available	Available	Available
Fujitsu OpenMP Library (-Nfjomplib option)	Available	Available	Not available

## 4.1 Overview of Multiprocessing

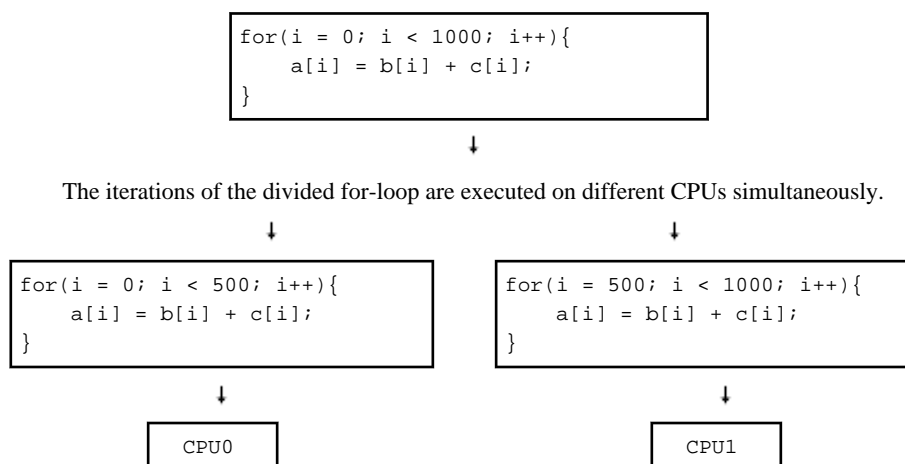
This section describes multiprocessing of a C++ program and parallelization using LLVM OpenMP Library.

### 4.1.1 What is Multiprocessing?

In this document, multiprocessing means that one program is executed on two or more CPUs that can work independently at the same time. As used here, it does not mean executing two or more programs simultaneously.

"[Figure 4.1 Multiprocessing](#)" illustrates multiprocessing by using two CPUs simultaneously.

Figure 4.1 Multiprocessing




## 4.1.2 Effect of Multiprocessing

---

The effect of multiprocessing is to reduce elapsed execution time by using two or more CPUs simultaneously. For instance, if a `for` loop can be executed in parallel by dividing it as shown in "Figure 4.1 Multiprocessing", the execution time of this `for` loop may be cut in half (See "Figure 4.2 Reducing elapsed time using multiprocessing").

Figure 4.2 Reducing elapsed time using multiprocessing

---



Although the elapsed time will usually decrease with multiprocessing, the total CPU time required to execute the program may increase. This is because the total CPU time is at least as large as the CPU time when the program is executed on a single processor, and the overhead time for multiprocessing may increase the total CPU time.

The performance of a multiprocessing program is usually evaluated by the reduction of the elapsed time. The total CPU time of a program using multiprocessing is larger than that required by serial processing. This is because of the overhead time required by multiprocessing.

## 4.1.3 Requirements for Effective Multiprocessing

---

A computer environment that can use two or more CPUs simultaneously is necessary to reduce elapsed time by multiprocessing. A multiprocessing program can be executed on hardware with only a single CPU; however, the elapsed time will not be less than the execution time for a comparable program written without multiprocessing features. Moreover, even if the program is executed on hardware with two or more CPUs, shortening the elapsed time is difficult when other jobs are executing on the same CPUs. The reason for this is that the probability that two or more CPUs can be allocated at the same time decreases.

To achieve effective multiprocessing, it is necessary to run a multiprocessing program on hardware with multiple CPUs and within a system environment that has room for multiprocessing.

To reduce the relative overhead of multiprocessing, it is necessary that a loop have many iterations and many statements in its body.

## 4.1.4 Parallelization

---

This system offers automatic parallelization. "`for`" loops, "`while`" loops, "`do-while`" loops and "`if-goto`" loops are parallelized without any modifications made to the program. Therefore, it is easy to migrate source programs to other processing systems as long as the program conforms to the C++ standard.

However, the automatic parallelization is not applied to the loops that the control jumps outside of the loop or jumps in from outside of the loop.

Moreover, this system offers the optimization control line (OCL), which helps automatic parallelization. The optimization control line is used by the programmer to identify constructs that may be executed in parallel.

For more information about automatic parallelization, see Section "4.2 Automatic Parallelization".

This system also offers parallelization by the OpenMP specification, see Section "4.3 Parallelization by OpenMP Specification" for details.

## 4.2 Automatic Parallelization

---

This section describes automatic parallelization.

## 4.2.1 Compilation

---

To activate automatic parallelization, specify the following options:

```
-Kparallel [-Nlibomp]
```

### 4.2.1.1 Compiler Option for Automatic Parallelization

The compiler option for automatic parallelization is:

```
-K{parallel|parallel_strong|visimpact}  
[,array_private,dynamic_iteration,instance=N,loop_part_parallel,ocl,optmsg=2,parallel_fp_precision,parallel_iteration=N,reduction,region_extension] [-Nlibomp]
```

See Section "2.2 Compiler Options" for details on how to use each option.

## 4.2.2 Execution Process

---

When a multiprocessing program is executed, the number of threads to be used may be specified by the environment variable `OMP_NUM_THREADS`. The size of the thread stack may be specified by the environment variable `OMP_STACKSIZE`. The method of synchronizing threads can be changed using the environment variable `OMP_WAIT_POLICY`.

Except these environment variable settings, the procedure for execution is same as for a serial-processing program.

### 4.2.2.1 Number of Threads

The number of threads is determined with the following priority.

1. The value of the environment variable `OMP_NUM_THREADS`
2. The number of CPUs that can be used in the system

The following environment variable for OpenMP specifications can be used in automatic parallelization. See OpenMP specifications for the detail.

`OMP_NUM_THREADS`

Sets the number of threads to use during execution.

### 4.2.2.2 Stack Size on Execution

The stack area size for each thread can be specified using the following environment variable.

`OMP_STACKSIZE`

The stack area size for each thread can be specified using the environment variable `OMP_STACKSIZE` in byte, Kbytes, Mbytes, Gbytes, or Tbytes. The default size is 8M bytes.

### 4.2.2.3 Synchronization Process

The synchronization process can be controlled with the environment variable `OMP_WAIT_POLICY`.

`OMP_WAIT_POLICY`

ACTIVE	Uses spin wait until all threads are synchronized.
PASSIVE	Uses no spin wait but suspending wait.

The default value is "PASSIVE".

Select ACTIVE to give priority to the elapsed time. Select PASSIVE to give priority to the CPU time.

## 4.2.2.4 CPU Binding for Thread

Environmental variable `GOMP_CPU_AFFINITY` or `OMP_PROC_BIND` can control the CPU binding of threads. `GOMP_CPU_AFFINITY` takes precedence over `OMP_PROC_BIND`.

The environment variable `GOMP_CPU_AFFINITY`

Threads are bound to CPUs in order of the specified `cpuid` list.

When the number of specified CPUs is exceeded, it is repeatedly used from the beginning of the list.

The `cpuid` list shall be separated by comma (',' ) or space (' ').

The `cpuid` list can have the next form that has range with increment.

```
cpuid1[-cpuid2[:inc]]
```

*cpuid1*: `cpuid` of the beginning of the range. ( $0 \leq cpuid1 < CPU\_SETSIZE$ )

*cpuid2*: `cpuid` of the end of the range. ( $0 \leq cpuid2 < CPU\_SETSIZE$ )

*inc*: increment ( $1 \leq inc < CPU\_SETSIZE$ )

In addition, it is necessary to be the following.

```
cpuid1 <= cpuid2
```

It becomes equivalent to the case where all CPUs for every increment value *inc* in the range from *cpuid1* to *cpuid2* are specified.

The `cpuid` can be used the above-mentioned value.

However, `cpuid` which can actually be assigned becomes only within the limits of CPU affinity of the process at the start of execution.

See `CPU_SET(3)` about details of `CPU_SETSIZE`.



### Note

If `cpuid` is outside the CPU affinity of the process at the start of execution, an error will be output and the program will terminate. Correct the setting value. The CPU affinity of the process at the start of execution can be checked by specifying `FLIB_HPCFUNC_INFO=TRUE` or system commands such as `taskset` and `numactl`. See "[Appendix H Using High-Speed Facility](#)" about the environment variable `FLIB_HPCFUNC_INFO`. For system commands, refer to each man manual.



### Example

Examples of using environment variable `GOMP_CPU_AFFINITY`

- Example 1:

```
$ export GOMP_CPU_AFFINITY="12,14,13,15"
```

The thread is bound to CPU in order of 12, 14, 13, and 15.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

- Example 2:

```
$ export GOMP_CPU_AFFINITY="12-19"
```

The thread is bound to CPU in order of 12, 13, 14, 15, 16, 17, 18, and 19.

When the number of threads is nine or more, it is repeatedly used from the beginning of the list.

- Example 3:

```
$ export GOMP_CPU_AFFINITY="12-19:2"
```

The thread is bound to CPU in order of 12, 14, 16, and 18.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

- Example 4:

```
$ export GOMP_CPU_AFFINITY="12-16:2,13,19"
```

The thread is bound to CPU in order of 12, 14, 16, 13, and 19.

When the number of threads is six or more, it is repeatedly used f from the beginning of the list.

The environment variable `OMP_PROC_BIND`

The environment variable `OMP_PROC_BIND` can control the thread affinity. Either `true`, `false` or comma separated list of `master`, `close`, or `spread` can be specified to this environment variable. The values of the list sets the thread affinity policy used by a `parallel` region corresponding to the nest level.

This environment variable is set to `close` by default.

If this environment variable is set to `false`, thread affinity is disabled, and the `proc_bind` clause on `parallel` construct is ignored.

Otherwise, thread affinity is enabled and the initial thread is bound to first place in place list.

The `master` thread affinity policy indicates that all threads are bound to same place as master thread.

The `close` thread affinity policy indicates that threads are bound to next to the place where master thread is bound.

The `spread` thread affinity policy indicates that the threads are bound to the one of sub-partition places divided from master thread's place partition.

The effect of `true` is same as `spread`.

See environment variable `OMP_PLACES` for place.

The environment variable `OMP_PLACES`

The environment variable `OMP_PLACES` defines place list.

The explicit place list is defined by ordered set of comma separated non-negative numbers enclosed by braces. The numbers represents the smallest unit of CPU resource in the system.

The value of `OMP_PLACES` is either abstract name or an explicit place list. The value can be specified using following format.

```
{ abstract-name[(num-places)] | place-list }
```

**abstract-name**

The following abstract names can be set to this environment variable. The default value is `cores`.

Abstract Name	Meaning
<code>threads</code>	Each place corresponds to a single hardware thread, which represents smallest unit of CPU resource in the system. The effect of <code>threads</code> is equivalent to <code>cores</code> .
<code>cores</code>	Each place corresponds to a single hardware core, which represents smallest unit of CPU resource in the system. The effect of <code>cores</code> is equivalent to <code>threads</code> .
<code>sockets</code>	Each place corresponds to a single socket, which represents NUMA node.

*num-places*

The length of the place list. Positive integer.

*place-list*

The explicit place list can be specified using following format.

```
{ place:length[:stride] | [!]place[,place-list] } (*1)
```

*place*

```
"{"place1"}" (*2)
```

*place1*

```
{ cpuid:length[:stride] | [!]cpuid[,place1] }
```

*length*

The length of elements. Positive integer.

*stride*

The value of increment or decrement. Positive integer. Default value is 1.

*cpuid*

Smallest unit of CPU resource in the system.

\*1) An exclusion operator "!" exclude the number or place immediately following the operator.

\*2) "{" and "}" are braces.



## Example

Example of environmental variable OMP\_PLACES

- Example 1:

```
$ export OMP_PLACES="cores"
```

This example defines place list of cores abstract name.

- Example 2:

```
$ export OMP_PLACES="cores(4)"
```

This example defines place list of cores abstract name of 4 places.

- Example 3:

```
$ export OMP_PLACES="{12,13,14},{15,16,17},{18,19,20},{21,22,23}"
$ export OMP_PLACES="{12:3},{15:3},{18:3},{21:3}"
$ export OMP_PLACES="{12:3}:4:3"
```

All above examples define same place of 12 to 14, 15 to 17, 18 to 20, and 21 to 23.

## 4.2.3 Example of Compilation and Execution

Example 1:

```
$ FCCpx -Kparallel,reduction,ocl -Nlibomp test1.cc
$ ./a.out
```

In this example, reduction parallelization optimization and the optimization control lines (OCL) are in effect during compilation.

The number of threads using at the runtime is defined by the execution environment. For details about the number of threads, see Section "4.2.2.1 Number of Threads".

Example 2:

```
$ FCCpx -Kparallel -Nlibomp test2.cc
$ OMP_NUM_THREADS=2
$ export OMP_NUM_THREADS
$ ./a.out
$ OMP_NUM_THREADS=4
```



```
$ export OMP_NUM_THREADS
$ ./a.out
```

The environment variable `OMP_NUM_THREADS` is set to 2 and the program executes with two threads.

Next, the environment variable `OMP_NUM_THREADS` is set to 4 and the program executes with four threads.

## 4.2.4 Performance Tuning

---

Tuning information for application programs is gathered by the execution time sampling function provided by the Profiler. A programmer can use this information to identify high cost statements in the program. To execute a program at high speed, the programmer needs to tune the program so that the system is able to parallelize the high cost statements. Refer to the "Profiler User's Guide" for further details about the Profiler.

## 4.2.5 Feature Details on Automatic Parallelization

---

This section describes automatic parallelization.


### 4.2.5.1 Targets for Automatic Parallelization

Target statements of automatic parallelization are loops (including nested loops) and pointer operations.

### 4.2.5.2 Loop Slicing

Automatic parallelization may slice a loop into several pieces. The elapsed execution time is reduced by executing the loop slices in parallel. "Figure 4.3 Loop slicing" illustrates loop slicing.

Figure 4.3 Loop slicing




### 4.2.5.3 Automatic Loop Slicing by Compiler

The system parallelizes a loop if the order of data reference will be the same as with serial execution, and if the system is certain that the result of the parallel loop will be the same as if the loop were processed serially.

"Figure 4.4 A loop that is not a candidate for loop slicing" is an example of a loop that is not amenable to loop slicing. In this `for` loop, when the loop variable `i` is 5000, the value of array element `a[4999]` should be available on the corresponding thread. However, this value is not available and the loop will not be sliced.

Figure 4.4 A loop that is not a candidate for loop slicing



```
for(i = 1; i < 5000; i++){
    a[i] = a[i-1] + b[i];
}
```

```
for(i = 5000; i < 10000; i++){
    a[i] = a[i-1] + b[i];
}
```

#### 4.2.5.4 Loop Interchange and Automatic Loop Slicing

When nested loops are sliced, the system attempts to parallelize the outermost loop if it can. Therefore, the system selects the loop that can be sliced most cheaply and interchanges it with the outermost loop. The purpose of this is to reduce the overhead of multiprocessing and improve the execution performance.

Loop interchange optimization is performed if both `-Kparallel` and `-Keval` are set.

"Figure 4.5 Loop interchange on a nested "for" loop" illustrates the loop slicing after performing loop interchange on a nested "for" loop.

Before performing parallelization for the "for" loop variable `j`, loop interchange is performed so that the parallelization overhead is reduced.

Figure 4.5 Loop interchange on a nested "for" loop

```
for(i = 0; i < 10000; i++){
    for(j = 1; j < 10; j++){
        a[j][i] = a[j][i] + b[j][i];
    }
}
```

↓

```
for(j = 1; j < 10; j++){          <- Interchange
    for(i = 0; i < 10000; i++){ <-
        a[j][i] = a[j][i] + b[j][i];
    }
}
```

↓

↓

```
for(j = 1; j < 5; j++){
    for(i = 0; i < 10000; i++){
        a[j][i] = a[j][i] + b[j][i];
    }
}
```

```
for(j = 5; j < 10; j++){
    for(i = 0; i < 10000; i++){
        a[j][i] = a[j][i] + b[j][i];
    }
}
```


#### 4.2.5.5 Loop Distribution and Automatic Loop Slicing

In "Figure 4.6 Distribution of loop and loop slicing", the references to array `a` cannot be sliced, because the order of data references in parallel execution would be different from the order of data references in serial execution. The references to array `b` can be sliced, because the order of parallel data references is the same as the order of serial data references. In this example, the statement where array `a` is defined and the statement where array `b` is defined are separated into two loops, and the loop where array `b` is defined is parallelized.

The partial automatic parallelization by the distribution of loop is performed when the `-Kloop_part_parallel` option is effective and the compiler assumes that the distribution of loop is effective.

"Figure 4.6 Distribution of loop and loop slicing" shows the example of loop separation and automatic loop slicing.

Figure 4.6 Distribution of loop and loop slicing



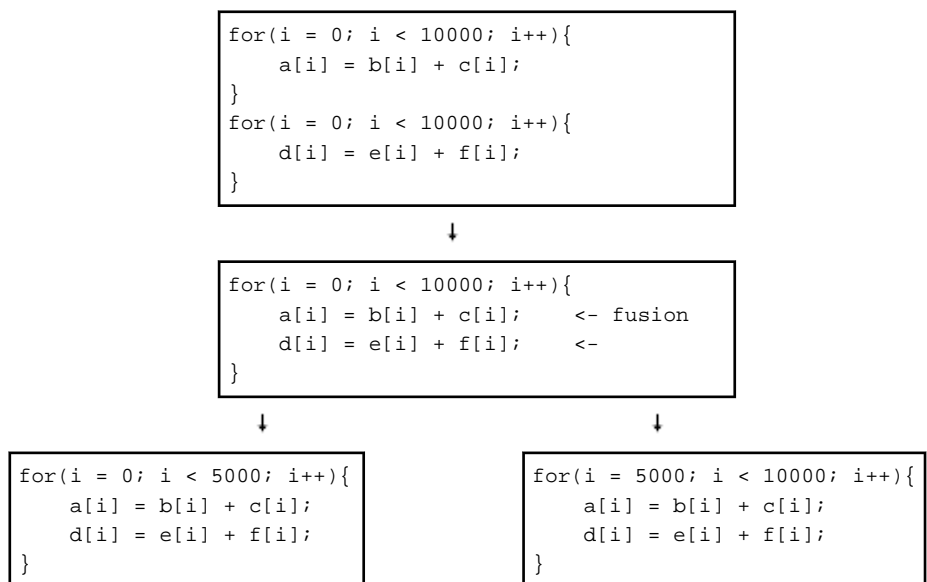
### 4.2.5.6 Loop Fusion and Automatic Loop Slicing

In "Figure 4.7 Fusion of loops and loop slicing", there are loops in sequence having the same loop control.

In this case, the overhead of the loop control and the overhead of multiprocessing can be reduced by merging those two loops into a single loop.

"Figure 4.7 Fusion of loops and loop slicing" shows the example of loop fusion and automatic loop slicing.

Figure 4.7 Fusion of loops and loop slicing



### 4.2.5.7 Loop Reduction

If both the `-Kparallel` and `-Kreduction` options are specified, loop reduction takes place. This optimization slices the loop, and changes the order of the reduction operation (addition, multiplication). Loop reduction may cause calculation errors in the execution results.

The loop reduction optimization is applied if there is one of the following operations in the loop:

- Sum

Example: `s = s + a[i];`

- Product

Example: `p = p * a[i];`

- Dot product

Example: `p = p + a[i] * b[i];`

- Min

Example: `x = min(x, a[i]);`

- Max

Example: `y = max(y, a[i]);`

- Bit or

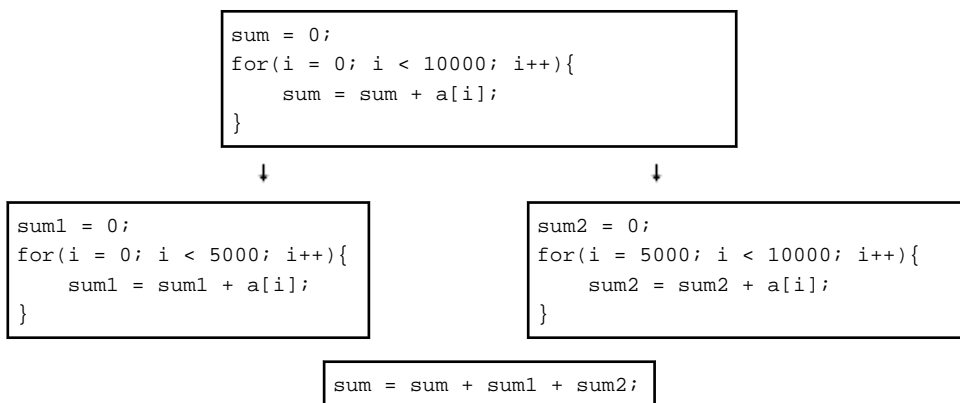
Example: `n = n | a[i];`

- Bit and

Example: `m = m & a[i];`

"Figure 4.8 Loop reduction" shows an example of loop reduction and automatic loop slicing.

Figure 4.8 Loop reduction



#### 4.2.5.8 Restrictions on Loop Slicing

Loops cannot be sliced under the following circumstances.

- Loop slicing would not reduce the elapsed time.
- The loop contains a function reference.
- The loop is complicated.
- The loop contains a standard library function.
- The order of data references in parallel would be different from that of the data references in serial execution.

Loop slicing would not reduce the elapsed time.

If the loop iteration count is small or the number of operations in the loop is small, execution performance of the multiprocessing program may decrease due to the overhead of parallel execution. Therefore, the system does not slice loops when it is clear that the performance will not improve.

"[Figure 4.9 Small loop iteration count and small number of operations](#)" shows an example of a loop for which the loop iteration count and the number of operations is small.

Figure 4.9 Small loop iteration count and small number of operations

---

```
for(i = 0; i < 10; i++){      <- Not parallelized because of small loop iteration count and
    a[i] = a[i] + b[i];      small number of operations
}
```

---

The loop contains a function reference.

A loop containing a function reference cannot be sliced. However, parallelization can be invoked by an optimization control line (OCL). See Section "[4.2.6 Optimization Control Line](#)" for details.

"[Figure 4.10 Loop contains a function call](#)" shows an example of a loop that contains a function call.

Figure 4.10 Loop contains a function call

---

```
for(j = 0; j < 10; j++){
    for(i = 0; i < 10000; i++){      <- Not parallelized because of the function call
        a[j][i] = a[j][i] + b[j][i];
        func(a);
    }
}
```

---

The loop is complicated

The following loops cannot be sliced because they are too complicated:

- There is a branch from inside the loop to outside the loop.

"[Figure 4.11 Jumping out from inside the loop](#)" shows an example of jumping from the inside to the outside of a loop.

Figure 4.11 Jumping out from inside the loop

---

```
for(j = 0; j < 10; j++){
    for(i = 0; i < 10000; i++){      <- Not parallelized because of jumping of the loop
        a[j][i] = a[j][i] + b[j][i];
        if(a[j][i]<0) goto out;
    }
}
out:;
```

---

The loop contains a standard library function.

Some standard library functions do not prevent loop slicing. The compiler produces a diagnostic message during compilation if a loop is not parallelized because of a standard library function reference.

The order of data references in parallel would be different from that of the data references in serial execution.

If the order of data references would be different from that of serial execution, the loop is not suitable for loop slicing as shown in the example of "[Figure 4.4 A loop that is not a candidate for loop slicing](#)".

### 4.2.5.9 Displaying the State of Automatic Parallelization

Specify both the `-Kparallel` and the `-Koptmsg=2` options to check whether automatic parallelization has been performed, and also to confirm what has been parallelized or learn what has prevented parallelization if it has not been performed.

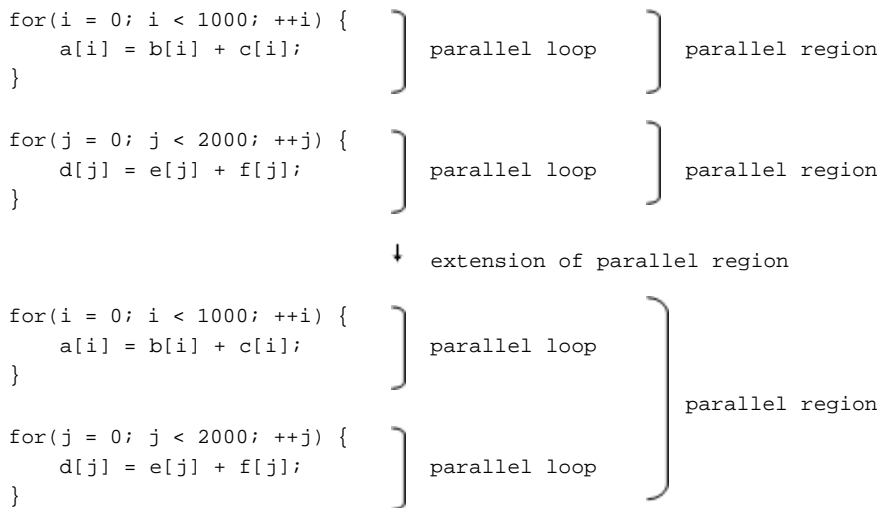
### 4.2.5.10 Parallel Region Extension

When the `-Kregion_extension` and `-Kparallel` options are effective, parallelization overhead may be reduced by extending the parallel region.

The parallel region is one executed by a team consisting of more than one thread. Usually, one parallelization loop is generated for one parallel region as in "[Figure 4.12 Example of parallel region Extension](#)". The cost of generating the parallel region is the overhead of parallel execution.

By applying a parallel region extension, these parallelization loops are generated in one parallel region. As a result, the parallel region is generated only once, and the cost of the overhead is reduced.


Figure 4.12 Example of parallel region Extension



The parallel region extension is also effective in reducing the overhead of the parallel region at the nested loops as in "[Figure 4.13 Example of extension of parallel region \(nested loop\)](#)".

In this case, the number of times the parallel region is generated is reduced from 1000 to 1 because the generating point is moved from the top of the inner loop to the top of the outer loop.

Figure 4.13 Example of extension of parallel region (nested loop)



```
}  
}
```

---

### 4.2.5.11 Block Distribution and Cyclic Distribution

The methods of loop slicing in the automatic parallelization are block distribution and cyclic distribution.

The block distribution is to allocate the block that the loop iteration count is distributed by the number of threads to each thread.

The cyclic distribution is to sequentially allocate the block that the loop iteration count is distributed by the arbitrary size to each thread.

The default method of loop slicing in the automatic parallelization is block distribution.

When `parallel_cyclic` optimization control specifier is specified, the loop is cyclically distributed.

Figure 4.14 Block distribution and cyclic distribution

---

```
for(i = 0; i < 20; i++) {  
    for(j = i; j < 20; j++){  
        a[i][j] = b[i][j];  
    }  
}
```

} parallel loop

---

When block distribution or cyclic distribution is applied to the program of "[Figure 4.14 Block distribution and cyclic distribution](#)", the loop iteration count is allocated to the each thread as follows:

#### - Block distribution

The blocks that the loop iteration count is divided by the number of threads (2) are allocated to each thread as follows:

Thread 0:	{0,1,2,3,4,5,6,7,8,9}
Thread 1:	{10,11,12,13,14,15,16,17,18,19}

#### - Cyclic distribution (block size: 2)

The blocks that the loop iteration count is divided by two iteration (block size: 2) are sequentially allocated to each thread as follows:

Thread 0:	{0,1}, {4,5}, {8,9}, {12,13}, {16,17}
Thread 1:	{2,3}, {6,7}, {10,11}, {14,15}, {18,19}

---

## 4.2.6 Optimization Control Line

This system provides the optimization control line (OCL) to guide automatic parallelization. The optimization control line (OCL) takes effect when both the `-Kparallel` and `-Kocl` options are set.

### 4.2.6.1 Optimization Control Specifier

The optimization control line (OCL) can have several functions depending on the optimization control specifier. The following table shows a list of the optimization control specifiers for automatic parallelization.

The specifier gives the compiler information on parallelization to the compiler.

Table 4.1 Optimization Identifiers that can be specified for parallelization

Optimization Control Specifiers	Explanation	Optimization control line that can be Specified[a]		
		global line	procedure line	loop line
<code>array_private</code>	Designates privatized arrays.	*	*	*
<code>noarray_private</code>	Designates no privatized arrays.	*	*	*
<code>independent [ext[,ext]...]<sup>[b]</sup></code>	Performs the loop slice that calls function ( <i>ext</i> ).	-	*	*
<code>loop_part_parallel</code>	Designates automatic parallelization that requires dividing loops.	*	*	*
<code>loop_nopart_parallel</code>	Suppresses automatic parallelization that requires dividing loops.	*	*	*
<code>serial</code>	Suppresses automatic parallelization.	*	*	*
<code>parallel</code>	Performs automatic parallelization.	*	*	*
<code>parallel_cyclic [n]</code>	Designates cyclic distribution in block size ( <i>n</i> )	-	-	*
<code>parallel_strong</code>	Designates to parallelize a loop with a small number of loop iterations or a loop that has a small number of operations.	*	*	*
<code>reduction</code>	Designates to apply reduction operation.	*	*	*
<code>noreduction</code>	Designates not to apply reduction operation.	*	*	*
<code>temp [var[,var]...]<sup>[d]</sup></code>	Specifies variable <i>var</i> that is used temporarily in loop.	*	*	*
<code>temp_private var [,var]...<sup>[d]</sup></code>	Designates to assign variable ( <i>var</i> ) to local area in thread. As a result, the data dependency between threads is solved, and the automatic parallelization is promoted.	-	-	*
<code>first_private var [,var]...<sup>[d]</sup></code>	Designates to assign variable ( <i>var</i> ) to local area in thread. And, designates to have a possibility of reference of the initial value of the variable. As a result, the data dependency between threads is solved, and the automatic parallelization is promoted.	-	-	*
<code>last_private var [,var]...<sup>[d]</sup></code>	Designates to assign variable ( <i>var</i> ) to local area in thread. And, designates to have a possibility of reference of the variable after executing the loop.  As a result, the data dependency between threads is solved, and the automatic parallelization is promoted.	-	-	*
<code>var1 op var2</code> or <code>var1 op const</code>	Designates a relationship between variable <i>var1</i> and variable <i>var2</i> or between variable <i>var1</i> and constant <i>const</i> .	*	*	*

[a]

\* : Specifies an optimization control specifier to an optimization control line.

- : Cannot specify an optimization control specifier to an optimization control line.

[b] *ext* should be declared before use.



[c] `ary` should be declared before use.

[d] `var` should be declared before use.

## 4.2.6.2 Automatic Parallelization and Optimization Control Specifiers

An optimization control specifier becomes ineffective for a loop that cannot be sliced, even if the optimization control specifier for automatic parallelization is specified. See Section "4.2.5.8 Restrictions on Loop Slicing" for the loops that cannot be sliced.

## 4.2.6.3 Optimization Control Specifiers for Automatic Parallelization

Optimization control specifiers for automatic parallelization are described as follows.

### `array_private` specifier

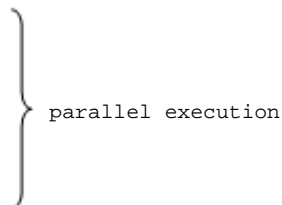
The `array_private` specifier instructs that arrays in a loop are privatized.

The `array_private` specifier is in effect for the code shown in "Figure 4.15 Usage of the `array_private` specifier", and array `a` in the loop is privatized and the loop is parallelized.

Figure 4.15 Usage of the `array_private` specifier

---

```
#pragma loop array_private
for(i = 0; i < 100; i++) {
    for(j = 0; j < 1000; j++) {
        a[j] = i + d[j];
        b[i][j] = a[1] + c[j];
    }
}
```



parallel execution

---

### `noarray_private` specifier

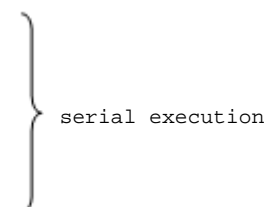
The `noarray_private` specifier instructs that arrays in a loop are not privatized.

The `noarray_private` specifier is in effect for the code shown in "Figure 4.16 Usage of the `noarray_private` specifier", and array `a` in the loop is not privatized and the loop is not parallelized.

Figure 4.16 Usage of the `noarray_private` specifier

---

```
#pragma loop noarray_private
for(i = 0; i < 100; i++) {
    for(j = 0; j < 1000; j++) {
        a[j] = i + d[j];
        b[i][j] = a[1] + c[j];
    }
}
```



serial execution

---

### `independent` specifier

The `independent` specifier instructs that parallel execution will give the same results as serial execution even if a function is called in the loop. As a result, the loop that contains the function is suitable for loop slicing.

A function name which does not inhibit loop slicing can be specified to the `independent` specifier. If a function name is omitted, the `independent` specifier is set for all functions within the range of the loop. A function that is a target of the `independent` specifier should be compiled with the `-Kparallel` option.

In the program of "[Figure 4.17 Example program without independent specifier](#)", because the `func` called in the loop may have side effects, the system cannot determine whether the loop can be parallelized.

Figure 4.17 Example program without independent specifier

---

```
#include <math.h>
main()
{
    int i;
    double a[10000],j;
    double func(double);
    for(i = 0; i < 10000; i++){
        j = 1;
        a[i] = func(j);
    }
    ...
}
double func(double j)
{
    return sqrt(pow(j, 2) + 3 * j + 6);
}
```

---

If it is known that the function `func` has no side effects, slicing of the loop can be promoted using the `independent` specifier, as shown in "[Figure 4.18 Usage of the independent specifier](#)".

Figure 4.18 Usage of the independent specifier

---

```
#include <math.h>
main()
{
    int i;
    double a[10000],j;
    double func(double);

    #pragma loop independent func
    for(i = 0; i < 10000; i++){
        j = 1;
        a[i] = func(j);
    }
    ...
}
double func(double j)
{
    return sqrt(pow(j, 2) + 3 * j + 6);
}
```

} parallel execution

---

## Note

This system may produce incorrect results if a procedure that cannot be a target of loop-slicing is specified for the `independent` specifier.

An example of such a procedure is:

- a procedure that has data dependency with other procedures

### `loop_part_parallel` specifier

The `loop_part_parallel` specifier instructs to perform automatic parallelization that requires dividing loops.

The `loop_part_parallel` specifier is in effect for the code shown in "[Figure 4.19 Usage of the `loop\_part\_parallel` specifier](#)", then the loop is divided and automatic parallelization is applied to the part of the divided loop without the `-Kloop_part_parallel` option.

Figure 4.19 Usage of the `loop_part_parallel` specifier

---

```
#pragma loop loop_part_parallel
for(i = 1; i < n; i++){
    a[i] = a[i] - b[i] + log(c[i]);      ] parallel execution
    d[i] = d[i-1] + a[i];              ] serial execution
}
```

---

### `loop_nopart_parallel` specifier

The `loop_nopart_parallel` specifier instructs to suppress automatic parallelization that requires dividing loops.

The `loop_nopart_parallel` specifier is in effect for the code shown in "[Figure 4.20 Usage of the `loop\_nopart\_parallel` specifier](#)", then the loop is not divided and automatic parallelization is not performed to the loop.

Figure 4.20 Usage of the `loop_nopart_parallel` specifier

---

```
#pragma loop loop_nopart_parallel
for(i = 1; i < n; i++){
    a[i] = a[i] - b[i] + log(c[i]);
    d[i] = d[i-1] + a[i];
}
```

} serial execution

---

### `serial` specifier

The `serial` specifier instructs to inhibit loop slicing.

For instance, if the programmer knows that serial execution of a loop is faster than parallel execution, perhaps because the iteration count will always be small, the `serial` specifier may be specified for the loop.

In the program of "[Figure 4.21 Example program without the `serial` specifier](#)", if `loop2` should not be sliced, the slicing can be inhibited by specifying `serial` specifier as shown in "[Figure 4.22 Usage of the `serial` specifier](#)".

Figure 4.21 Example program without the `serial` specifier

---

```

for(j = 0; j < 10; j++){
    for(i = 0; i < l; i++){      <- loop1
        a1[j][i] = a1[j][i] + b1[j][i];
    }
}
...
for(j = 0; j < 10; j++){
    for(i = 0; i < m; i++){      <- loop2
        a2[j][i] = a2[j][i] + b2[j][i];
    }
}
...
for(j = 0; j < 10; j++){
    for(i = 0; i < n; i++){      <- loop3
        a3[j][i] = a3[j][i] + b3[j][i];
    }
}

```

} parallel execution  
} parallel execution  
} parallel execution

---

Figure 4.22 Usage of the serial specifier

---

```

for(j = 0; j < 10; j++){
    for(i = 0; i < l; i++){      <- loop1
        a1[j][i] = a1[j][i] + b1[j][i];
    }
}
...
#pragma loop serial
for(j = 0; j < 10; j++){
    for(i = 0; i < m; i++){      <- loop2
        a2[j][i] = a2[j][i] + b2[j][i];
    }
}
...
for(j = 0; j < 10; j++){
    for(i = 0; i < n; i++){      <- loop3
        a3[j][i] = a3[j][i] + b3[j][i];
    }
}

```

} parallel execution  
} serial execution  
} parallel execution

---

#### parallel specifier

The `parallel` specifier instructs to reverse the effect of the `serial` specifier and to allow loop slicing for specific loops.

For example, it may be necessary to slice only `loop2` in the program of "[Figure 4.23 Example program without the parallel specifier](#)". This can be achieved by specifying a combination of the `parallel` and `serial` specifiers as shown in "[Figure 4.24 Example program with combination of the parallel and serial specifiers](#)".

Figure 4.23 Example program without the parallel specifier

---

```

void foo(){
    ...
    for(j = 0; j < m1; j++){          <- loop1
        for(i = 0; i < 100; i++){
            a1[j][i] = a1[j][i] + b1[j][i];
        }
    }
    ...

    for(j = 0; j < m2; j++){          <- loop2
        for(i = 0; i < 100; i++){
            a2[j][i] = a2[j][i] + b2[j][i];
        }
    }
    ...

    for(j = 0; j < m3; j++){          <- loop3
        for(i = 0; i < 100; i++){
            a3[j][i] = a3[j][i] + b3[j][i];
        }
    }
    ...
}

```

} parallel execution

} parallel execution

} parallel execution

---

Figure 4.24 Example program with combination of the parallel and serial specifiers

---

```

#pragma global serial
void foo(){
    ...

    for(j = 0; j < m1; j++){          <- loop1
        for(i = 0; i < 100; i++){
            a1[j][i] = a1[j][i] + b1[j][i];
        }
    }
    ...

    #pragma loop parallel
    for(j = 0; j < m2; j++){          <- loop2
        for(i = 0; i < 100; i++){
            a2[j][i] = a2[j][i] + b2[j][i];
        }
    }
    ...

    for(j = 0; j < m3; j++){          <- loop3
        for(i = 0; i < 100; i++){
            a3[j][i] = a3[j][i] + b3[j][i];
        }
    }
    ...
}

```

} serial execution

} parallel execution

} serial execution

### parallel\_cyclic specifier

The `parallel_cyclic` specifier instructs to perform cyclic distribution for a loop that automatic parallelization can be analyzed by compiler.  $n$  following the `parallel_cyclic` specifier is a decimal value ranging from 1 to 10000, and specifies the block size. If  $n$  is omitted, the default block size is 1.

When this specifier is applied, the loop is parallelized without estimating the effects of the parallelization.

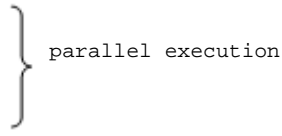
For details about cyclic distribution, see Section "4.2.5.11 Block Distribution and Cyclic Distribution".

In "Figure 4.25 Usage of the `parallel_cyclic` specifier", cyclic distribution in block size 2 is applied to the loop.

Figure 4.25 Usage of the `parallel_cyclic` specifier

---

```
#pragma loop parallel_cyclic 2
for (i = 0; i < 1000; i++) {
    for (j = i; j < 1000; j++) {
        a[i][j] = b[i][j];
    }
}
```



---

### parallel\_strong specifier

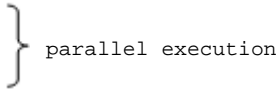
The `parallel_strong` specifier instructs to indicate compelling parallelization for loops that might not usually be parallelized, for example those with minimal iteration or with few operations.

Usually, for the code shown in "Figure 4.26 Usage of the `parallel_strong` specifier", parallelization is suppressed to the loop because of the minimal iteration, but the `parallel_strong` specifier forces parallelization to the loop.

Figure 4.26 Usage of the `parallel_strong` specifier

---

```
#pragma loop parallel_strong
for (i = 0; i < 10; i++) {
    a[i] = a[i] + b[i] + c[i];
}
```



---

### reduction specifier

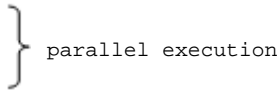
The `reduction` specifier instructs that the reduction operation in a loop is parallelized.

The `reduction` specifier is in effect for the code shown in "Figure 4.27 Usage of the `reduction` specifier", and the reduction operation in the loop is parallelized without the `-Kreduction` option.

Figure 4.27 Usage of the `reduction` specifier

---

```
#pragma loop reduction
for(i = 0; i < 5000; i++){
    s = s + a[i];
}
```



---

### noreduction specifier

The `noreduction` specifier instructs that the reduction operation in a loop is not parallelized.

The `noreduction` specifier is in effect in "Figure 4.28 Usage of the `noreduction` specifier", and the reduction operation in the loop is not parallelized.

Figure 4.28 Usage of the `noreduction` specifier

---

```
#pragma loop noreduction
for(i = 0; i < 5000; i++){
    s = s + a[i];
}
```

} serial execution

---

#### `temp` specifier

The `temp` specifier instructs that the variables listed are used only within the loop. As a result, the execution performance of the parallelized loop can be improved.

Figure 4.29 Example program without a `temp` specifier

---

```
int t;
main()
{
    ...

    for(j = 0; j < 50; j++){
        for(i = 0; i < 1000; i++){
            t = a[j][i] + b[j][i];
            c[j][i] = t + d[j][i];
        }
    }
    ...

    func();
}
```

} parallel execution

---

In this case, if it is known that the value of `t` at the end of the loop is not needed in the function `func`, the programmer may assign the `temp` specifier with `t` as shown in "Figure 4.30 Usage of the `temp` specifier". As a result, the execution performance improves, because the instruction which corrects the value of `t` becomes unnecessary at the end of the loop.

Figure 4.30 Usage of the `temp` specifier

---

```
int t;
main()
{
    ...

    #pragma loop temp t
    for(j = 0; j < 50; j++){
        for(i = 0; i < 1000; i++){
            t = a[j][i] + b[j][i];
            c[j][i] = t + d[j][i];
        }
    }
    ...
}
```

} parallel execution

---

```
func();  
}
```

---

## Note

Usage of this specifier is the programmer's responsibility. If a variable is specified in a `temp` specifier, and is referred to later in the code, the results will be unpredictable.

### `temp_private` specifier

In the target loop of automatic parallelization, the `temp_private` specifier instructs to treat a specified variable `var` as a local variable in each thread.

When the optimization message like the following is output, specify this specifier to assign the target data to local area of each thread. As a result, the data dependency is solved, and the automatic parallelization is promoted.

```
jwd5208p-i "file", line n: This loop is not parallelized because the uncertain order of the  
definition and reference to variable 'var' may cause different results from serial execution.
```

This specifier is only effective for the subsequent loop of this specifier.

The variable that can be specified for this specifier is as follows.

- The variable of integer type or float type.
- The variable of array type with integer type or float type in element, and array size has been fixed at compilation.

When this specifier is applied, the following message is output.

```
jwd5013p-i "file", line n: The temp_private specifier in Optimization Control Line (OCL) is applied  
to variable 'var'.
```

In "[Figure 4.31 Example program without `temp\_private` specifier](#)", this loop is not parallelized, because the order of definition and the reference of array 'a' is uncertain in an outside loop.

### Figure 4.31 Example program without `temp_private` specifier

```
for (i = 0; i < 100; i++) {           <- Target loop for automatic parallelization.  
  for (j = 0; j < m; j++) {         <- Because the value of variable 'm' is uncertain, the range  
    a[j] = b[j];                   where array 'a' is defined is uncertain.  
  }  
  for (j = 0; j < n; j++) {         <- Because the value of variable 'n' is uncertain, the range  
    c[i][j] = a[j];                 where array 'a' is referred is uncertain.  
  }  
}
```

In "[Figure 4.32 Usage of `temp\_private` specifier](#)", data dependency of array 'a' is solved by the `temp_private` specifier. As a result, the outside loop is parallelized.

### Figure 4.32 Usage of `temp_private` specifier



```

#pragma loop temp_private a
for(i = 0; i < 100; i++){
    for(j = 0; j < m; j++){
        a[j] = b[j];
    }
    for(j = 0; j < n; j++){
        c[i][j] = a[j];
    }
}

```

} parallel execution

## Note

When `temp_private` specifier is specified incorrectly, errors in execution results may occur.

### `first_private` specifier

In the target loop of automatic parallelization, the `first_private` specifier instructs to treat a specified variable `var` as a local variable in each thread, and instructs that the initial value of the variable is referred at the loop entry of each thread.

When the optimization message like the following is output, specify this specifier to assign the target data to local area of each thread. As a result, the data dependency is solved, and the automatic parallelization is promoted.

```

jwd5208p-i "file", line n: This loop is not parallelized because the uncertain order of the
definition and reference to variable 'var' may cause different results from serial execution.

```

This specifier is only effective for the subsequent loop of this specifier.

The variable that can be specified for this specifier is as follows.

- The variable of integer type or float type.
- The variable of array type with integer type or float type in element, and array size has been fixed at compilation.

When this specifier is applied, the following message is output

```

jwd5014p-i "file", line n: The first_private specifier in Optimization Control Line (OCL) is
applied to variable 'var'.

```

In "[Figure 4.33 Example program without first\\_private specifier](#)", this loop is not parallelized, because the order of definition and the reference of array 'a' is uncertain in an outside loop.

### Figure 4.33 Example program without `first_private` specifier

```

for (i = 0; i < 100; i++) {          <- Target loop for automatic parallelization.
    for (j = 0; j < m; j++) {      <- Because the value of variable 'm' is uncertain, the range
        b[i][j] = a[j];           where array 'a' is defined is uncertain.
    }
    for (j = n; j < 100; j++) {    <- Because the value of variable 'n' is uncertain, the range
        a[j] = b[i][j] + d[j];     where array 'a' is referred is uncertain.
        c[i][j] = a[j];
    }
}

```

In "[Figure 4.34 Usage of first\\_private specifier](#)", data dependency of array 'a' is solved by the `first_private` specifier. As a result, the outside loop is parallelized.

Figure 4.34 Usage of `first_private` specifier

```
#pragma loop first_private a
for(i = 0; i < 100; i++){
    for(j = 0; j < m; j++){
        b[i][j] = a[j];
    }
    for(j = n; j < 100; j++){
        a[j] = b[i][j] + d[j];
        c[i][j] = a[j];
    }
}
```

} parallel execution

### Note

When `first_private` specifier is specified incorrectly, errors in execution results may occur.

### `last_private` specifier

In the target loop of automatic parallelization, the `last_private` specifier instructs to treat a specified variable `var` as a local variable in each thread, and instructs that the value of the variable is referred after executing the loop.

When the optimization message like the following is output, specify this specifier to assign the target data to local area of each thread. As a result, the data dependency is solved, and the automatic parallelization is promoted.

```
jwd5208p-i "file", line n: This loop is not parallelized because the uncertain order of the
definition and reference to variable 'var' may cause different results from serial execution.
```

This specifier is only effective for the subsequent loop of this specifier.

The variable that can be specified for this specifier is as follows.

- The variable of integer type or float type.
- The variable of array type with integer type or float type in element, and array size has been fixed at compilation.

When this specifier is applied, the following message is output.

```
jwd5015p-i "file", line n: The last_private specifier in Optimization Control Line (OCL) is applied
to variable 'var'.
```

In "Figure 4.35 Example program without `last_private` specifier", this loop is not parallelized, because the order of definition and the reference of array 'a' is uncertain in an outside loop.

Figure 4.35 Example program without `last_private` specifier

```
for (i = 0; i < 100; i++) {      <- Target loop for automatic parallelization.
    for (j = 0; j < m; j++) {
        b[i][j] = i;
    }
    for (j = n; j < 100; j++) { <- Because the value of variable 'n' is uncertain, the range
        a[j] = b[i][j] + d[j];   where array 'a' is referred is uncertain.
        c[i][j] = a[j];
    }
}
foo(a, c);
```

In "Figure 4.36 Usage of `last_private` specifier", data dependency of array 'a' is solved by the `last_private` specifier. As a result, the outside loop is parallelized.

Figure 4.36 Usage of `last_private` specifier

---

```
#pragma loop last_private a
for(i = 0; i < 100; i++){
    for(j = 0; j < m; j++){
        b[i][j] = i;
    }
    for(j = n; j < 100; j++){
        a[j] = b[i][j] + d[j];
        c[i][j] = a[j];
    }
}
foo(a, c);
```

} parallel execution

---

### Note

When `last_private` specifier is specified incorrectly, errors in execution results may occur.

*var1 op var2* specifier

*var1 op const* specifier

The "*var1 op var2*" and "*var1 op const*" specifiers indicate the relationship between variables or between a variable and a constant.

*var1* and *var2* should be variables. *const* should be an integer constant. *op* should be the relational operators (<, <=, >, >=) or the equality operators (==, !=).

In "Figure 4.37 Usage of the *var1 op const* specifier", since the relationship between values of the subscripts of array elements is not known unless the optimization control line is described, data dependency of array a in parallel execution may vary with that in serial execution, and the loop is not parallelized. In this case, by specifying the "*var1 op const*" specifier, it turns out that data dependency in parallel execution is equivalent to that in serial execution, and the loop is parallelized.

Figure 4.37 Usage of the *var1 op const* specifier

---

```
int i, m;
float a[10000], b[10000];
...

#pragma loop m > 2000
for (i = 0; i < 2000; i++){
    a[i] = a[i + m] + b[i];
}
```

} parallel execution

---

## 4.2.7 Notes on Automatic Parallelization

This section provides notes on the automatic parallelization feature.

## 4.2.7.1 Multiprocessing of Nested Loops

This system can only parallelize one level of nesting in nested loops. If there is a call in a parallelized loop to a function which contains itself another parallelized loop, a nest of parallelized loops through function boundaries is generated. A program that contains these types of nested loops will give unpredictable results when compiled with the `-Kparallel,instance=N` option. It should be compiled only with the `-Kparallel` compiler option.

"[Figure 4.38 Multiprocessing of nested loops](#)" shows an example in which one of the parallelized loops should be executed in serial. If a source program that contains such a nest of loops is compiled with the `-Kparallel,instance=N` option, the result may be incorrect.

Figure 4.38 Multiprocessing of nested loops

```
#include <stdio.h>
#define M 4000
double a[M][M], b[M][M];
void f1(int);
void f2(int);

main()
{
    f1(M);
    printf("%e %e\n", a[0][0], b[0][0]);
}

void f1(int n)
{
    int i;
    #pragma loop independent
    for(i = 0; i < n; i++){    <- Executed in parallel
        f2(i);
    }
}

void f2(int i)
{
    int j;
    for(j = 0; j < i + 1; j++){    <- Should be executed in serial
        a[i][j] = 3.8;
    }
    for(j = 0; j < i + 1; j++){    <- Should be executed in serial
        b[i][j] = 4.8;
    }
}
```

The result may be incorrect if the source program "a.cc" is compiled as follows.

```
$ FCCpx -Kparallel,instance=4 -Nlibomp a.cc (invalid use)
```

To prevent this type of incorrect execution, serialize the loop in the function by specifying the `serial` specifier as follows.

```
void f2(int i)
{
    int j;
    #pragma procedure serial
    for(j = 0; j < i + 1; j++){    <- Serialization
        a[i][j] = 3.8;
    }
    for(j = 0; j < i + 1; j++){    <- Serialization
        b[i][j] = 4.8;
    }
}
```

```
}  
}
```


### 4.2.7.2 Side Effect of Using -Kparallel,reduction

When `-Kparallel,reduction` is specified as a compiler option, the result of parallel execution may be different from the result of serial execution. The reason for this is that the order of operations in the parallel execution may be different from the order of operations in the serial execution.

"[Figure 4.39 Calculation error with the reduction optimization](#)" is an example illustrating the loop reduction optimization in Section "[4.2.5.7 Loop Reduction](#)". The variable `sum` accumulates the values `a[0]` to `a[9999]` in order with serial execution. In parallel execution, `sum1` accumulates the values `a[0]` to `a[4999]`, and `sum2` accumulates the values `a[5000]` to `a[9999]` simultaneously. After that, the sum of `sum1` and `sum2` is assigned to `sum`.

Loop reduction optimization may cause a calculation error in the execution result, because the order of adding the array elements is different between parallel and serial execution.

Figure 4.39 Calculation error with the reduction optimization



### 4.2.7.3 Examples of Invalid Usage of an Optimization Control Line

This section shows examples of invalid usage of the optimization control lines `norecurrence`, `temp`, `independent`, `temp_private`, and `last_private`. The system will perform incorrect loop slicing because of the incorrect optimization control lines.

In "[Figure 4.40 Invalid usage of the norecurrence specifier](#)", the `norecurrence` specifier is specified incorrectly for array `a`. The execution result will be unpredictable when array `a` is sliced, because the order of data references for array `a` in parallel execution is different from the order of data references in serial execution.

Figure 4.40 Invalid usage of the `norecurrence` specifier

```
#pragma loop norecurrence a  
for(i = 1; i < 10000; i++){  
    a[i] = a[i-1] + b[i];  
}
```

In "[Figure 4.41 Invalid usage of the temp specifier](#)", the `temp` specifier is specified incorrectly for variable `t`. There will not be a correct value assigned to variable `t`, because the system cannot guarantee a correct value of variable `t` at the end of the loop.

Figure 4.41 Invalid usage of the `temp` specifier

```

#pragma loop temp t
for(i = 0; i < 1000; i++){
    t = a[i] + b[i];
    c[i] = t + d[i];
}
last = t;

```

In "Figure 4.42 Invalid usage of the independent specifier", the independent specifier is specified incorrectly for function func. The execution result is unpredictable when array a is sliced, because the order of the data references for array a in parallel execution is different from the order of data references in serial execution.

Figure 4.42 Invalid usage of the independent specifier

```

int a[1000], b[1000];

main()
{
    int i;
    void func(int);
#pragma procedure independent func
    for(i = 1; i < 1000; i++){
        a[i] = b[i] + 1;
        func(i-1);
    }
}

void func(int j)
{
    a[j] = a[j] + 1;
}

```

In "Figure 4.43 Invalid usage of the temp\_private specifier", the temp\_private specifier is specified incorrectly for variable a. Because an uncertain value is referred in each thread, the execution result is incorrect.

Figure 4.43 Invalid usage of the temp\_private specifier

```

#pragma loop temp_private a
for(i = 0; i < 1000; i++){
    b[i] = a;
}

```

In "Figure 4.44 Invalid usage of the last\_private specifier", the last\_private specifier is specified incorrectly for array a. Because the array a may not be defined by the last thread, the execution result may be incorrect.

Figure 4.44 Invalid usage of the last\_private specifier

```

#pragma loop last_private a
for(j = 0; j < 1000; j++){
    for(i = j; i < 500; i++){
        a[i] = b[i];
    }
}

```

```
}  
foo(a);
```

#### 4.2.7.4 Standard Library Function References

If there is a standard library function reference that is not suitable for loop slicing in a function called in a parallelized loop, execution of the program will produce incorrect results. The execution performance of the multiprocessing program may decrease due to the overhead of parallel execution.

#### 4.2.7.5 Note on C++ Programs

Under the following conditions, a loop may not be parallelized.

- A loop in a function which contains an exception
- A loop in a constructor or destructor
- A loop in range of `#pragma {global|procedure|loop} tmp` in a virtual function

When a loop is in a function that is called from a constructor or destructor, or when the loop is parallelized in global operator new or operator delete, the execution result may be incorrect.

### 4.3 Parallelization by OpenMP Specification

This section explains parallelization as described in the OpenMP Specification. See the "OpenMP Architecture Review Board" website for details about OpenMP specification.

This system supports following specifications except the directives, the environmental variables and the runtime routines for the device.

Compiler Mode	Supported Specifications
Trad Mode	OpenMP 3.1
	Part of the OpenMP 4.0 <sup>[a]</sup>
	Part of the OpenMP 4.5 <sup>[a]</sup>
Clang Mode	OpenMP 4.5 <sup>[b]</sup>
	Part of the OpenMP 5.0 <sup>[c]</sup>

[a] The following functions can be used:

- `simd construct`
- `declare simd construct`
- `proc_bind` clause of `parallel construct`
- `depend` clause of `task construct`
- `taskgroup construct`

[b] The following functions cannot be used:

- `declare simd construct`
- `linear` clause of `taskloop simd construct`

[c] The following functions can be used:

- `in_reduction` clause of `task construct`
- `task_reduction` clause of `taskgroup construct`
- `reduction` and `in_reduction` clauses of `taskgroup construct`

## 4.3.1 Compilation

For compiling source programs and linking, specify the following options to the compile command.

The specified option is different by whether to perform parallelization by the OpenMP specification.

Parallelization by the OpenMP specification	Specified option	
	Compile time	Link time
When using the parallelization and the SIMD Extensions	-Kopenmp	-Kopenmp -Nlibomp
When using the SIMD Extensions only, not using the parallelization	-Kopenmp_simd	-

-: There is no specified option.

### 4.3.1.1 Compiler Option for OpenMP C++ Program

This section describes the compilation options used to compile OpenMP programs.

#### -Kopenmp

This option enables OpenMP directives and compiles the source programs.

The -Kopenmp option should also be specified at the linkage phase of an object file, when the object file has previously been compiled with this option.



#### Example

```
$ FCCpx a.cc -Kopenmp -c
$ FCCpx a.o -Kopenmp -Nlibomp
```

#### -Kopenmp\_simd

-Kopenmp\_simd option enables only the `simd` construct and the `declare simd` construct of OpenMP and compiles the source programs.

#### -Nlibomp

This option specifies to use LLVM OpenMP Library for OpenMP.

-Nlibomp option is required at linking.

### 4.3.1.2 Obtaining of Optimization Information

This section describes Trad Mode compiler options obtaining optimization Information.

#### -Nsrc, -Nlst

When the -Kopenmp option is specified with -Nsrc or -Nlst simultaneously for the compilation of an OpenMP C++ program, the following letters are considered optimization information and are noted as such in the statements in the compilation information:

- 'p' is marked on each statement that can be executed in parallel. This is not marked on a statement executed redundantly. Statements directly enclosed by a construct of "for" directive or "sections" directive meet this condition.
- 's' is marked on each statement that will be executed on a single thread simultaneously. Statements directly enclosed by a construct of a "master" directive, a "single" directive, a "critical" directive, or "ordered" directive meet this condition. A statement specified with an "atomic" directive is marked with 's' when the whole statement is executed with a single thread exclusively.
- 'm' is marked on each statement that can be executed partially in parallel and partially in serial.

For nested parallelism, the mark is related to the innermost `parallel` region enclosing the statement. The output is not affected whether or not the `parallel` region being processed is executed in parallel or serially.



-Koptmsg=2

The diagnostic message related to the execution performance such as SIMD Extensions and collapse of the OpenMP specification is output.

### 4.3.1.3 Restriction of OpenMP programs

In Trad Mode, user defined functions that include OpenMP directives are not expanded inline.

## 4.3.2 Execution Process

---

The execution process of an OpenMP C++ program is the same as the procedure for a serial-processing program.

### 4.3.2.1 Environment Variable at Execution

Environment Variable `OMP_WAIT_POLICY`

<code>ACTIVE</code>	Uses spin wait until all threads are synchronized.
<code>PASSIVE</code>	Uses no spin wait but suspending wait.

The default value is "PASSIVE".

Select `ACTIVE` to give priority to the elapsed time. Select `PASSIVE` to give priority to the CPU time.

Environment Variable `OMP_PROC_BIND`

The environment variable `OMP_PROC_BIND` can control the thread affinity. Either `true`, `false` or comma separated list of `master`, `close`, or `spread` can be specified to this environment variable. The values of the list sets the thread affinity policy used by a `parallel` region corresponding to the nest level.

This environment variable is set to `close` by default.

If this environment variable is set to `false`, thread affinity is disabled, and the `proc_bind` clause on `parallel` construct is ignored.

Otherwise, thread affinity is enabled and the initial thread is bound to first place in place list.

The `master` thread affinity policy indicates that all threads are bound to same place as master thread.

The `close` thread affinity policy indicates that threads are bound to next to the place where master thread is bound.

The `spread` thread affinity policy indicates that the threads are bound to the one of sub-partition places divided from master thread's place partition.

The effect of `true` is same as `spread`.

See environment variable `OMP_PLACES` for place.

Environment Variable `OMP_PLACES`

The environment variable `OMP_PLACES` defines place list.

The explicit place list is defined by ordered set of comma separated non-negative numbers enclosed by braces. The numbers represents the smallest unit of CPU resource in the system.

The value of `OMP_PLACES` is either abstract name or an explicit place list. The value can be specified using following format.

{ <i>abstract-name</i> [( <i>num-places</i> )]   <i>place-list</i> }
--

*abstract-name*

The following abstract names can be set to this environment variable. The default value is `cores`.

Abstract Name	Meaning
<code>threads</code>	Each place corresponds to a single hardware thread, which represents smallest unit of CPU resource in the system.  The effect of <code>threads</code> is equivalent to <code>cores</code> .

Abstract Name	Meaning
<code>cores</code>	Each place corresponds to a single hardware core, which represents smallest unit of CPU resource in the system. The effect of <code>cores</code> is equivalent to <code>threads</code> .
<code>sockets</code>	Each place corresponds to a single socket, which represents NUMA node.

### *num-places*

The length of the place list. Positive integer.

### *place-list*

The explicit place list can be specified using following format.

```
{ place:length[:stride] | [!]place[,place-list] } (*1)
```

### *place*

"{"*place*"}" (\*2)

### *place1*

{ *cpuid*:*length*[:*stride*] | [!]*cpuid*[,*place*] }

### *length*

The length of elements. Positive integer.

### *stride*

The value of increment or decrement. Positive integer. Default value is 1.

### *cpuid*

Smallest unit of CPU resource in the system.

\*1) An exclusion operator "!" exclude the number or place immediately following the operator.

\*2) "{" and "}" are braces.



## Example

Example of environmental variable `OMP_PLACES`

- Example 1:

```
$ export OMP_PLACES="cores"
```

This example defines place list of `cores` abstract name.

- Example 2:

```
$ export OMP_PLACES="cores(4)"
```

This example defines place list of `cores` abstract name of 4 places.

- Example 3:

```
$ export OMP_PLACES="{12,13,14},{15,16,17},{18,19,20},{21,22,23}"
$ export OMP_PLACES="{12:3},{15:3},{18:3},{21:3}"
$ export OMP_PLACES="{12:3}:4:3"
```

All above examples define same place of 0 to 12 to 14, 15 to 17, 18 to 20, and 21 to 23.

## Environment Variable OMP\_STACKSIZE

A user can specify the stack area size for each thread using the environment variable OMP\_STACKSIZE in Bytes, Kbytes, Mbytes, Gbytes, or Tbytes. The default size is 8M bytes.

## Environment Variable GOMP\_CPU\_AFFINITY

Threads are bound to CPUs in order of the specified cpuid list.

When the number of specified CPUs is exceeded, it is repeatedly used from the beginning of the list.

The cpuid list shall be separated by comma (',' ) or space (' ').

The cpuid list can have the next form that has range with increment.

```
cpuid1[-cpuid2[:inc]]
```

*cpuid1*: cpuid of the beginning of the range. ( $0 \leq cpuid1 < CPU\_SETSIZE$ )

*cpuid2*: cpuid of the end of the range. ( $0 \leq cpuid2 < CPU\_SETSIZE$ )

*inc*: increment ( $1 \leq inc < CPU\_SETSIZE$ )

In addition, it is necessary to be the following.

```
cpuid1 <= cpuid2
```

It becomes equivalent to the case where all CPUs for every increment value *inc* in the range from *cpuid1* to *cpuid2* are specified.

The cpuid can be used the above-mentioned value. However, cpuid which can actually be assigned becomes only within the limits of CPU affinity of the process at the start of execution.

See CPU\_SET(3) about details of CPU\_SETSIZE.

## Note

If cpuid is outside the CPU affinity of the process at the start of execution, an error will be output and the program will terminate. Correct the setting value. The CPU affinity of the process at the start of execution can be checked by specifying FLIB\_HPCFUNC\_INFO=TRUE or system commands such as taskset and numactl. See "[Appendix H Using High-Speed Facility](#)" about the environment variable FLIB\_HPCFUNC\_INFO. For system commands, refer to each man manual.

## Example

Examples of using environment variable GOMP\_CPU\_AFFINITY

- Example 1:

```
$ export GOMP_CPU_AFFINITY="12,14,13,15"
```

The thread is bound to CPU in order of 12, 14, 13, and 15.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

- Example 2:

```
$ export GOMP_CPU_AFFINITY="12-19"
```

The thread is bound to CPU in order of 12, 13, 14, 15, 16, 17, 18, and 19.

When the number of threads is nine or more, it is repeatedly used from the beginning of the list.

- Example 3:

```
$ export GOMP_CPU_AFFINITY="12-19:2"
```

The thread is bound to CPU in order of 12, 14, 16, and 18.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

- Example 4:

```
$ export GOMP_CPU_AFFINITY="12-16:2,13,19"
```

The thread is bound to CPU in order of 12, 14, 16, 13, and 19.

When the number of threads is six or more, it is repeatedly used from the beginning of the list.

.....

### 4.3.2.2 Environment Variable for OpenMP Specifications

The user can set the following environment variables, which are in accordance with the OpenMP Specifications.

For environment variables available on OpenMP 4.0 and later, the specifications they are supported on are enclosed in a pair of parentheses after each environment variable.

For the details, refer to the OpenMP specifications.

Environment Variable `OMP_SCHEDULE`

Sets the schedule type and chunk size for directives that have the schedule type at runtime.

Environment Variable `OMP_NUM_THREADS`

Sets the number of threads to use during execution.

Environment Variable `OMP_DYNAMIC`

Enables or disables dynamic thread adjustment.

Environment Variable `OMP_PROC_BIND`

Specifies whether threads are bound to CPUs.

Environment Variable `OMP_NESTED`

Enables or disables nested parallelism.

Environment Variable `OMP_STACKSIZE`

Specifies the stack size for each thread.

Environment Variable `OMP_WAIT_POLICY`

Specifies the behavior of the standby thread.

Environment Variable `OMP_MAX_ACTIVE_LEVELS`

Specifies the maximum number of nested active parallel regions.

Environment Variable `OMP_THREAD_LIMIT`

Specifies the maximum number of threads performed on an Open MP program.

Environment Variable `OMP_PLACES` (OpenMP 4.0 and later)

This specifies the list of places for thread affinity.

Environment Variable `OMP_CANCELLATION` (OpenMP 4.0 and later)

This enables or disables the cancellation.

Environment Variable `OMP_MAX_TASK_PRIORITY` (OpenMP 4.5 and later)

This specifies maximum task priority.

Environment Variable `OMP_DISPLAY_ENV` (OpenMP 4.0 and later)

This enables or disables to display the OpenMP version number and the initial value of the internal control variables.

The OpenMP version number is 201611.

### 4.3.2.3 Notes during Execution

When executing programs created on this system, note the following.

### Variable allocation at execution

If the size of the stack area for every thread is defined as a specific size, the environment variable `OMP_STACKSIZE` can be specified.

See Section "4.3.2.1 Environment Variable at Execution" for information about environment variables `OMP_STACKSIZE`.

### Limits on the Number of CPUs

The number of CPUs that can be used in the system.

### Number of Threads

The number of threads used for OpenMP code is determined with the following priority (while the one for automatic parallelization is described in Section "4.2 Automatic Parallelization"):

- The value specified with the `num_threads` clause of the "parallel" directive
- The value specified with the `omp_set_num_threads` function
- The value of the environment variable `OMP_NUM_THREADS`
- The number of CPUs that can be used with jobs

When the dynamic thread adjustment feature is enabled, the number of threads is adjusted to be less or equal to the number of CPUs.

When the dynamic thread adjustment feature is disabled, the thread number determined by the above priority will be used and multiple threads can be executed on the same processor via time-sharing. Such parallel execution is not effective because the overhead of synchronization among threads is very high. Therefore, it is recommended to set the number of threads to be less than the limit on the number of CPUs.

### CPU binding for thread

Environmental variable `OMP_PROC_BIND` or `GOMP_CPU_AFFINITY` can control thread affinity. `GOMP_CPU_AFFINITY` takes precedence over `OMP_PROC_BIND`.

## 4.3.3 Implementation-Dependent Specifications

---

The OpenMP specification includes the following processor-dependent specification. For details on each item, refer to the OpenMP Specifications.

### Memory Model

If a variable is longer than 4 bytes or crosses a 4-byte boundary:

- The value of the variable written from two threads without synchronization may be undefined.
- The value read by a thread from a variable that is written from another thread without synchronization may be undefined.

### Internal Control Variables

The initial values for each of the internal control variables are as follows:

- `nthreads-var`: Maximum number of CPUs.
- `dyn-var`: FALSE.
- `run-sched-var`: `static` without chunk size.
- `def-sched-var`: `static` without chunk size.
- `bind-var`: FALSE.
- `stacksize-var`: the stack area size for threads. Default size is 8M bytes.
- `wait-policy-var`: `PASSIVE`.
- `thread-limit-var`: 2147483647.
- `max-active-levels-var`: 2147483647.
- `place-partition-var`: Default value is `cores`, which is equivalent to threads in the system.

## Dynamic Thread Adjustment Features

This system supports the dynamic thread adjustment described in Section "4.3.2.3 Notes during Execution".

By default, the dynamic adjustment features are off.

## Loop Statement

The type used to calculate the iteration count of a loop that is not nested is `long int`.

When `auto` is set for the internal control variable `run-sched-var`, the effect of the `schedule(runtime)` construct is set to `schedule(static)`.

## sections Construct

Assignment to a thread of a structured block in a `sections` construct is performed in the same way as a static schedule.

## single Construct

A single region is executed by the thread that encounters the region first.

## taskloop Construct

When `grainsize` and `num_tasks` clause are not specified, these values will be set appropriately considering the loop iteration.

An eight byte integer type is used to calculate the iteration count of a collapsed loop.

## critical Construct

This implementation do not support lock hint. Therefore there is no the effect of using a hint clause.

## atomic Construct

Two `atomic` regions will be executed independently (not exclusively), when a variable type to be updated is different. When the types match, it may be executed exclusively even if the address is different.

- When updating a logical type, complex type, 1 byte integer type, 2 byte integer type, or quadruple precision (16 byte) real type variable
- When updating array element and index expression is not the same
- The target expression has explicit or implicit type conversion

## omp\_set\_num\_threads Function

If the argument is not a positive integer, the effect of `omp_set_num_threads` function is same as the argument is 1. A value that exceeds the number of threads supported by the system must not be specified.

## omp\_set\_schedule Function

There are no schedule types that are implementation-dependent.

## omp\_set\_max\_active\_levels Function

A call to the `omp_set_max_active_levels` function will be ignored when it is performed from an explicit `parallel` region. It will also be ignored when the argument is an integer that less than 0.

## omp\_get\_max\_active\_levels Function

The `omp_get_max_active_levels` function can be called from anywhere in the program and it returns the value of the internal control variable `max-active-levels-var`.

## omp\_get\_place\_proc\_ids Function

The `omp_get_place_proc_ids` function returns identifiers of the available processors in the specified place. The definition of processor is the CPU controlled by operating system.

## omp\_init\_lock\_with\_hint and omp\_init\_nest\_lock\_with\_hint Function

This implementation do not support lock hint. Therefore there is no the effect of using these functions.

## Environment Variable OMP\_SCHEDULE

The behavior is different by compiler mode.

When the schedule type specified for `OMP_SCHEDULE` is invalid, the schedule type is ignored, and the following default schedule is used.

Compiler mode	Default value
Trad Mode	<code>static</code> without chunk size.
Clang Mode	<code>static</code> with chunk size 1.

When the schedule type specified for `OMP_SCHEDULE` is `static`, `dynamic`, or `guided`, and the chunk size is not a positive number, the chunk size will be as follows:

Compiler mode	Schedule type	Chunk size
Trad Mode	<code>static</code>	no chunk size
	<code>dynamic</code>	1
	<code>guided</code>	
Clang Mode	<code>static</code>	1
	<code>dynamic</code>	
	<code>guided</code>	

#### Environment Variable `OMP_NUM_THREADS`

When a value equal to 0 is specified for the list of `OMP_NUM_THREADS`, it works in the same ways as when 1 is specified. When a value less than 0 is specified for the list of `OMP_NUM_THREADS`, the number of threads is determined by "[4.3.2.3 Notes during Execution](#)" procedure. A value that exceeds the number of threads supported by the system must not be specified.

#### Environment Variable `OMP_PROC_BIND`

If the value does not conform to the specified format for `OMP_PROC_BIND`, the value is ignored, and the default value (`close`) is used.

If all of the following conditions are met, the excess threads are assigned to places evenly as possible.

- `spread` or `close` is specified to `OMP_PROC_BIND`.
- T (number of threads) is greater than P (number of places).
- P does not divide T evenly.

#### Environment Variable `OMP_PLACES`

If the value does not conform to the specified format, the value is ignored, and the default value (`cores`) is used.

#### Environment Variable `OMP_DYNAMIC`

When a value other than `TRUE` or `FALSE` is specified for `OMP_DYNAMIC`, the value is ignored and the default value (`FALSE`) is used.

#### Environment Variable `OMP_NESTED`

When a value other than `TRUE` or `FALSE` is specified for `OMP_NESTED`, it is ignored and the default value (`FALSE`) is used.

#### Environment Variable `OMP_STACKSIZE`

When the value specified for `OMP_STACKSIZE` does not meet the defined format, it is ignored and the default value (8M bytes) is used.

#### Environment Variable `OMP_WAIT_POLICY`

`ACTIVE` performs spin wait. `PASSIVE` performs suspend wait.

#### Environment Variable `OMP_MAX_ACTIVE_LEVELS`

When the value specified for `OMP_MAX_ACTIVE_LEVELS` is an integer that is less than 0, it is ignored and the default value (2147483647) is used.

#### Environment Variable `OMP_THREAD_LIMIT`

When the value specified for `OMP_THREAD_LIMIT` is not a positive integer, it is ignored and the default value (2147483647) is used.

## 4.3.4 Notes on OpenMP Programming

This section provides notes on OpenMP programming.

### 4.3.4.1 Implementation of `parallel` Region and Explicit `task` Region

The structured block within a `parallel` construct or a `task` construct is compiled as an internal function.

The name of the internal function generated from the `parallel` construct or the `task` construct is different depending on the compiler mode. The name of an internal function is shown below.

Compiler mode	Type of internal function	Internal function name
Trad Mode	Internal function generated from <code>parallel</code> construct	The name " <code>._OMP_id-numberA</code> " is added
	Internal function generated from <code>task</code> construct	The name " <code>._TSK_id-numberB</code> " is added
Clang Mode	Internal function generated from <code>parallel</code> construct	<code>.omp_outlined.[_debug__][.id-numberC](*1)(*2)</code>
	Internal function generated from <code>task</code> construct	<code>.omp_task_entry.[.id-numberC](*3)</code>

*id-numberA*: Unique number that identifies `parallel` construct in a source file.

*id-numberB*: Unique number that identifies `task` construct in a source file.

*id-numberC*: Unique number in a source file. When plurals of `parallel` constructs and `task` constructs exist in the source file, a unique number to identify the `parallel` construct and the `task` construct collectively is assigned to "`.id-numberC`".

\*1) When `-g` option is specified, "`_debug__`" is added.

\*2) When two or more `parallel` constructs are specified in the source file, "`.id-numberC`" is added.

\*3) When two or more `task` constructs are specified in the source file, "`.id-numberC`" is added.

### 4.3.4.2 Implementation of `threadprivate` Variable

The `threadprivate` variable is implemented by using the Thread-Local Storage. When the size of the Thread-Local Storage exceeds the range of the offset, an error occurs at link time.

Example of error message:

```
relocation truncated to fit: R_AARCH64_TLSLE_ADD_TPREL_HI12 against symbol 'varname'
```

When the error occurs, specify the size of an appropriate offset by `-Ktls_size={12|24|32|48}` option for Trad Mode or `-mfj-tls-size={12|24|32|48}` option for Clang Mode. For details about the `-Ktls_size` option, see Section "2.2.2.6 -K Option". For details about the `-mfj-tls-size` option, see Section "9.1.2.2.6 Options for Code Generation".

### 4.3.4.3 Automatic Parallelization for OpenMP Programs

The `-Kopenmp` and `-Kparallel` options can be specified at the same time.

If both options are specified, automatic parallelization is not applied to the following loops:

- A loop within the OpenMP directives.
- A loop that statically includes OpenMP directives.
- In the case that a loop is parallelized by OpenMP directives,
  - If the loop itself is parallelized by OpenMP directives.
  - If the loop is lexically contained in a loop that is parallelized by OpenMP directives.



### 4.3.4.4 OpenMP Directive

If the OpenMP directive is specified in the constructor, destructor, operator `new()`, or operator `delete()`, a software malfunction may occur.

### 4.3.4.5 Creating Shared Libraries in Clang Mode

To create the following shared libraries in Clang Mode, specify the compiler option `-fopenmp`. If created without the compiler option `-fopenmp`, the program linking the shared library may terminate abnormally.

- A shared library that links a shared library created with the compiler option `-fopenmp`
- A shared library that is dynamically opened with a function such as the function `dlopen` and used from a program created with the compiler option `-fopenmp`

#### Example

- Example 1

```
$ fccpx shraed1.c -Nclang -fopenmp -shared -o libsampl_e_linked.so
$ fccpx shared2.c -Nclang -fopenmp -shared -o libsampl_e.so -L. -lsampl_e_linked
```

The linked `libsampl_e_linked.so` was created with the compiler option `-fopenmp`, so specify the compiler option `-fopenmp` when creating `libsampl_e.so` as well.

- Example 2

```
$ fccpx main.c -Nclang -fopenmp -o main.exe
$ fccpx shared.c -Nclang -fopenmp -shared -o libsampl_e.so
```

The program `main.exe` that dynamically opens `libsampl_e.so` was created with the compiler option `-fopenmp`, so specify the compiler option `-fopenmp` when creating `libsampl_e.so` as well.

## 4.4 Runtime Messages

LLVM OpenMP Library outputs runtime messages to standard error as following format.

```
OMP: type message
```

The meaning of runtime messages of LLVM OpenMP Library is explained as follows.

OMP:

Message prefix of LLVM OpenMP library.

*type*

Message type. The one of the following values is output.

Value	Meanings
Hint	Hint about OpenMP.
Info	Information about OpenMP.
Warning	Warning about OpenMP. Not fatal error. Even if this message is displayed, execution continues.
Error	Fatal error about OpenMP. If this message is displayed, execution is aborted.
System error	Fatal error about the system. If this message is displayed, execution is aborted.

*message*

Content of message



## Example

---

```
$ export OMP_STACKSIZE=1Z
$ ./a.out
OMP: Warning #80: OMP_STACKSIZE="1Z": value too large.
OMP: Info #107: OMP_STACKSIZE value "9223372036854775807" will be used.
OMP: Error #34: System unable to allocate necessary resources for OMP thread:
OMP: System error #11: Resource temporarily unavailable
OMP: Hint Try decreasing the value of OMP_NUM_THREADS.
```

The above example is series of messages when unacceptable value in FX system is set to environmental value OMP\_STACKSIZE as stack size.

---



## Note

---

The following environment variable is invalidated for the control of the runtime message of LLVM OpenMP Library. For details about the environment variable, see Section "[2.5.1 Environment Variable for Execution](#)".

- Environment variable FLIB\_C\_MESSAGE
-

# Chapter 5 Emitting Information

This chapter provides information on C++ language program compilation and execution.

## 5.1 Emitting Information at Compilation

This section explains the information emitted by this system at compilation.

### 5.1.1 Header

When either the `-Nlst`, `-Nlst_out=file`, `-Nsrc` or `-Nsta` compilation option is specified, a header is emitted for each type of information.

#### 5.1.1.1 Output Format

The header is emitted as shown in the following format.

Header Format

```
Fujitsu C/C++ Version version date
```

<i>version</i>	Language processing system version.
<i>date</i>	Compilation date and time in the format of the <code>asctime</code> function.

### 5.1.2 Source List

When the program is compiled, the source list is put out according to the following options.

- `-Nlst[={p|t}]`

Specifies to output source list and statistics list to file(s).

- `-Nlst_out=file`

Specifies the filename to output source list and statistics list.

When this option is specified, the `-Nlst=p` option is also effective.

- `-Nsrc`

Specifies to output source list to the standard output.

Note that when the `-Nlst` or `-Nlst_out=file` option is specified, the source list is output to the *file*.

The source list includes symbols indicating optimization and parallelization if they are performed.

#### 5.1.2.1 Output Format

The source list is emitted as shown in the following format.

Source List Format

```
Compilation information
Current directory : directory-name
Source file      : source-file-name
(line-no.)(optimize)
nnnnnnnn pi nnnnnnv source
...
                <<< Loop-information Start >>>
                <<< details optimization information
                <<< Loop-information End >>>
nnnnnnnn pi nnnnnnv source
...
```

<i>directory-name</i>	The name of directory where source file is stored
<i>source-file-name</i>	Source file name
<i>nnnnnnnn</i>	Line number of source (variable length)
<i>p</i>	Symbols for parallelization
<i>i</i>	Symbols for inline expansion
<i>mmmm</i>	The number of loop unrolling (variable length)
<i>v</i>	Symbols for using SIMD Extensions
<i>source</i>	Source line
<<< Loop-information Start >>>(a)	Details of optimization and parallelization information header
<i>details optimization information(a)</i>	Details of optimization and parallelization information which has been performed to the next statement.
<<< Loop-information End >>>(a)	Details of optimization and parallelization information footer

a) Emitted only if the -Nlst=t option is effective.

## 5.1.2.2 Information Included in Source List

### 5.1.2.2.1 Line Number of Source

Line number of source is emitted.

### 5.1.2.2.2 Symbols for Parallelization

The symbols for parallelization are emitted by the following conditions:

- Symbols for loop control statement are emitted at the line including loop control statements.  
The loop control statement is "for" statement, "while" statement, "do-while" statement, and "if-goto" statement.
- If the parallelization contains more than one loop control statement, the symbols for the left most loop control line are emitted.

Symbols are classified in two different types; one for loop control statements and one for anything other than loop control statements.

#### Symbols for loop control statements

<i>p</i>	Indicates all statements in this loop have been parallelized. It becomes "pp" for -Nlst=t in the first line within the parallelization is performed.
<i>m</i>	Indicates that only some of the statements in this loop have been parallelized.
<i>s</i>	Indicates none of the statements in this loop have been parallelized.
Blank	Indicates none of the statements in this loop are targeted for parallelization.

#### Symbols for anything other than loop control statements

<i>p</i>	Indicates all statements in this loop have been parallelized. <sup>[a]</sup>
<i>m</i>	Indicates that only some of the statements in this loop have been parallelized. <sup>[b]</sup>
<i>s</i>	Indicates none of the statements in this loop have been parallelized.
Blank	Indicates none of the statements in this line are targeted for parallelization.

[a] When the mark "s" is displayed for a loop control statement, the parallelization is not applied because parallelization is available for statements in the loop but worthless.

[b] When the mark "s" is displayed for a loop control statement, the parallelization is not applied because parallelization is available for some of statements in the loop but worthless.

### 5.1.2.2.3 Symbols for Inline Expansion

If inline expansion is performed, "i" is emitted.

Otherwise, blank is emitted.

### 5.1.2.2.4 Number of Loop Unrolling

If loop unrolling is performed, multiplicity of loop unrolling is emitted.

If the loop full unrolled, "f" is emitted instead of multiplicity of loop unrolling.

Otherwise, blank is emitted.

### 5.1.2.2.5 Symbols for Using SIMD Extensions

The symbols for the optimization that uses SIMD Extensions are emitted by the following conditions:

- Symbols for the optimization that uses SIMD Extensions are emitted at the line including loop control statements. The loop control statement is "for" statement, "while" statement, "do-while" statement, and "if-goto" statement.
- If the optimization that uses SIMD Extensions contains more than one loop control statement, the symbols for the left most loop control line are emitted.

Symbols are classified in two different types; one for loop control statements and one for anything other than loop control statements.

Symbols for loop control statements

v	Indicates all statements in this loop have been optimized using SIMD Extensions.
m	Indicates that only some of the statements in this loop have been optimized using SIMD Extensions.
s	Indicates none of the statements in this loop have been optimized using SIMD Extensions.
Blank	Indicates none of the statements in this loop are targeted for optimization using SIMD Extensions.

Symbols for anything other than loop control statements

v	Indicates all statements in this loop have been optimized using SIMD Extensions. <sup>[a][c]</sup>
m	Indicates that only some of the statements in this loop have been optimized using SIMD Extensions. <sup>[b]</sup> <sub>[c]</sub>
s	Indicates none of the statements in this loop have been optimized using SIMD Extensions.
Blank	Indicates none of the statements in this line are targeted for optimization using SIMD Extensions.

[a] When the mark "s" is displayed for a loop control statement, SIMD optimization is not applied because SIMD optimization is available for statements in the loop but worthless.

[b] When the mark "s" is displayed for a loop control statement, SIMD optimization is not applied because SIMD optimization is available for some of statements in the loop but worthless.

[c] When SIMD optimization is applied to a loop control statement and the mark "v" or "m" is displayed for the statement, the mark "v" could be displayed for statements in the loop though SIMD extensions are not used for them. The mark "s" is displayed for a statement preventing SIMD optimization, which is a clue to tune programs.

### 5.1.2.2.6 Details Optimization Information

When the -Nlst=t option is specified, the source list containing details of optimization and parallelization information which has been performed are written to file(s).

The following optimization information are put just before each loop.

## About Automatic Parallelization

- Standard iteration count: *N*

When the iteration count of a loop is *N* or more, the loop is brought to parallel execution. When less than *N*, it is brought to serial execution.

## Optimization Information In Loop Units

- INTERCHANGED(*nest: nest-no*)

It means that loop was interchanged.

- *nest-no* is the nesting level of the loop that was moved by loop interchange optimization.
- (*nest:*) may not be indicated if the other optimization like loop collapsing was applied.

- FUSED(*lines: num1, num2[ , num3] . . .*)

It means that loops were fused.

Loops at line *num2* and subsequent lines were fused into a loop at line *num1*.

- (*lines:*) is not indicated at line *num2* and subsequent lines, that were fused to line *num1*.

- FISSION(*num: N*)

It means that loop was split.

- *N* is the number of loops after fission.

- COLLAPSED

It means that nested loops were collapsed into a single loop.

- SOFTWARE PIPELINING(*IPC: ipc, ITR: itr, MVE: mve, POL: pol*)

It means that software pipelining was applied to the loop.

- *ipc* is the expectation value of the number of executed instructions per cycle (Instructions Per Cycle) when the software pipelining is applied to the loop.
- *itr* is the minimum iteration count required to choose the software pipelined loop. It is the same as the value output by the optimization message jwd8205o-i.
- *mve* is the loop expansion number by the software pipelining.
- *pol* is the applied instruction scheduling algorithm.  
S means that the algorithm fit for a small loop is applied.  
L means that the algorithm fit for a large loop is applied.  
These algorithms correspond to the ones applied by -Kswp\_policy=small option and -Kswp\_policy=large option, respectively.

- SIMD

It means that SIMD instructions were generated for the loop. It is displayed by one of the following.

- SIMD(*VL: length[ , length] . . .*)  
It means that SIMD instructions were generated for the loop. A SIMD instruction treats *length* elements of array. When loop fission is applied to the loop and the values of *length* for each loop are different, two or more *lengths* are displayed.
- SIMD(*VL: AGNOSTIC; VL: length[ , length] . . . in 128-bit*)  
It means that SIMD instructions were generated for the loop without regarding the SVE vector register size as a specific size. A SIMD instruction treats *length* elements of array if the size of the SVE vector register is 128-bit. When loop fission is applied to the loop and the values of *length* for each loop are different, two or more *lengths* are displayed.

- STRIPING

It means that loop was performed loop striping.

- MULTI-OPERATION FUNCTION

It means that loop includes multi-operation function.

- PATTERN MATCHING (matmul)

It means that loop was changed to library function call (matmul).

- UNSWITCHING

It means that loop was performed loop unswitching.

- FULL UNROLLING

It means that loop was fully unrolled.

- CLONE

It means that clone optimization was applied to the loop.

- LOOP VERSIONING

It means that loop versioning was applied to the loop.

### Information That Relates To The Prefetch

#### Hardware-prefetch

- PREFETCH(HARD) Expected by compiler:

It indicates that the compiler expects hardware-prefetch and does not generate the prefetch instruction for array data accessed sequentially within a loop.

- array name,...

It indicates an array name expecting to use hardware-prefetch. The array which was generated by compiler is shown as "(unknown)".

#### Prefetch instruction

- PREFETCH(SOFT): *N*

It indicates the number of all prefetch instructions that exist in the loop.

In addition, the number and the array name of the prefetch instruction of each access type of the prefetch are shown in the form of the following.

```
access type: N
array name: N,...
```

- access type

- SEQUENTIAL: *N*

It indicates the number of prefetch instructions for array data accessed sequentially within a loop.

- STRIDE: *N*

It indicates the number of prefetch instructions for array data that is accessed with a stride larger than the cache line size used in the loop.

- INDIRECT: *N*

It indicates the number of prefetch instructions for array data that is accessed indirectly (list access) within a loop.

- SPECIFIED: *N*

It indicates the number of prefetch instructions generated by Built-in Functions `__builtin_prefetch`.

- array name

- array name: *N*,...

It indicates the number of prefetch instructions for the array name. The array which was generated by compiler is shown as "(unknown)".

### Information That Relates To The zfill Optimization

- ZFILL

Indicates the zfill optimization is applied.

- *array name*: ...

Indicates arrays names to which the zfill optimization is applied. An array generated by compiler is shown as "(unknown)".

### Information That Relates To The Register

- SPILLS :

Indicates the number of saving and restoring instructions for registers that exist in the innermost loop shown for each register type.

- GENERAL : SPILL *N* FILL *N*

Indicates the number of saving and restoring instructions for general registers to and from the memory.

- SIMD&FP : SPILL *N* FILL *N*

Indicates the number of saving and restoring instructions for SIMD and floating-point registers to and from the memory.

- SCALABLE : SPILL *N* FILL *N*

Indicates the number of saving and restoring instructions for scalable vector registers to and from the memory.

- PREDICATE : SPILL *N* FILL *N*

Indicates the number of saving and restoring instructions for predicate registers to and from the memory.

## 5.1.2.3 Example of Source Listing

Example of source listing is shown in following list.

### Example of Source Listing 1

```
Compilation information
Current directory : directory-name
Source file      : source-file-name
(line-no.)(optimize)
 1          #include <cstdio>
 2
 3          float sub(int i);
 4
 5          int main() {
 6              int i;
 7              float a[1000];
 8
 9  p      2v  for(i = 0; i < 1000; i++) {
10  p      2v  a[i] = i;
11  p      2v  }
12
13          printf("%lf\n", a[0]);
14
15          i = 0;
16          s  _loop:
17          s  if(i < 1000) {
18              goto _next;
19          }
20          s  a[i] = i;
21          s  i++;
22          s  goto _loop;
23          s  _next:
24
25          printf("%lf\n", a[0]);
26
27          4  for(i = 0; i < 1000; i++) {
```



```

28      4      a[i] = sub(i);
29      4      }
30
31      f      for(i = 0; i < 2; i++) {
32      f          a[i] = i;
33      f      }
34
35          printf("%lf\n", a[0]);
36          return 0;
37      }
38
39      float sub(int j) {
40          return (float)j;
41      }

```

In this sample, following information can be read.

- The statements included in the "for"-statement at line-9 are parallelized.
- The statements included in the "for"-statement at line-9 are optimized using SIMD Extensions.
- The statements included in the "for"-statement at line-9 and line-27 are performed with loop unrolling.
- SIMD Extensions is not used at the loop constructed with "if-goto" statement from line-16 to line-23.
- The statements included in the "for"-statement at line-31 are full unrolled.

#### Example of Source Listing 2

```

Compilation information
Current directory : directory-name
Source file      : source-file-name
(line-no.)(optimize)
 1      #include <cstdio>
 2
 3      float sub(int i);
 4
 5      int main() {
 6          int i;
 7          float a[10000];
 8
          <<< Loop-information Start >>>
          <<< [PARALLELIZATION]
          <<< Standard iteration count: 1000
          <<< [OPTIMIZATION]
          <<< SIMD(VL: 16)
          <<< SOFTWARE PIPELINING(IPC: 3.00, ITR: 80, MVE: 5, POL: S)
          <<< Loop-information End >>>
 9  pp    2v    for(i = 0; i < 10000; i++) {
10  p     2v    a[i] = i;
11  p     2v    }
12
13          printf("%lf\n", a[0]);
14
15          i = 0;
16          s    _loop:
17          s    if(i < 10000) {
18              goto _next;
19          }
20          s    a[i] = i;
21          s    i++;
22          s    goto _loop;
23          s    _next:
24
25          printf("%lf\n", a[0]);

```

```

26          <<< Loop-information Start >>>
          <<< [PARALLELIZATION]
          <<<   Standard iteration count: 1000
          <<< [OPTIMIZATION]
          <<<   SIMD(VL: 16)
          <<<   SOFTWARE PIPELINING(IPC: 3.00, ITR: 80, MVE: 5, POL: S)
          <<< Loop-information End >>>
27 pp      v   for(i = 0; i < 10000; i++) {
28 pi      v   a[i] = sub(i);
29          }
30
          <<< Loop-information Start >>>
          <<< [PARALLELIZATION]
          <<<   Standard iteration count: 1000
          <<< [OPTIMIZATION]
          <<<   SIMD(VL: 8)
          <<< Loop-information End >>>
31 s      v   for(i = 0; i < 8; i++) {
32 p      v   a[i] = i;
33 p      v   }
34
35          printf("%lf\n", a[0]);
36          return 0;
37          }
38
39          float sub(int j) {
40              return (float)j;
41          }
Total prefetch num: 0

```

In this sample, following information can be read.

- The loop by "for"-statement at line-9 is parallelized if its iteration count is greater than or equal to 1000.
- The loop by "for"-statement at line-9 is optimized with software pipelining.
- The loop by "for"-statement at line-9 is parallelized.
- The statements included in the "for"-statement at line-9 is optimized using SIMD Extensions.
- The statements included in the "for"-statement at line-9 is optimized with loop unrolling.
- The statements at line-16, line-17, line-20, line-21, line-22, and line-23 are not optimized using SIMD Extensions.
- The loop by "for"-statement at line-27 is parallelized if its iteration count is greater than or equal to 1000.
- The loop by "for"-statement at line-27 is optimized with software pipelining.
- The loop by "for"-statement at line-27 is parallelized.
- The statements included in the "for"-statement at line-27 is optimized using SIMD Extensions.
- The function called at line-28 is optimized with inline expansion.
- The loop by "for"-statement at line-31 is parallelized if its iteration count is greater than or equal to 1000.
- The loop by "for"-statement at line-31 is optimized using SIMD Extensions.
- The statement at line-32 is not parallelized because the parallelization is not effective.
- The statement at line-33 is not parallelized because the parallelization is not effective.

#### 5.1.2.4 Notes on Information at Compilation (Source List)

The -Nlst=p option, -Nlst=t option and -Nsrc option could output wrong compilation information (optimization information) in the following cases.

- When a function is inline expanded, different optimizations could be applied to each line. In this case, the compiler could output information incorrectly as follows.
    - Output messages redundantly.
    - Output optimization messages inconsistent with compilation information (optimization information).
    - Output no compilation information (optimization information).
- In addition, PREFETCH: *N*, which is the number of prefetch instructions in the inlined function, includes all prefetch instructions expanded in multiple lines.
- The compiler applies multiple optimizations such as SIMDization, automatic parallelization, loop fusion, and loop distribution to one loop. In this case, the compiler could output information incorrectly as follows.
    - Output optimization information on line number to an incorrect line.
    - Output compilation information (optimization information) inconsistent with optimization messages.
  - Loop unswitching optimization creates a loop in each "TRUE" and "FALSE" blocks of the "if" statement, but the compiler outputs optimization message for only one of the loops. In addition, optimization information on line number may not be output.
  - When multiple loops are written in the same line in a program, the compiler outputs optimization information for only one of the loops. Please make sure to write one loop in one line in order to output optimization information.
  - Detailed optimization information for a loop which consists of "if-goto" statements is not output. Optimization information on line number could be output to an incorrect line.
  - The compiler may generate loops under the following condition. The optimization messages may be output for the generated loops, and compilation information (optimization information) may be output for the generated loops.
    - There is a loop parallelized automatically by -Karray\_private option or optimization control specifier array\_private, first\_private or last\_private.

### 5.1.3 Parallelization Messages

---

This parallelization messages indicate why the parallelization failed, or how the parallelization and optimization are performed.

Using the -Koptmsg=2 option with the -Kparallel option, the parallelization messages can be emitted to the standard error. An example of the parallelization message is shown in following.

Example of Parallelization Messages

```
Parallelization messages
jwd5001p-i  "source-filename", line n1: This loop with loop variable 'i' is parallelized.
jwd6001s-i  "source-filename", line n1: SIMD conversion is applied to this loop with the loop
variable 'i'.
jwd8202o-i  "source-filename", line n1: Loop unrolling expanding 8 times is applied to this loop.
jwd5122p-i  "source-filename", line n2: This loop is not parallelized because a function call
that prevents parallelization exists in the loop.
```

### 5.1.4 Statistics List

---

When the program is compiled, the statistics list is put out according to the following options.

- -Nlst[={p|t}]
 

Specifies to output source list and statistics list to file(s).
- -Nlst\_out=*file*

Specifies the filename to output source list and statistics list.

When this option is specified, the -Nlst=p option is also effective.
- -Nsta
 

Specifies to output statistics list to the standard output.

Note that when the `-Nlst` or `-Nlst_out=file` option is specified, the statistics list is output to the *file*.

Only effective optimization options are emitted in the statistics list. The options which are not emitted are invalid.

### 5.1.4.1 Output Format

The statistics list is emitted as shown in the following format.

#### Source List Format

```
Statistics information
Option information
  Profile file           : Options that specified in compilation profile file[a]
  Environment variable  : Options that specified in environment variables to set compiler options[a]
  Command line options  : Options that specified in compile command
  Effective options     : Effective optimization options
```

[a] The output is on one line in the order of "<mode common> <mode specific>".

Note that the "Profile file" and "Environment variable" lines are not emitted in the following cases:

- The compilation profile file does not exist, or the compilation profile file size is zero
- Environment variables to set compiler options do not exist

### 5.1.4.2 Options Which Are Not Output

The statistics list outputs only effective options for optimization.

Options for `ld` command are not output.

- Options which are not output as follows:

```
-# / -### / -A / -Aname[(tokens)] / -B{dynamic|static} / -C / -Dname[=tokens] / -E / -H / -Idir /
-Ldir / -Nlst_out=file / -P / -S / -Uname / -Wtool,arg1[,arg2]... / -Yitem,dir / -V / -c / {-help|--
help} / -j / -lname / -o pathname / -shared / Message options (except -w) / Language specifications
options / PCH options
```

### 5.1.4.3 Options Which Are Interpreted by Compiler

The following options are interpreted by the compiler as follows.

For more the details about interpretation, see Section "[2.2.2.6 -K Option](#)".

- `-Kfast`
- `-Knoprefetch`
- `-Kvisimpact`

## 5.2 Information at Execution

---

This section describes the information emitted during execution.

### 5.2.1 Runtime Messages

---

For details about runtime messages displayed by this system:

- When the message starts with "jwe", refer to "Fortran/C/C++ Runtime Messages".
- For all other messages refer to the reference manual provided with this system.

## 5.2.2 Trace Back Information

---

When an error or trap occurs in execution, a trace back map is emitted to standard error if either the environment variable `FLIB_C_MESSAGE` is not set to "NO\_MESSAGE" or the `-NRtrap` option is specified.

Using this feature, it is possible to find the path from the main program to the program unit where the error occurred.

This information is obtained from the debug information stored in the heap memory allocated at startup. If the information is not output, the information may be output by setting the size of the heap memory larger than the default value using the environment variable `FLIB_TRACEBACK_MEM_SIZE`. If the information cannot be obtained due to insufficient heap memory, a diagnostic message `jwe1531i-i` will be output. This diagnostic message is output during error processing, so the order in which the diagnostic messages are output may be incorrect.

For details about the environment variables `FLIB_C_MESSAGE` and `FLIB_TRACEBACK_MEM_SIZE`, see Section "[2.5.1 Environment Variable for Execution](#)".

For details about the `-NRtrap` option, see Section "[2.2.2.7 -N Option](#)".

For details about Trace Back Information, refer to "Fortran User's Guide".

# Chapter 6 Language Specifications

This chapter describes the language specifications based on the C++ language standard and extended under this processing system, and the support status of the language specifications.

Note that the following syntax description symbols are used in this chapter.

- The nonterminal symbols (syntactical elements) are indicated by italic text. Terminal symbols (literals and elements of character sets) are represented by roman text, including special symbols.
- A colon (:) appearing after a nonterminal symbol indicates that the text after the colon is a definition of that nonterminal symbol.
- Except when prefaced by "one of the following", selectable definitions appear on separate lines.
- Optional symbols that may be omitted are indicated with the subscript *opt*. Therefore, in {*expression*<*opt*>}, the curly brackets indicate the range that may be omitted.

## 6.1 Language Specifications Extended in This System

### 6.1.1 long long Type

In addition to the four signed integer types of `signed char`, `short int`, `int` and `long int`, the `long long int` type is added.

Suffixes

*integer-suffix:*

*unsigned-suffix long-word-suffix*<*opt*>

*long-word-suffix unsigned-suffix*<*opt*>

*unsigned-suffix long-long-word-suffix*<*opt*>

*long-long-word-suffix unsigned-suffix*<*opt*>

*unsigned-suffix:*

One of the following

- `u`

- `U`

*long-long-word-suffix:*

*long-word-suffix long-word-suffix*

*long-word-suffix:*

One of the following

- `l`

- `L`

Meaning Rules

The compiler uses the first data type in the following list in which the value can be represented, depending on the size of the constant.

Decimal number with no suffix:

`int`, `long int`, `unsigned long int`, `long long int`, `unsigned long long int`

Octal or hexadecimal number with no suffix:

`int`, `unsigned int`, `long int`, `unsigned long int`, `long long int`, `unsigned long long int`

If the letter `u` or `U` is added as a suffix:

`unsigned int`, `unsigned long int`, `unsigned long long int`

If the letter `l` or `L` is added as a suffix:

```
long int, unsigned long int, long long int, unsigned long long int
```

If both the letter `u` or `U` and the letter `l` or `L` are added as suffixes:

```
unsigned long int, unsigned long long int
```

If both the letter `l` or `L` and the letter `l` or `L` are added as suffixes:

```
long long int, unsigned long long int
```

### Limitations of Type Specifiers

The various lists of type specifiers must belong to one of the following groups. If there are two or more groups in one item, they are delimited with commas. The type specifiers may appear in any order, and may be mixed with other declaration specifiers.

- `long long`, `signed long long`, `long long int`, `signed long long int`
- `unsigned long long`, `unsigned long long int`

Code Samples:

```
long long int lli=10LL;
```

```
unsigned long long int ulli=10ULL;
```

## 6.1.2 The pragma Directive

---

The `pragma` directive directs the operation of the processing system.

Meaning

`#pragma ident string-literal new-line`

Directs the processing system to add the string in *string-literal* to the object file as a comment. This is a `#pragma` directive that has the same meaning as the `#ident` directive. Typically, this is used to add version control information to the object file.

`#pragma weak identifier new-line`

Directs the processing system to define *identifier* in the object file as a weak symbol.

`#pragma weak identifier1=identifier2 new-line`

Directs the processing system to define *identifier1* in the object file as a weak symbol having the same value and type as *identifier2*.

`#pragma unknown_control_flow (list-of-identifiers) new-line`

Informs the processing system that the function specified in *list-of-identifiers* has control flow unknown to the processing system.

`#pragma redefine_extname <old function name> <new function name> new-line`

Directs the processing system to replace the old external function name with the new external function name in the object code.

## 6.1.3 The ident Directive

---

The `ident` directive specifies that comments are to be added to the C program.

Meaning

`#ident string-literal new-line`

Directs the processing system to add the specified *string-literal* to the ".comment section" of the object file. The ".comment section" is not loaded into memory when the program is executed.

`#ident preprocessor-token new-line`

Directs the processing system to process the specified *preprocessor-token* as normal text, and the various identifiers currently defined as macro names are replaced with a list of their replacements. All `#ident` directives after replacement must match their format prior to replacement.

## 6.1.4 The assert Directive

The `assert` directive defines predicate names and assigns assertion tokens to predicate names.

### Meaning

`#assert identifier assertion-specifier<opt> new-line`

Defines that the *identifier* following the `#assert` directive is a predicate name. In addition, the list of identifiers in the *assertion-specifier* following the predicate name are defined to be assertion tokens.

The `#assert` directive defines predicate names and assigns assertion tokens to predicate names.

An `#assert` directive with no list of identifiers is effective until a corresponding `#unassert` directive appears.

### Various Identifiers of the # Operator Expression

The various identifiers of the `#` operator expression correspond to the predicate names and assertion tokens specified in the `#assert` directive, respectively. If an assertion token is assigned to a predicate name, it evaluates to 1; otherwise it evaluates to 0. For example, `#system(unix)` becomes 1.

The following assertions are predefined by the processing system (predefined assertions):

- `#assert system(unix)`

### Code Samples:

```
#assert langlevel(ansi)
#assert langlevel(sysv)
```

The above directives define the predicate `langlevel` and assign the assertions `ansi` and `sysv` to the predicate.

```
#assert float
```

The above directive defines the predicate `float`. It has no assertions.

## 6.1.5 The unassert Directive

The `unassert` directive cancels the assignment of predicate names and assertion tokens.

### Meaning

`#unassert identifier assertion-specifier<opt> new-line`

Cancels the assignment of assertion tokens. An `#unassert` directive with no list of *identifiers* (*assertion-specifiers*) cancels the assignment of all assertion tokens to predicate names.

If the specified *identifier* is not a predicate name, the `#unassert` directive is ignored. Also, if an unassigned assertion token is specified, the `#unassert` directive is ignored.

## 6.1.6 Predefined Macro Names

Some of the predefined macros are changed by the compiler option `-std=level`.

The following tables show the values of some predefined macros. To display all predefined macros, specify the compiler options `-E` and `-dM`.

Table 6.1 Predefined Macros whose value is changed by the compiler option `-std=level`

Identifier	The level of language specifications specified with the compiler option <code>-std=level</code>							
	c++98 c++03	gnu++98 gnu++03	c++11	gnu++11	c++14	gnu++14	c++17	gnu++17
<code>__STDCPP_THREADS__</code>	-		1		1		1	
<code>__STRICT_ANSI__</code>	1	-	1	-	1	-	1	-
<code>__cplusplus</code>	199711L		201103L		201402L		201703L	



Identifier	The level of language specifications specified with the compiler option <code>-std=level</code>							
	c++98 c++03	gnu++98 gnu++03	c++11	gnu++11	c++14	gnu++14	c++17	gnu++17
linux	-	1	-	1	-	1	-	1
unix	-	1	-	1	-	1	-	1

-: Undefined

Table 6.2 Predefined Macros whose value is not changed by the compiler option `-std=level`

Identifier	Value
<code>_BOOL</code>	1
<code>_GNU_SOURCE</code>	1
<code>_LP64</code>	1
<code>_OPENMP (*1)</code>	201107
<code>_REENTRANT (*2)</code>	1
<code>_WCHAR_T</code>	1
<code>__ARM_ARCH</code>	8
<code>__ARRAY_OPERATORS</code>	1
<code>__DATE__</code>	The date of compilation (in the format of the <code>asctime</code> function)
<code>__ELF__</code>	1
<code>__EXCEPTIONS (*3)</code>	1
<code>__FCC_major__</code>	The major version number of the compiler
<code>__FCC_minor__</code>	The minor version number of the compiler
<code>__FCC_patchlevel__</code>	The patch level of the compiler
<code>__FCC_version__</code>	The string which represents the version of the compiler
<code>__FILE__</code>	Source file name
<code>__FJ_RUNTIME_USES_NAMESPACES</code>	1
<code>__FMTL_PARALLEL (*4)</code>	1
<code>__FMTL_SERIAL (*5)</code>	1
<code>__FMTL_SSL2 (*6)</code>	1
<code>__FMTL_SSL2BLAMP (*7)</code>	1
<code>__FUJITSU</code>	1
<code>__GNUC__</code>	6
<code>__GNUC_MINOR__</code>	1
<code>__GNUC_PATCHLEVEL__</code>	0
<code>__GNUG__</code>	6
<code>__LINE__</code>	Line number of source file
<code>__LP64__</code>	1
<code>__MT__ (*8)</code>	1
<code>__OPTIMIZE__ (*9)</code>	1
<code>__PIC__</code>	1 (*10)

Identifier	Value
	2 (*11)
__PLACEMENT_DELETE	1
__PRAGMA_REDEFINE_EXTNAME	1
__PTRDIFF_TYPE__	long int
__RTTI	1
__SIGNED_CHARS__	1 (*12)
	- (*13)
__SIZE_TYPE__	long unsigned int
__STDC__	1
__STDC_HOSTED__	1
__TIME__	The time of compilation (in the format of the <code>asctime</code> function)
__WCHAR_TYPE__	int (*14)
__aarch64__	1
__linux	1
__linux__	1
__pic__	1 (*10)
	2 (*11)
__unix	1
__unix__	1

-: Undefined

\*1) When the compiler option `-Kopenmp` or `-fopenmp` is set

\*2) When the compiler option `-mt` or `-pthread` is set

\*3) When the compiler option `-Nexceptions` or `-fexceptions` is set

\*4) When the compiler option `-Nfmtl=parallel` is set

\*5) When the compiler option `-Nfmtl=serial` is set

\*6) When the compiler option `-Nfmtl=SSL2` is set

\*7) When the compiler option `-Nfmtl=SSL2BLAMP` is set

\*8) When the compiler option `-mt` is set

\*9) When the compiler option `-O1` or higher is set

\*10) When the compiler option `-fpic` or `-fpie` is set

\*11) When the compiler option `-fPIC` or `-fPIE` is set

\*12) When the compiler option `-fsigned-char` is set

\*13) When the compiler option `-funsigned-char` is set

\*14) This value is different from Clang Mode, and unintended behavior may occur when you execute a program linking an object generated in Trad Mode and an object generated in Clang Mode.

## 6.2 Support Status of Language Specifications

This processing system supports the following C++ language standard:

- C++03
- C++11
- C++14
- Part of C++17

## 6.2.1 C++03 Specifications

---

Following language specifications are not supported in this processing system.

- `export` keyword in template

## 6.2.2 C++11 Specifications

---

The C++11 specifications supported by this system are as follows.

Table 6.3 Support Status of C++11 Specifications (Language Features)

Language Features	Support
Rvalue references	Yes
Rvalue references for <code>*this</code>	Yes
Initialization of class objects by rvalues	Yes
Non-static data member initializers	Yes
Variadic templates	Yes
Extending variadic template template parameters	Yes
Initializer lists	Yes
Static assertions	Yes
auto-typed variables	Yes
Multi-declarator auto	Yes
Removal of auto as a storage-class specifier	Yes
New function declarator syntax	Yes
Lambda expressions	Yes
Declared type of an expression	Yes
<code>decltype</code> and call expressions	Yes
Right angle brackets	Yes
Default template arguments for function templates	Yes
Solving the SFINAE problem for expressions	Yes
Template aliases	Yes
Extern templates	Yes
Null pointer constant	Yes
Strongly-typed enums	Yes

Language Features	Support
Forward declarations for enums	Yes
Standardized attribute syntax	Yes [a]
Generalized constant expressions	Yes
Alignment support	Yes
Delegating constructors	Yes
Inheriting constructors	Yes
Explicit conversion operators	Yes
New character types	Yes
Unicode string literals	Yes
Raw string literals	Yes
Universal character names in literals	Yes
User-defined literals	Yes
Standard Layout Types	Yes
Defaulted and deleted functions	Yes
Extended friend declarations	Yes
Extending sizeof	Yes
Inline namespaces	Yes
Unrestricted unions	Yes
Local and unnamed types as template arguments	Yes
Range-based for	Yes
Explicit virtual overrides	Yes
Minimal support for garbage collection and reachability-based leak detection	N/A
Allowing move constructors to throw [noexcept]	Yes
Defining move special member functions	Yes

[a] The carries\_dependency attribute can be compiled, but it is ineffectual.

**Table 6.4 Support Status of C++11 Specifications (Concurrency)**

Concurrency	Support
Sequence points	Yes
Atomic operations	Yes
Strong Compare and Exchange	Yes
Bidirectional Fences	Yes

Concurrency	Support
Memory model	Yes
Data-dependency ordering: atomics and memory model	Yes
Propagating exceptions	Yes
Allow atomics use in signal handlers	Yes
Thread-local storage	Yes
Dynamic initialization and destruction with concurrency	Yes

Table 6.5 Support Status of C++11 Specifications (C99 Features in C++11)

C99 Features in C++11	Support
<code>__func__</code> predefined identifier	Yes
C99 preprocessor	Yes
<code>long long</code>	Yes
Extended integral types	Yes

For the C++ Standard Library, the C++11 specifications except the following feature are supported.

- `<cuchar>`

## 6.2.3 C++14 Specifications

The C++14 specifications supported by this system are as follows.

Table 6.6 Support Status of C++14 Specifications

Language Features	Support
Tweak to certain C++ contextual conversions	Yes
Binary literals	Yes
<code>decltype(auto)</code>	Yes
Return type deduction for normal functions	Yes
Initialized lambda captures	Yes
Generic lambda expressions	Yes
Variable templates	Yes
Relaxing requirements on <code>constexpr</code> functions	Yes
Member initializers and aggregates	Yes
Clarifying memory allocation	N/A
<code>[[deprecated]]</code> attribute	Yes
Single-quotation-mark as a digit separator	Yes
C++ Sized Deallocation	Yes

## 6.2.4 C++17 Specifications

---

The following specifications of the C++17 language standard are available.

- New auto rules for direct-list-initialization
- typename in a template template parameter
- Attributes for namespaces and enumerators

# Chapter 7 Notes on Linking with Different Languages and Trad/Clang Modes

This chapter describes notes about linking with object programs in different programming languages and Trad/Clang Modes.

## 7.1 Compile Commands and Required Options when Linking

In this system, there are two types of compile commands for the Fortran, C, and C++ languages, a cross compiler and a native compiler. In addition, the compiler commands are used depending on whether the program uses the MPI library as shown in the table below.

The following description uses the compile command of the cross compiler (No use of MPI libraries). When using the MPI library or native compiler, replace the compile commands in the table below.

Programming Language	MPI Library	Cross Compiler	Native Compiler
Fortran	Not Use	frtpx	frt
	Use	mpifrtpx	mpifrt
C	Not Use	fccpx	fcc
	Use	mpifccpx	mpifcc
C++	Not Use	FCCpx	FCC
	Use	mpiFCCpx	mpiFCC

Object programs can be linked in a combination of Fortran, C Trad Mode, C Clang Mode, C++ Trad Mode, and C++ Clang Mode.

Depending on the combination of object programs, some compilation commands can not use at linking. For object programs combinations, and compilation commands/required options at linking, see "[Table 7.1 Object programs combinations, compilation commands, and required options](#)". For more information on the compile commands and options, see the user's guide for each language.

Table 7.1 Object programs combinations, compilation commands, and required options

Object to Link						COARRAY	Compile Commands/Required Options
Fortran	C trad	C clang	C++ trad	C++ clang			
				Use libc++ for C++ Standard Library	Use libstdc++ for C++ Standard Library		
*	*	-	-	-	-	Not Use	frtpx or fccpx --linkfortran
						Use	frtpx -Ncoarray or fccpx --linkcoarray
*	-	*	-	-	-	Not Use	frtpx or fccpx -Nclang --linkfortran
						Use	frtpx -Ncoarray or fccpx -Nclang --linkcoarray
*	-	-	*	-	-	Not Use	frtpx --linkstl=libfjcpp or FCCpx --linkfortran

Object to Link						COARRAY	Compile Commands/Required Options	
Fortran	C trad	C clang	C++ trad	C++ clang				
				Use libc++ for C++ Standard Library	Use libstdc++ for C++ Standard Library			
						Use	frtpx --linkstl=libfjc++ -Ncoarray or FCCpx --linkcoarray	
*	-	-	-	*	-	Not Use	frtpx --linkstl=libc++ or FCCpx -Nclang --linkfortran -stdlib=libc++	
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray	
				-	*	Not Use	frtpx --linkstl=libstdc++ or FCCpx -Nclang --linkfortran	
						Use	frtpx --linkstl=libstdc++ -Ncoarray or FCCpx -Nclang --linkcoarray	
-	*	*	-	-	-	-	fccpx -Nclang	
-	*	-	*	-	-	-	-	FCCpx
-	*	-	-	*	-	-	-	FCCpx -Nclang
				-	*			FCCpx -Nclang -stdlib=libstdc++
-	-	*	*	-	-	-	-	FCCpx
-	-	*	-	*	-	-	-	FCCpx -Nclang
				-	*			FCCpx -Nclang -stdlib=libstdc++
-	-	-	*	*	-	-	-	FCCpx -Nclang -stdlib=libc++ (Refer to "Note 1")
				-	*			Can not to be linked. (Refer to "Note 2")
*	*	*	-	-	-	Not Use	frtpx or fccpx -Nclang --linkfortran	
						Use	frtpx -Ncoarray or fccpx -Nclang --linkcoarray	
*	*	-	*	-	-	Not Use	frtpx --linkstl=libfjc++ or FCCpx --linkfortran	
						Use	frtpx --linkstl=libfjc++ -Ncoarray or FCCpx --linkcoarray	
*	*	-	-	*	-	Not Use	frtpx --linkstl=libc++ or	



Object to Link						COARRAY	Compile Commands/Required Options
Fortran	C trad	C clang	C++ trad	C++ clang			
				Use libc++ for C++ Standard Library	Use libstdc++ for C++ Standard Library		
							FCCpx -Nclang --linkfortran -stdlib=libc++
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray
				-	*	Not Use	frtpx --linkstl=libstdc++ or FCCpx -Nclang --linkfortran
						Use	frtpx --linkstl=libstdc++ -Ncoarray or FCCpx -Nclang --linkcoarray
*	-	*	*	-	-	Not Use	frtpx --linkstl=libfjc++ or FCCpx --linkfortran
						Use	frtpx --linkstl=libfjc++ -Ncoarray or FCCpx --linkcoarray
				*	-	Not Use	frtpx --linkstl=libc++ or FCCpx -Nclang --linkfortran -stdlib=libc++
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray
				-	*	Not Use	frtpx --linkstl=libstdc++ or FCCpx -Nclang --linkfortran
						Use	frtpx --linkstl=libstdc++ -Ncoarray or FCCpx -Nclang --linkcoarray
				*	-	Not Use	frtpx --linkstl=libc++ or FCCpx -Nclang --linkfortran -stdlib=libc++ (Refer to "Note 1")
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray (Refer to "Note 1")
				-	*	Not Use	Can not to be linked.
						Use	(Refer to "Note 2")

Object to Link						COARRAY	Compile Commands/Required Options
Fortran	C trad	C clang	C++ trad	C++ clang			
				Use libc++ for C++ Standard Library	Use libstdc++ for C++ Standard Library		
-	*	*	*	-	-	-	FCCpx
-	*	*	-	*	-	-	FCCpx -Nclang -stdlib=libc++
				-	*		FCCpx -Nclang -stdlib=libstdc++
-	*	-	*	*	-	-	FCCpx -Nclang -stdlib=libc++ (Refer to "Note 1")
				-	*		Can not to be linked. (Refer to "Note 2")
-	-	*	*	*	-	-	FCCpx -Nclang -stdlib=libc++ (Refer to "Note 1")
				-	*		Can not to be linked. (Refer to "Note 2")
*	*	*	*	-	-	Not Use	frtpx --linkstl=libfjc++ or FCCpx --linkfortran
						Use	frtpx --linkstl=libfjc++ -Ncoarray or FCCpx --linkcoarray
*	*	*	-	*	-	Not Use	frtpx --linkstl=libc++ or FCCpx -Nclang --linkfortran -stdlib=libc++
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray
				-	*	Not Use	frtpx --linkstl=libstdc++ or FCCpx -Nclang --linkfortran
						Use	frtpx --linkstl=libstdc++ -Ncoarray or FCCpx -Nclang --linkcoarray
*	*	-	*	*	-	Not Use	frtpx --linkstl=libc++ or FCCpx -Nclang --linkfortran -stdlib=libc++ (Refer to "Note 1")
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray (Refer to "Note 1")
				-	*	Not Use	Can not to be linked. (Refer to "Note 2")
						Use	Can not to be linked. (Refer to "Note 2")

Object to Link						COARRAY	Compile Commands/Required Options
Fortran	C trad	C clang	C++ trad	C++ clang			
				Use libc++ for C++ Standard Library	Use libstdc++ for C++ Standard Library		
*	-	*	*	*	-	Not Use	frtpx --linkstl=libc++ or FCCpx -Nclang --linkfortran -stdlib=libc++ (Refer to "Note 1")
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray (Refer to "Note 1")
						Not Use	Can not to be linked. (Refer to "Note 2")
-	*	*	*	-	*	Use	Can not to be linked. (Refer to "Note 2")
				*	-	-	FCCpx -Nclang -stdlib=libc++ (Refer to "Note 1")
*	*	*	*	*	-	Not Use	frtpx --linkstl=libc++ or FCCpx -Nclang --linkfortran -stdlib=libc++ (Refer to "Note 1")
						Use	frtpx --linkstl=libc++ -Ncoarray or FCCpx -Nclang -stdlib=libc++ --linkcoarray (Refer to "Note 1")
						Not Use	Can not to be linked. (Refer to "Note 2")
-	*	*	*	-	*	Use	Can not to be linked. (Refer to "Note 2")
				*	-	-	FCCpx -Nclang -stdlib=libc++ (Refer to "Note 1")

### 7.1.1 Linking C++ Trad Mode with C++ Clang Mode

Whether object programs generated in C++ Trad Mode and object programs generated in C++ Clang Mode can be linked depends on C++ Standard Library used at compilation.

C++ Trad Mode	C++ Clang Mode	Compile Commands/Required Options
(libc++ use only)	Use libc++ for C++ Standard Library	FCCpx -Nclang -stdlib=libc++ (Refer to "Note 1")
	Use libstdc++ for C++ Standard Library	Can not to be linked. (Refer to "Note 2")

#### Note 1

When all of the following conditions are met, the operation may become indefinite and can not be guaranteed.

1. An object generated in Clang Mode and an object generated in Trad Mode have an interface between functions.
2. An argument or return value of the function in condition 1 is a class type.

## Note 2

The default C++ Standard Library that C++ Clang Mode uses is `libstdc++`.

Specify `libc++` instead of the default `libstdc++` for C++ Standard Library used in Clang Mode.

## 7.1.2 Linking C++ Clang Mode with C++ Clang Mode

Whether object programs generated in C++ Trad Mode and object programs generated in C++ Clang Mode can be linked depends on C++ Standard Library used at compilation.

C++ Standard Library used in compilation		Compile Commands/Required Options
<code>libc++</code>	<code>libc++</code>	<code>FCCpx -Nclang -stdlib=libc++</code>
<code>libc++</code>	<code>libstdc++</code>	Can not to be linked.
<code>libstdc++</code>	<code>libstdc++</code>	<code>FCCpx -Nclang -stdlib=libstdc++</code> (Note that the <code>-stdlib = libstdc++</code> option can be omitted because it is the default.)

## 7.1.3 Linking MPI Object Programs and Use Profiler

Specify the following option at linkage:

- Use `mpifrtpx` command : `-lfjprofmpi` option
- Use `mpifccpx` command or `mpiFCCpx` command : `-lfjprofmpif` option

## 7.1.4 Objects Generated in Clang Mode with `-flto` Option

Objects generated with the `"-Nclang -flto"` options are in the format used internally by LLVM (different format from normal objects) and can only be linked in Clang Mode.

When including an object generated with the `"-Nclang -flto"` options at mixed language programming, specify the `"-Nclang -flto"` options at linking.

## 7.2 Linking with C

Note the following:

- To call a C++ function from a C function, the C++ function must be declared with linkage to C specified (via `extern "C"`) in the C++ program.
- To call a C function from a C++ function, the C function must be declared with linkage to C specified (via `extern "C"`) in the C++ program.
- The user can specify any type permitted under C for the arguments and return values of a function to be called in a C specified function specified with `extern "C"`.

Example 1: Passing control to the C++ program first

C++ source: `cplusmain.cc`

```
#include <iostream>

extern "C" {
    void cfunc(short int);
    int cpfunc(int);
}

int cpfunc(int i) {
    std::cout << "C++: cpfunc called, i=" << i << std::endl;
    return i;
}
```

```
int main() {
    std::cout << "C++ main()" << std::endl;
    cfunc(10);
    return 0;
}
```

C source: csub.c

```
#include <stdio.h>

int cfunc(int);

void cfunc(short int si) {
    printf("C:  cfunc called, si=%d\n", si);
    cfunc(si);
}
```

Compiling and linking:

```
$ fccpx -c csub.c
$ FCCpx csub.o cplusplusmain.cc
```

Execution:

```
$ ./a.out
```

Result:

```
C++ main()
C:  cfunc called, si=10
C++: cfunc called, i=10
```

Example 2: Passing control to the C program first

C++ source: cplusplus.cc

```
#include <iostream>

extern "C" {
    void cfunc(short int);
    int cfunc(int);
}

int cfunc(int i) {
    std::cout << "C++: cfunc called, i=" << i << std::endl;
    cfunc(i);
    return i;
}
```

C source: cmain.c

```
#include <stdio.h>

int cfunc(int);

int main() {
    printf("C main()\n");
    cfunc(10);
    return 0;
}

void cfunc(short int si) {
    printf("C:  cfunc called, si=%d\n", si);
}
```

Compiling and linking:

```
$ fccpx -c cmain.c
$ FCCpx cmain.o cplussub.cc
```

Execution:

```
$ ./a.out
```

Result:

```
C main()
C++: cpfunc called, i=10
C:   cfunc called, si=10
```

## 7.3 Linking with Fortran

---

Note the following:

- To call C++ function from Fortran, the function must be declared with linkage to C specified (`extern "C"`) in C++ program.
- To call a Fortran function from C++, the Fortran function must be declared with linkage to C specified in the C++ program.
- The user can specify any type permitted under C for the arguments and return values of a function to be called in a C specified function.
- If the program to which control is first passed is in C++, the function name must be `"MAIN__"` or `"main"`.
- In the following cases, the return value of Fortran is set the return value of executable program. For detail of the return value of Fortran, refer to "Fortran User's Guide".
  - The program to which control is first passed is Fortran program.
- Append `"_"` to the end of the C++ function name that is called from Fortran and then also at the end of the Fortran function name called from C++.
- Other rules are the same as those for interlanguage linkage between C++ and Fortran. Please refer to the "Fortran User's Guide".

When using the COARRAY specification, specify the `--linkcoarray` option instead of the `--linkfortran` option. For more information on the COARRAY specification, see "Fortran User's Guide Additional Volume COARRAY".

Example 1: Passing control to the C++ program first

C++ source: `cplusmain.cc`

```
#include <iostream>

extern "C" {
    int cpfunc_(int *);
    int func_(int *);
}

int cpfunc_(int *p) {
    std::cout << " C++: cpfunc_ called, *p=" << *p << std::endl;
    return *p;
}

int MAIN__() {
    std::cout << " C++ MAIN__()" << std::endl;
    int i=10;
    func_(&i);
    return 0;
}
```

Fortran source: `fortransub.f95`

```
integer function func(x)
integer*4 x
```

```
print *, "Fortran: func() called"
call cfunc(x)
func=x
end
```

Compiling and linking:

- When using the compile command for the Fortran language at linking

```
$ FCCpx -c cplusmain.cc
$ frtpx --linkstl=libfjc++ fortransub.f95 cplusmain.o
```

- When using the compile command for the C++ language at linking

```
$ frtpx -c fortransub.f95
$ FCCpx --linkfortran fortransub.o cplusmain.cc
```

Execution:

```
$ ./a.out
```

Result:

```
C++ MAIN__()
Fortran: func() called
C++: cfunc_ called. *p=10
```

Example 2: Passing control to the Fortran program

C++ source: cplussub.cc

```
#include <iostream>

extern "C" {
    int cfunc_(int *);
}

int cfunc_(int *p) {
    std::cout << " C++: cfunc_ called, *p=" << *p << std::endl;
    return *p;
}
```

Fortran source: fortranmain.f95

```
print *, "Fortran: main program"
call cfunc(10)
end
```

Compiling and linking:

- When using the compile command for the Fortran language at linking

```
$ FCCpx -c cplussub.cc
$ frtpx --linkstl=libfjc++ fortranmain.f95 cplussub.o
```

- When using the compile command for the C++ language at linking

```
$ frtpx -c fortranmain.f95
$ FCCpx --linkfortran fortranmain.o cplussub.cc
```

Execution:

```
$ ./a.out
```

Result:

```
Fortran: main program  
C++: cpfunc_ called, *p=10
```



# Chapter 8 Debugging Functions

A program does not work as expected, for example, it needs to check the cause and to correct.

This chapter describes the debugging functions provided in this system that help checking the cause.

## 8.1 Functions for Debugging

This section describes the check functions used to debug C++ programs.

The `-Nquickdbg` option is used for debugging during program execution. It provides the following checks:

- Subscript range checks (see Section "[8.1.1 Subscript Range Checks \(subchk Function\)](#)")
- Heap checks (see Section "[8.1.2 Heap Memory Checks \(heapchk Function\)](#)")

The `-Nquickdbg=inf_detail` and the `-Nquickdbg=inf_simple` options output the diagnostic messages when any error is detected. For the `subchk` function, these options control execution performance and contents of the diagnostic messages.

When an error is detected, if the `-Nquickdbg=inf_detail` option is enabled, a message to identify the cause, line number where the error occurred and the variable name that caused the error are displayed as diagnostic message and the execution is continued.

When an error is detected, if the `-Nquickdbg=inf_simple` option is enabled, a message indicating the error and the line number where the error occurred are displayed as diagnostic message and the execution is terminated. However, if the compiler option `-Nnoline` is enabled, the value of line number is not guaranteed.

By limiting the information included in diagnostic messages, the `-Nquickdbg=inf_simple` option has less impact on execution performance than the `-Nquickdbg=inf_detail` option.

For Heap Memory Checks, even if either of the `-Nquickdbg=inf_detail` or the `-Nquickdbg=inf_simple` option is specified, the message and the line number where the error occurred are displayed as diagnostic messages and the execution is continued. Even if which option is specified, it becomes an equal execution performance.

See Section "[2.2.2.7 -N Option](#)", for details of `-Nquickdbg`.

### 8.1.1 Subscript Range Checks (subchk Function)

If the `-Nquickdbg=subchk` option is set, the validity of the referenced subscript range is checked against the declared array size when arrays are referenced. If the subscript is not within the range that applied when the array was declared, a diagnostic message is output. When `-Nquickdbg=inf_detail` is set, the diagnostic message is `jwe1601i-w`. When `-Nquickdbg=inf_simple` is set, the diagnostic message is `jwe1606i-u`.

This check is not performed in the following cases:

- An array specified as a parameter of function.
- An array which size is not determined at compilation (Variable-Length Array, Flexible Array, Arrays of Length Zero)
- An expression used as index which has apprehension of occurrence of side effect to refer array.
- Statements and declarations in an expression <sup>[a]</sup> used as index which has apprehension of occurrence of side effect to refer array.

[a] "Statements and Declarations in Expressions" is a part of GNU C++ Extensions.

### 8.1.2 Heap Memory Checks (heapchk Function)

If the `-Nquickdbg=heapchk` option is set, the memory allocated by the `malloc`, `calloc`, `realloc`, `valloc`, `posix_memalign`, and `memalign` functions is checked as follows.

- The memory release check and the buffer overrun check are performed when heap memory is released.
- The memory leak check is performed when the program ends.

The check items are explained below.



If memory allocated using functions other than malloc, calloc, realloc, valloc, posix\_memalign, and memalign is released, diagnostic message jwe1602i-w or jwe1603i-w may be output.

### 8.1.2.1 Memory Release Check

When heap memory is released, the following checks are performed for the memory being released:

- Release of memory that has already been released
- Invalid memory release

If a duplicate release of memory is detected, a diagnostic message (jwe1602i-w) is output. If an invalid memory release is detected, a diagnostic message (jwe1603i-w) is output.

### 8.1.2.2 Buffer Overrun Check

When heap memory is released, the function checks whether or not memory outside the allocated memory block is being written to.

The buffer overrun check performs a write check for the followed 8 bytes of heap memory. If buffer overrun is detected, a diagnostic message (jwe1604i-w) is output.

However, if the followed 8 byte areas of heap memory have the following value, an invalid area write is not detected even if buffer overrun has occurred:

- 0x8b8b8b8b8b8b8b8b

The above value can be changed by specifying the FLIB\_HEAPCHK\_VALUE environment variable.

```
FLIB_HEAPCHK_VALUE hex
```

For buffer overrun checks, the values used for the check are changed. By specifying hex, hexadecimal values can be specified. The valid range of values is 00<=hex<=FF. If the environment variable is not specified or if hex values are disabled, the system default value takes effect. The default value is 8B.

An example of specifying the FLIB\_HEAPCHK\_VALUE environment variable is shown below.

```
$ FLIB_HEAPCHK_VALUE=8C
$ export FLIB_HEAPCHK_VALUE
```

In this example, 0x8c8c8c8c8c8c8c8c is applied as the check value used in the outside area write check.

### 8.1.2.3 Memory Leak Check

When a program ends, the function checks whether or not heap memory remains without being released. If unreleased heap memory remains, a diagnostic message (jwe1605i-w) is output.

## 8.2 Debugging Programs for Abend

When the -NRtrap option is specified, the signal is caught if an abnormal termination event (hereafter called an abend) occurs during execution of a program. And then the information is output to help the user determine the cause of the abend.

When the -NRnotrap option is enabled, the information is not output.

For details about the -NRtrap and -NRnotrap option, see Section "2.2.2.7 -N Option".

### 8.2.1 Causes of Abend

"Table 8.1 caught signals and corresponding signal codes" lists the caught signals (\*) and the corresponding signal codes.

\*: Only signals supported by the Fujitsu CPU A64FX with Arm architecture are targeted. Therefore, enabling the -NRtrap option will not detect unsupported signals. See Section "2.2.2.7 -N Option" for more information about the -NRtrap option.

If the -NRtrap option is not specified, these signals are not caught.

Table 8.1 caught signals and corresponding signal codes

Signal number	Meaning of signal number	Signal code	Meaning of signal code	
SIGILL(04)	Incorrect instruction executed	1	ILL_ILLOPC	Illegal opcode
		2	ILL_ILLOPN	Illegal operand
		3	ILL_ILLADR	Illegal addressing mode
		4	ILL_ILLTRP	Illegal trap
		5	ILL_PRVOPC	Illegal privileged opcode
		6	ILL_PRVREG	Illegal privileged register
		7	ILL_COPROC	Coprocessor error
		8	ILL_BADSTK	Internal stack error
SIGFPE(08)	Floating-point exception <sup>[a]</sup>	3	FPE_FLTDIV	Floating-point divide by zero
		4	FPE_FLTOVF	Floating-point overflow
		5	FPE_FLTUND	Floating-point underflow <sup>[b]</sup>
		7	FPE_FLTINV	Invalid floating-point operation
		-	-	Floating-point exception
SIGBUS(10)	Storage protection exception	1	BUS_ADRALN	Invalid address alignment
		2	BUS_ADRERR	Non-existent physical address
		3	BUS_OBJERR	Object-specific hardware error
SIGSEGV(11)	Segmentation exception	1	SEGV_MAPERR	Address not mapped to object
		2	SEGV_ACCERR	Invalid permissions for mapped object
SIGXCPU(30)	CPU time interrupt	-	-	-

[a]: Note that FPE\_INTDIV (the integer divide by zero) is not detected in the Fujitsu CPU A64FX with Arm architecture. See Section "B.2.10 Integer Division Exception when the Divisor Is Zero" for details.

[b] This signal is caught when the -NRtrap option and environment variable FLIB\_EXCEPT=u is specified.

## 8.2.2 Information Generated when an Abend Occurs

Depending on the signal, the information is written to the standard error as explained in the following sections.

### 8.2.2.1 Information Generated for a General Abend

The information generated for a general abend is follows. This information is generated when SIGILL, SIGBUS, or SIGSEGV is caught. A trace back map is printed after this information.

```
jwe0019i-u The program was terminated abnormally with signal number SSSSSS.
      signal identifier = NNNNNNNNNN, (Detailed information.)
```

SSSSSS	SIGILL, SIGBUS, or SIGSEGV
NNNNNNNNNN	Signal code for the abend cause (hexadecimal)
(Detailed information.)	Detailed information of Signal code NNNNNNNNNN

When it is strong possibility that the cause of SIGBUS or SIGSEGV is stack-overflow, add the following information.

```
The cause of this exception may be stack-overflow.
```

## 8.2.2.2 Information of SIGXCPU

The output when SIGXCPU is detected is follows.

```
jwe0017i-u The program was terminated with signal number SIGXCPU.
```

## 8.2.2.3 Information when an Abend Occurs Again during Abend Processing

The output when an abend occurs again during abend processing is follows.

```
jwe0020i-u An error was detected during an abnormal termination process.
```

# 8.3 Hook Function

---

This section describes the hook function which calls the user-defined function by specified location in a C++ program, or by regular time interval.

A user program can be monitored via the user-defined function which can output trace information and specific variables values.

## 8.3.1 User-Defined Function Format

---

The user-defined function name is fixed (`user_defined_proc`).

Define the user-defined function in the format shown below.

- format:

```
#include "fjhook.h"
void user_defined_proc(int *FLAG, char *NAME, int *LINE, int *THREAD)
```

- Arguments:

- FLAG: Indicate the calling source for the user-defined function.

= 0: Program entry

= 1: Program exit

= 2: function entry

= 3: function exit

= 4: Parallel region (OpenMP or automatic parallelization) entry

= 5: Parallel region (OpenMP or automatic parallelization) exit

= 6: Regular time interval

= 7 to 99: Reserved for system

>= 100: Can be used by the user

- NAME: Indicate the calling source function name.

NAME can be referenced only if the FLAG is 2, 3, 4, 5, or 100 or greater.

- LINE: Indicate the calling source line number.

LINE can be referenced only if the FLAG is 2, 3, 4, 5, or 100 or greater.

- THREAD: With thread parallelization using OpenMP or automatic parallelization, indicate the number of the thread that calls the user-defined function.

THREAD can be referenced only if the FLAG is 2, 3, 4, 5, or 100 or greater.

## 8.3.2 Notes on Hook Function

---

The following notes apply when the hook function is used:

- The function name "user\_defined\_proc" must not be used for any other purpose.
- Operation is not guaranteed if the user-defined function is not defined.
- Values must not be set in the user-defined function arguments.
- If the -Nnoline compile option is set, the LINE argument value is not guaranteed.
- In thread parallelized programs, the user-defined function may be called from multiple threads.
- If the setjmp or longjmp function is called, or if an exception is caught, the user-defined function may not be called.
- The entry and exit of user-defined function except in the case of calling from any location in the program must not be calling source for the user-defined function.
- Operation is not guaranteed if the exit ( 3 ) function is executed directly or indirectly in the user-defined function.
- The OpenMP specifications which are available in the user-defined function are the following OpenMP functions.
  - omp\_get\_active\_level
  - omp\_get\_ancestor\_thread\_num
  - omp\_get\_dynamic
  - omp\_get\_level
  - omp\_get\_max\_active\_levels
  - omp\_get\_max\_threads
  - omp\_get\_nested
  - omp\_get\_num\_procs
  - omp\_get\_num\_threads
  - omp\_get\_proc\_bind
  - omp\_get\_schedule
  - omp\_get\_thread\_limit
  - omp\_get\_thread\_num
  - omp\_get\_wtick
  - omp\_get\_wtime
  - omp\_in\_final
  - omp\_in\_parallel
  - omp\_set\_dynamic
  - omp\_set\_max\_active\_levels
  - omp\_set\_nested
  - omp\_set\_num\_threads
  - omp\_set\_schedule

## 8.3.3 Calling the User-Defined Function from a Specified Location

---

If the -Nhook\_func option is set when the object program and executable program are created, the user-defined function can be called from the following locations:

- Program entry and exit
- Function entry and exit

If the `-Kopenmp` or `-Kparallel` option is enabled, user-defined function can also be called from the following locations:

- Parallel region (OpenMP or automatic parallelization) entry and exit

When the `-Nhook_func` is set, operation is as follows for Fortran, and C++ linked programs:

- If user-defined subroutine/function is defined within the Fortran program and the C++ program respectively, the user-defined subroutine defined in the Fortran program is called from the Fortran program, and the user-defined function defined in the C++ program is called from the C++ program.
- If user-defined subroutine/function is defined within only the Fortran program or only the C++ program, the defined user-defined subroutine/function is called from both the Fortran program and the C++ program.

An example using the hook function with location specified is shown below.

User-defined function example program:

```
#include <stdio>
#include "fjhook.h"

void sub() {
    printf("SUB\n");
}

int main() {
    printf("HELLO\n");
    sub();
}

void user_defined_proc(int *FLAG, char *NAME, int *LINE, int *THREAD) {
    switch(*FLAG){
        case 0:
            printf("PROGRAM START\n");
            break ;
        case 1:
            printf("PROGRAM END\n");
            break ;
        case 2:
            printf("PROC START: %s LINE: %d\n",NAME,*LINE);
            break ;
        case 3:
            printf("PROC END: %s LINE: %d\n",NAME,*LINE);
            break ;
    }
}
```

Executable program creation

```
$ FCCpx -Nhook_func test.cc
```

Execution results

```
PROGRAM START
PROC START: main LINE: 8
HELLO
PROC START: sub LINE: 4
SUB
PROC END: sub LINE: 6
PROC END: main LINE: 11
PROGRAM END
```

### 8.3.3.1 Notes on Calling from a Specified Location

The following notes apply when the `-Nhook_func` option is set:

- If a function with few computations is called repeatedly, calling the user-defined function from a function entry/exit may affect execution performance.
- If a parallel region with few computations is executed repeatedly, calling the user-defined function from a parallel region entry/exit may affect execution performance.

## 8.3.4 Calling the User-Defined Function at Regular Time Interval

---

If the `-Nhook_time` option is set at linking, the user-defined function is called by regular user time interval.

The calling interval is specified using the environment variable `FLIB_HOOK_TIME`. If it is not specified, the user-defined function is called once every minute.

See Section "[2.5.1 Environment Variable for Execution](#)" for information about the environment variable `FLIB_HOOK_TIME`.

### 8.3.4.1 Notes on Calling by Regular Time Interval

Note the following when the `-Nhook_time` option is set:

- The environment variable `FLIB_HOOK_TIME` must not be changed while program execution is in progress.
- Operation is not guaranteed if a value outside the range 0 to 2147483647 is specified in the environment variable `FLIB_HOOK_TIME`.
- Calling the user-defined function may affect execution performance.

Moreover, the following restrictions are applied because calling by regular interval uses asynchronization signal to call the user-defined function.

- No functions excluding `async-signal-safe` functions can be used in the user-defined function.
- No signal handler which handles `SIGVTALRM` must be defined if the `-Nhook_time` option is set.
- When the `-Nhook_time` option is set, `setitimer(2)` must not be used.
- Only volatile `sig_atomic_t` type variables are guaranteed for global variable references by the user-defined function.
- The profiler must not be used for the executable program created with the `-Nhook_time` option.

## 8.3.5 Calling from Any Location in the Program

---

The user-defined function can be called from any location in a program by specifying arguments (`FLAG:>=100, NAME, LINE, THREAD`).

# Chapter 9 Clang Mode

This chapter describes Clang Mode and the difference with Trad Mode.

## 9.1 From Compilation to Execution

### 9.1.1 Compile Command

This section describes the compile command.

#### 9.1.1.1 Syntax of the Compile Command

Clang Mode is available when `-Nclang` option is set.

The format of the compile command is as follows:

Table 9.1 Format of the Compile Command

Command Name		Operands
Cross compiler	FCCpx	<code>[_option-list-including-Nclang]_filename-list</code>
Native compiler	FCC	<code>[_option-list-including-Nclang]_filename-list</code>

\_: At least one blank is required.

#### Information

There is no restriction on the order in which `option-list-including-Nclang` and `filename-list` are specified.

For example, `option-list-including-Nclang` may be specified together.

#### 9.1.1.2 Input Files for the Compile Command

The file types that can be specified as input files for the compile command are as follows:

Table 9.2 Format of Input Files

File Type	File Suffix	Passed To
Headers	.h	Preprocessor <sup>[a]</sup>
	.H	
C++ source file	.cc	Preprocessor and Compiler
	.CC	
	.cpp	
	.CPP	
	.cxx	
	.CXX	
	.c++	
	.C++	
.C		



File Type	File Suffix	Passed To
	.cp	
	.ii	
Assembler source file	.s	Assembler
Assembler source file that must be preprocessed	.S	Preprocessor and Assembler
Object file	.o	Linker

[a] When -E option is effective.

### 9.1.1.3 Return Value of the Compile Command

The possible return values of the compile command are as follows:

Table 9.3 Return Values of the Compile Command

Return Value	Meaning
0	Terminated normally.
1	A compiler error or linker error has occurred.

### 9.1.1.4 Notes on Compilation

This section provides notes on compilation.

#### 9.1.1.4.1 Stack Size on Compilation

When compiling programs with remarkably many operations in loops, the segmentation exception may occur due to insufficient stack area in the compiler.

In this case, the stack area of the compilation environment needs to be extended to appropriate size.

The stack area of the process can be set by `ulimit` (bash built-in command), etc.

## 9.1.2 Compiler Options

This section describes the format of compiler options and their meanings.

The compiler options can be specified as follows:

- Option arguments of the compile command
- Environment variable (See Section "9.1.3 Environment Variable for Compile Command")
- Compilation profile file (See Section "9.1.4 Compilation Profile File")

See Section "2.4 Compilation Profile File" for the priority of the specified compiler options.



#### Note

- Almost all Clang/LLVM options are available. However, if you specify an option not described in this chapter, the result is not guaranteed. The result is not guaranteed when using the options not listed in this chapter.
- Trad Mode compilation options not listed in "Table 9.6 Options replaced in Clang Mode" cannot be specified in Clang Mode. If specified, it outputs an error message and aborts the compilation.

### 9.1.2.1 Syntax of the Compile Command

Compiler options can be specified as operands of the compile command. The format is given below.

## General compiler options

```
[ -### ] [ -C ] [ -Dname[=tokens] ] [ -E ] [ -I $dir$  ] [ -L $dir$  ] [ -M ] [ -MD ] [ -MF  $filename$  ] [ -MM ] [ -MMD ] [ -MP ] [ -MT  $target$  ] [ -Nopt ] [ -P ] [ -S ] [ -SSL2 ] [ -SSL2BLAMP ] [ -Uname ] [ -W $tool, arg1[, arg2]...$  ] [ -c ] [ -fopt ] [ {-g|-g0} ] [ -lname ] [ -o  $pathname$  ] [ -shared ] [ -v ] [ --coverage ] [ --verbose ]
```

## Options for messages

```
[ -Rpass=.* ] [ -Rpass-analysis=.* ] [ -Rpass-missed=.* ] [ -w ] [ -Koptmsg=2 ]
```

## Options for optimization

```
[ -O $n$ ] [ -m{omit-leaf-frame-pointer|no-omit-leaf-frame-pointer} ] [ -msve-vector-bits={512|scalable} ] [ -fopt ] [ -Kopt ]
```

## Options for language specifications

```
[ -std= $name$  ] [ -stdlib= $name$  ] [ --linkcoarray ] [ --linkfortran ]
```

## Options for CPU/architecture

```
[ -march= $arch[+features]...$  ] [ -Karch ] [ -mcpu= $cpuname[+features]...$  ] [ -Kcpuname ]
```

## Options for code generation

```
[ -fopt ] [ -mcmmodel= $name$  ] [ -mfj-tls-size={12|24|32|48} ] [ -Kopt ]
```

## 9.1.2.2 Description of Compiler Options

A description of the compiler options is given below.

### 9.1.2.2.1 General Options for Compiler

The following options are general options for the compiler.

**-###**

Shows commands passed by the compile command, but does not execute.

**-C**

This option executes the preprocessing phase to retain all comments other than those on preprocessing directive lines.

**-Dname[=tokens]**

Assigns *name* to the specified *tokens* as in the `#define` preprocessor directive.

**-E**

Performs only preprocessing on the specified C++ source code and sends the results to the standard output.

**-Nopt**

When *-Nopt* option is specified, it is replaced with the corresponding *-fopt* option.

For more detail, see "[Table 9.6 Options replaced in Clang Mode](#)".

**-I $dir$**

Sets the search path for headers with names that do not start with `'/'` (i.e. headers specified with an absolute path name). The directory specified in *dir* is searched first and then the directory-installed standard headers are searched.

If multiple directories are specified in multiple *-I* options, the directories are searched in the specified order.

**-L $dir$**

Adds *dir* to the list of directories through which the linker searches for libraries.

This option is passed to the linker.

**-M**

Specifies to output the dependency of the sources in the format recognized by the `make` command.

When this option is valid, only the preprocessing is executed.

**-MD**

Specifies to output the dependency of the sources in the format recognized by the `make` command.

By default, the result is output to the file with a ".d" suffix.

When `-E` and `-o` options are specified, the result is output to the file specified by `-o` option.

When `-MF` option is specified simultaneously, the result is output to the file specified by `-MF` option.

**-MF *filename***

Specifies the *filename* to output the result of the dependency of the source using `-M`, `-MM`, `-MD` or `-MMD` options.

**-MM**

Specifies to output the dependency of the sources in the format recognized by the `make` command. However, a standard header is not contained in the output.

When this option is valid, only the preprocessing is executed.

**-MMD**

Specifies to output the dependency of the sources in the format to recognized by the `make` command. However, a standard header is not contained in the output.

By default, the result is output to the file with a ".d" suffix.

When `-E` and `-o` options are specified, the result is output to the file specified by `-o` option.

When `-MF` option is specified simultaneously, the result is output to the file specified by `-MF` option.

**-MP**

Specifies to append the phony target of the dependency to the output of the result produced by `-M`, `-MM`, `-MD` or `-MMD` option.

**-MT *target***

Specifies to change the *target* of the dependency in the output of the result produced by `-M`, `-MM`, `-MD` or `-MMD` option.

**-P**

Specifies not to include linemarkers in the output of `-E` option.

**-S**

Compiles the specified C++ source file and leaves the assembly-language output in a corresponding file with the suffix ".s".

When this option and `-flto` option are specified simultaneously, the generated file includes the internal information of the compiler. Therefore, the file cannot be passed to Trad Mode or the assembler (`as` command).

When this option and either of `-g` option or `-ffj-line` option are enabled, the generated file includes information extended for Clang Mode. Therefore, the file cannot be passed to Trad Mode or the assembler (`as` command).

**-SSL2**

Specifies to link C-SSL II, C-SSL II thread-parallel capabilities, and BLAS/LAPACK.

**-SSL2BLAMP**

Specifies to link C-SSL II, C-SSL II thread-parallel capabilities, and BLAS/LAPACK thread-parallel versions.

This option just replaces the sequential BLAS/LAPACK from `-SSL2` with the corresponding thread-parallel versions.

**-U*name***

Invalidates the definition of *name* as specified in `#undef` preprocessor directive.

**-W*tool, arg1[, arg2]...***

Specified *arg1[, arg2]...* are passed as separate arguments to *tool*. An individual argument must be separated by a comma only.

*tool* can be one of the following characters:

Character specified as <i>tool</i>	Where <i>arg1[,arg2]...</i> are passed
p	Preprocessor
a	Assembler
l	Linker

Note that `-flt` option is invalidated when the `-Wa` and `-flt` options are specified simultaneously.



## Example

To indicate to the linker that delayed link should not be used for symbol resolution of shared library addresses, specify as follows:

```
-Wl,-z,now
```

**-c**

Suppresses the linking phase.

The files generated when this option and `-flt` option are specified simultaneously can be linked only when `-Nclang` option is set.

**-fopt**

For *opt*, specify one of the following:

```
{ {debug-info-for-profiling|no-debug-info-for-profiling} | {ffj-fjcex|ffj-no-fjcex} | {ffj-fjprof|ffj-no-fjprof} | {ffj-hook-time|ffj-no-hook-time} | {ffj-line|ffj-no-line} | ffj-lst[={p|t}] | ffj-lst-out=file | ffj-src | {cxx-exceptions|no-cxx-exceptions} | profile-dir=dir_name | sanitize={address|undefined} | {unwind-tables|no-unwind-tables} }
```

**-f{debug-info-for-profiling|no-debug-info-for-profiling}**

The compiler option `-fdebug-info-for-profiling` generates signatures for additional information available to the Profiler. The signature is the information used by overload resolution. This additional information is required to use the `-Minlined` option of the Instant Performance Profiler.

If the compiler option `-fdebug-info-for-profiling` is specified, the signature is generated as additional information. In this case, the Instant Performance Profiler displays the function cost per signature for inlined functions when the `-Minlined` option is effective.

If the compiler option `-fno-debug-info-for-profiling` is specified, the signature is not generated as additional information. In this case, the Instant Performance Profiler displays the function cost per function name for inlined functions when the `-Minlined` option is effective. Also, because no signatures are generated, the size of the object program is smaller than when the compiler option `-fdebug-info-for-profiling` is specified.

The compiler option `-fdebug-info-for-profiling` is set by default. The compiler options `-fdebug-info-for-profiling` and `-fno-debug-info-for-profiling` require that the compiler option `-ffj-line` is set.

See the "Profiler User's Guide" about the Profiler.

**-f{ffj-fjcex|ffj-no-fjcex}**

`-ffj-fjcex` specifies to use the Fujitsu Extended Functions in Clang Mode. `-ffj-no-fjcex` option invalidates `-ffj-fjcex` option. `-ffj-no-fjcex` is set by default.

`-ffj-fjcex` option must be set at both compilation and linking.

For details about the extended functions, see "[Appendix F Fujitsu Extended Functions](#)".

**-f{ffj-fjprof|ffj-no-fjprof}**

`-ffj-fjprof` option specifies to enable the Profiler function. The `-ffj-no-fjprof` option invalidates the `-ffj-fjprof` option. These options must be set at linking. The default is `-ffj-fjprof`.

See the "Profiler User's Guide" for information on the Profiler.

### -f{fj-hook-time|fj-no-hook-time}

-ffj-hook-time specifies to use a hook function that is called at regular time interval. -ffj-no-hook-time option invalidates -fj-hook-time option. -ffj-no-hook-time is set by default.

If -ffj-hook-time option is set, the user-defined function is called at regular time interval.

The interval time to call the user-defined function is specified by environment variable FLIB\_HOOK\_TIME. If environment variable FLIB\_HOOK\_TIME is not specified, the user-defined function is called at the interval of a minute.

For details about the environment variable FLIB\_HOOK\_TIME, see Section "9.1.6.1 Environment Variable for Execution".

This option must be set at linking.

See Section "8.3.4 Calling the User-Defined Function at Regular Time Interval", for information about the hook function.

### -f{fj-line|fj-no-line}

-ffj-line generates additional information required to use the execution time sampling function provided with the Profiler. -ffj-no-line option invalidates -ffj-line option. -ffj-line is set by default.

If the -flto option is specified with -ffj-line option, information about loops required by the Profiler's execution time sampling feature will not be generated. See the "Profiler User's Guide" about the Profiler.

### -ffj-lst[={p|t}]

Specifies to output compilation information to file(s) whose name is with suffix ".lst". When more than one C++ source files are specified, the file names are "*each-filename*.lst".

When the -ffj-lst option is specified without arguments, -ffj-lst=p is set.

#### -ffj-lst=p

Specifies to write compilation information that consists of source list.

The source list includes annotation of optimization which shows the situation of inline expansion and loop unrolling, etc.

#### -ffj-lst=t

Specifies to write detail optimization information to source list in addition to the function of -ffj-lst=p.



If the "*each-filename*.opt.yaml" file already exists in the directory where the compilation was performed, the information of compilation is not output.

### -ffj-lst-out=*file*

Specifies the filename to output compilation information.

When this option is specified, -ffj-lst=p option is also effective.

### -ffj-src

Specifies to output source list to the standard output.

The source list includes annotation of optimization which shows the situation of automatic parallelization, inline expansion, loop unrolling, etc.

Note that when -ffj-lst or -ffj-lst-out=*file* option is specified, the source list is output to the *file*.

### -f{cxx-exceptions|no-cxx-exceptions}

-fcxx-exceptions defines the \_\_EXCEPTIONS macro. -fcxx-exceptions is set by default.

### -fprofile-dir=*dir\_name*

This option specifies storage location directory of the .gcdc file which is necessary for the use of the code coverage. For *dir\_name*, the storage location directory name is specified by the relative path or the absolute path.

The .gcdc file is generated when you run an executable program that link an object file with information for code coverage. At runtime, the directory specified by *dir\_name* is created if it does not already exist.

This option is effective when the `--coverage` and, `-S` or `-c` options are enabled.

For information about the code coverage feature and the default storage location directory, see "[9.9 Code Coverage](#)".

`-fsanitize={address|undefined}`

`-fsanitize` option is used to debug C++ programs. The debugging function includes embedded information for debugging to a C++ object file, and checks are performed automatically during execution. Checks are performed if `-fsanitize=undefined` or `-fsanitize=address` option is set.

Multiple debugging features can be combined by specifying multiple `-fsanitize` options.

This option must be set when compiling and linking programs.

This option invalidates `-ffj-largepage` option.

`-fsanitize=address`

Specifies whether to check the heap for improper use of memory, buffer overrun, and memory leaks.

`-fsanitize=undefined`

Specifies whether to check the validity of the referenced subscript range against the declared array size when arrays are referenced.

`-f{unwind-tables|no-unwind-tables}`

`-funwind-tables` specifies to generate unwind tables for storing information necessary for debugging, such as stack trace information.

`-fno-unwind-tables` option invalidates `-funwind-tables` option.

`-funwind-tables` is set by default.

`{-g|-g0}`

`-g` option generates additional information used by the debugger.

`-g0` option invalidates `-g` option.

`-g0` is set by default.

`-lname`

Specifies to link a library named `libname.so` or `libname.a`.

The position of this option in the command line is important, since libraries are searched for in the order in which the other libraries and object files appear in the command line. This option must be specified after source files.

This option is passed to the linker.

`-o pathname`

Creates a file of the name specified in *pathname*.

`-shared`

Specifies to create a shared object. This option is passed to the linker.

`-v`

This option specifies to display the executing command which compile command called as compiler, assembler and linker.

`--coverage`

The `--coverage` option specifies to generate the information for the code coverage. `--coverage` option should be specified when programs are compiled and linked.

When the `--coverage` option is effective, the `-flto` option is invalid.

If the `--coverage` option is specified, the execution performance may decrease due to the instructions that measure the execution time are added in the object program.

For details about the code coverage, see section "[9.9 Code Coverage](#)".

`--verbose`

This option specifies to display the executing command which compile command called as compiler, assembler and linker.

## 9.1.2.2.2 Options for Messages

The options for messages are given below.

**-Rpass=.\***

**-Rpass=.\*** specifies to output the messages of applied optimizations.

### Example

Usage of **-Rpass=.\***

```
$ FCCpx -Nclang test.cpp -Ofast -Rpass=.* -S
test.cpp:4:6: remark: hoisting load [-Rpass=licm]
test.cpp:12:4: remark: peeled loop by 1 iterations [-Rpass=loop-unroll]
    while( j < mat.operator[](i).size() ){
    ^
test.cpp:4:6: remark: load of type double* eliminated [-Rpass=gvn]
void gemv(vector<double>          & ret,
```

**-Rpass-analysis=.\***

**-Rpass-analysis=.\*** specifies to output the reason why the compiler did not apply optimization.

### Example

Usage of **-Rpass-analysis=.\***

```
$ FCCpx -Nclang test.cpp -Ofast -Rpass-analysis=.* -S
b.cpp:13:23: remark: _ZNKSt6vectorIdSaIdEEixEm can be inlined into
    _Z4gemvRSt6vectorIdSaIdEERKS_IS1_SaIS1_EERKS1_ with cost=-35
    (threshold=375) [-Rpass-analysis=inline]
    double const v = vec[i];
    ^
b.cpp:15:13: remark: loop not vectorized: value that could not be identified as
    reduction is used outside the loop [-Rpass-analysis=loop-vectorize]
    ret[i] += m * v;
```

**-Rpass-missed=.\***

**-Rpass-missed=.\*** specifies to output that the optimization was not applied. However, the reason why the compiler did not apply optimization is not output.

### Example

Usage of **-Rpass-missed=.\***

```
$ FCCpx -Nclang test.cpp -Ofast -Rpass-missed=.* -S
b.cpp:13:23: remark: failed to move load with loop-invariant address because the
    loop may invalidate its value [-Rpass-missed=licm]
    double const v = vec[i];
    ^
b.cpp:12:4: remark: loop not vectorized [-Rpass-missed=loop-vectorize]
    while( j < mat.operator[](i).size() ){
```

**-w**

Suppresses the output of warning messages.

**-Koptmsg=2**

When **-Koptmsg=2** option is specified, it is replaced with **-Rpass=.\*** option.

### 9.1.2.2.3 Options for Optimization

The optimization options are described below.

#### `-O[n]`

In *n*, specify the level of optimization as 0, 1, 2, 3 or `fast`. If `-O` option is specified without *n*, `-O2` is set. As the optimization level increases, the execution time decreases and the compilation time increases. The higher levels of optimization functionally include the lower levels of optimization.

If `-O[n]` option is not specified, `-O0` is set.

- Optimization level 0

No optimization is performed.

- Optimization level 1

Basic optimization is performed. This may reduce the object size and the execution time compared with optimization level 0.

- Optimization level 2

In addition to the optimization of level 1, the following optimizations are performed:

- Create an object using SIMD extension instructions (`-fvectorize`)
- Loop unrolling (`-funroll-loops`)
- Loop fission (`-ffj-loop-fission`)
- Using prefetch instructions (`-ffj-prefetch-sequential`)
- Inline expansion (`-finline-functions`)
- Superword Level Parallelism(`-fslp-vectorize`)

- Repeated application of optimization functions

Repeated application of optimization functions means that the optimization functions performed in optimization level 1 are repeatedly performed until there is no room for further optimization.

- Optimization level 3

In addition to the optimization of level 2, the more highly optimization is performed.

The compilation time and the size of the object program may increase more than the optimization level 2.

- Optimization level `fast`

Specifies the optimizations for high speed execution on the target machine.

This option is equivalent to the following:

`-O3 -ffj-fast-matmul -ffast-math -ffp-contract=fast -ffj-fp-relaxed -ffj-ilfunc -fomit-frame-pointer -finline-functions`

When this option is set, side effects may occur in the execution results due to effect of `-ffj-fast-matmul`, `-ffast-math`, `-ffp-contract=fast`, `-ffj-fp-relaxed`, and `-ffj-ilfunc` options. See "[9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

#### `-m{omit-leaf-frame-pointer|no-omit-leaf-frame-pointer}`

Indicates whether or not to perform the optimization that the frame pointer register is not kept for leaf functions (functions that do not call any other functions).

When `-momit-leaf-frame-pointer` option is specified, this optimization is performed. Trace back information is not kept.

When `-mno-omit-leaf-frame-pointer` option is specified, this optimization is not performed and the execution performance may decrease by saving and restoring instructions for keeping frame pointer register in leaf functions.

`-momit-leaf-frame-pointer` option is set by default.

The following is the relationship to the `-f{omit-frame-pointer|no-omit-frame-pointer}` options.

- `-momit-leaf-frame-pointer` option is induced by `-fomit-frame-pointer` option.



- `-mno-omit-leaf-frame-pointer` option is induced and enabled only when `-fno-omit-leaf-frame-pointer` option is specified explicitly. Specifying only `-mno-omit-leaf-frame-pointer` option has no effect.

#### `-msve-vector-bits={512|scalable}`

Specifies the size of the SVE vector register. Units are bits.

When the `-msve-vector-bits=512` option is specified, optimizations are performed on the assumption that the vector register size is a fixed value specified as the option at compilation time. Therefore, optimizations are promoted and the improvement of the execution performance is expected.

However, the generated executable program works normally only on CPU architecture which has the same size of the SVE vector register as the size specified at compilation time. For details, see Section "[9.1.2.3.5 Notes on Specified SVE Vector Register Size](#)".

When the `-msve-vector-bits=scalable` option is specified, the SVE vector register is not considered to be a specific size, and the executable program decides the vector register size at execution time. The executable program does not depend on the SVE vector register size on the CPU architecture.

`-msve-vector-bits=scalable` is set by default.

This option is effective when `sve` is enabled at *features* of `-march` option.

Note that `-fslp-vectorize` option is invalidated when the `-msve-vector-bits=512` options is specified. And, the SIMD built-in functions cannot be used when the `-msve-vector-bits=512` option is specified.

If both the `-msve-vector-bits=512` and `-flto` options are specified, the `-msve-vector-bits=512` option is disabled and the `-msve-vector-bits=scalable` option is enabled.

#### `-fopt`

For *opt*, specify one of the following:

```
{ {fj-eval-concurrent|fj-no-eval-concurrent} | {fj-fast-matmul|fj-no-fast-matmul} | {fj-fp-precision|fj-no-fp-precision} | {fj-fp-relaxed|fj-no-fp-relaxed} | {fj-hpctag|fj-no-hpctag} | {fj-ilfunc[={loop|procedure}]|fj-no-ilfunc} | {fj-interleave-loop-insns[=N]|fj-no-interleave-loop-insns} | {fj-loop-fission|fj-no-loop-fission} | fj-loop-fission-threshold=N | fj-no-prefetch | {fj-ocl|fj-no-ocl} | {fj-optlib-string|fj-no-optlib-string} | {fj-preex|fj-no-preex} | fj-prefetch-cache-level={1|2|all} | {fj-prefetch-conditional|fj-no-prefetch-conditional} | fj-prefetch-iteration=N | fj-prefetch-iteration-L2=N | fj-prefetch-line=N | fj-prefetch-line-L2=N | {fj-prefetch-sequential[={auto|soft}]|fj-no-prefetch-sequential} | {fj-prefetch-stride|fj-no-prefetch-stride} | {fj-prefetch-strong|fj-no-prefetch-strong} | {fj-prefetch-strong-L2|fj-no-prefetch-strong-L2} | {fj-promote-licm-addressing|fj-no-promote-licm-addressing} | {fj-regalloc-using-latency|fj-no-regalloc-using-latency} | {fj-sched-insn-contiguous|fj-no-sched-insn-contiguous} | {fj-swp|fj-no-swp} | {fj-zfill[=N]|fj-no-zfill} | {builtin|no-builtin} | {fast-math|no-fast-math} | {finite-math-only|no-finite-math-only} | fp-contract={fast|on|off} | {inline-functions|no-inline-functions} | {lto|no-lto} | {omit-frame-pointer|no-omit-frame-pointer} | {openmp|no-openmp} | {openmp-simd|no-openmp-simd} | {reroll-loops|no-reroll-loops} | {signed-char|unsigned-char} | {slp-vectorize|no-slp-vectorize} | {strict-aliasing|no-strict-aliasing} | {unroll-loops|no-unroll-loops} | {vectorize|no-vectorize} }
```

#### `-f{fj-eval-concurrent|fj-no-eval-concurrent}`

`-ffj-eval-concurrent` specifies to improve the parallelism of floating-point arithmetic instructions using the optimization of tree-height-reduction. `-ffj-no-eval-concurrent` gives priority to the optimization using FMA instructions over tree-height-reduction. `-ffj-no-eval-concurrent` is set by default.

`-ffj-eval-concurrent` and `-ffj-no-eval-concurrent` options require that `-O1` option or higher is effective and the `-ffast-math` option is effective.

#### `-f{fj-fast-matmul|fj-no-fast-matmul}`

The compiler option `-ffj-fast-matmul` performs the optimization that converts the loop of matrix multiplication into a high speed library call. The compiler option `-ffj-no-fast-matmul` invalidates the compiler option `-ffj-fast-matmul`. The compiler option `-ffj-no-fast-matmul` is set by default.

When the compiler option `-ffj-fast-matmul` is set, side effects (calculation errors) may occur in the execution results. See "[9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

The compiler options `-ffj-fast-matmul` and `-ffj-no-fast-matmul` require that the compiler option `-O2` or higher is set.

The compiler option `-ffj-fast-matmul` cannot be specified simultaneously with `-shared` option.

The compiler option `-ffj-fast-matmul` is needed if an object program compiled with it exists in the command line as an input file.

#### `-f{fj-fp-precision|fj-no-fp-precision}`

`-ffj-fp-precision` specifies to induce the combination of optimization options that do not cause calculation errors in floating-point operations. `-ffj-no-fp-precision` option invalidates `-ffj-fp-precision` option. `-ffj-no-fp-precision` is set by default.

The `-ffj-fp-precision` option is equivalent to replacing the `-ffj-fp-precision` option with the following options:

`-fno-fast-math -ffp-contract=off -ffj-no-fast-matmul -ffj-no-fp-relaxed -ffj-no-ilfunc`

Therefore, if the `-ffj-fp-precision` option is enabled, some optimizations are limited and may decrease the performance.

The `-ffj-no-fp-precision` option does not affect the individual options induced by `-ffj-fp-precision` option, even if they are specified together.

#### `-f{fj-fp-relaxed|fj-no-fp-relaxed}`

`-ffj-fp-relaxed` specifies to convert floating point division operations and `sqrt` functions into the reciprocal approximation operation with the reciprocal approximation instructions and the Floating-Point Multiply-Add/Subtract instructions. This conversion is applied to the single-precision and double-precision real type calculation. `-ffj-no-fp-relaxed` option invalidates `-ffj-fp-relaxed` option. `-ffj-no-fp-relaxed` is set by default.

When `-ffj-fp-relaxed` is set, side effects may occur in the execution results. See "[9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

`-ffj-fp-relaxed` and `-ffj-no-fp-relaxed` options require that `-O1` option or higher is set.

#### `-f{fj-hpctag|fj-no-hpctag}`

`-ffj-hpctag` specifies to use the HPC tag address override feature of the A64FX processor. If the `-ffj-hpctag` option is specified, the HPC tag address override feature is used. The default is `-ffj-hpctag`.

Using the HPC tag address override feature, the hardware prefetch assistance feature become effective.

This option is effective when the `-mcpu=a64fx` option is set.

`-ffj-hpctag` and `-ffj-no-hpctag` options should be specified when programs are compiled and linked.

#### `-f{fj-ilfunc[={loop|procedure}]|fj-no-ilfunc}`

`-ffj-ilfunc` performs the inline expansion for the math functions. `-ffj-no-ilfunc` option invalidates `-ffj-ilfunc` option. `-fbuiltin` option is required for `-ffj-ilfunc` option to take effect. `-ffj-no-ilfunc` is set by default.

If the argument `={loop|procedure}` is omitted, `-ffj-ilfunc=procedure` is set.

When `-ffj-ilfunc` option is set, side effects may occur in the execution results. See "[9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

`-ffj-ilfunc` and `-ffj-no-ilfunc` options require that `-O1` option or higher is set.

The following shows the functions for which inline expansion can be performed.

- Single-precision and double-precision real type function

```
atan, atan2, cos, exp, log, log10, pow, sin, tan
```

- Single-precision and double-precision complex type function

```
abs, exp
```

#### `-ffj-ilfunc=loop`

Specifies to perform the inline expansion for the math functions in a loop. Inline expansion is applied when the optimization may be promoted by checking a data type, an existence of function call and so on.

#### `-ffj-ilfunc=procedure`

Specifies to perform the inline expansion for all the math functions in a function.

### -{ffj-interleave-loop-insns[=*N*]|fj-no-interleave-loop-insns}

-ffj-interleave-loop-insns specifies to apply loop interleaving to SIMD applied loops using SVE.

*N* is the interleaving count which should be from 2 to 10. If =*N* is omitted, the compiler determines a suitable value for *N*. The -ffj-no-interleave-loop-insns option invalidates the -ffj-interleave-loop-insns option. -ffj-no-interleave-loop-insns option is set by default.

This option is effective only when -fvectorize option is effective and sve is enabled at *features* of -march option.

Loop interleaving performs the same optimization as loop striping in Trad Mode.

### -{ffj-loop-fission|fj-no-loop-fission}

-ffj-loop-fission option specifies to perform the automatic loop fission optimization for the software pipelining, SIMD Extensions, and the resolution of register shortage. -ffj-no-loop-fission option invalidates -ffj-loop-fission option.

This loop fission optimization is performed for the loop with the loop loop\_fission\_target specifier when -ffj-loop-fission option and -ffj-ocl option are enabled. -ffj-no-loop-fission option is always applied when -O1 option is set. -ffj-loop-fission option is set by default when -O2 option or higher is set.

-ffj-loop-fission option works only when used in combination with the loop loop\_fission\_target specifier.

The algorithm of the loop fission optimization by -ffj-loop-fission option is correspond to the clustering algorithm of Trad Mode.

The loop fission optimization by -ffj-loop-fission option tends to use large stack area of the memory.

### -ffj-loop-fission-threshold=*N*

Specifies the threshold *N* to decide the granularity of loops after the automatic loop fission.

*N* is an integer value from 1 to 100. If the value of *N* is reduced, the fissioned loops tend to become small and the number of the fissioned loops tends to increase. -ffj-loop-fission-threshold=50 is set by default.

This option requires that -ffj-loop-fission option, -ffj-ocl option, and -O2 option or higher are set and the loop loop\_fission\_target specifier is specified in the optimization control line.

### -ffj-no-prefetch

Specifies to generate an object without a prefetch instruction.

### -{ffj-ocl|fj-no-ocl}

-ffj-ocl specifies that the optimization control line that is supported as Fujitsu-specific pragma supported by Clang Mode is effective.

-ffj-no-ocl option invalidates -ffj-ocl option. -ffj-no-ocl is set by default.

-ffj-ocl and -ffj-no-ocl options require that -O1 option or higher is set.

### -{ffj-stdlib-string|fj-no-stdlib-string}

Indicates whether or not to link the optimized version of the library in the string manipulation function. Statically links the optimized library if the -ffj-stdlib-string option is specified. If omitted, the -ffj-no-stdlib-string option applies.

The -ffj-stdlib-string option is enabled when specified with the -mcpu=a64fx[+sve] option.

The -ffj-stdlib-string option cannot be specified simultaneously with the -flto option when the -shared option is specified.

This option must be specified when compiling and linking the program.

The following string manipulation functions are supported:

bcopy, bzero, memchr, memcmp, memccpy, memcpy, memmove, memset, strcat, strcmp, strcpy, strlen, strncmp, strncpy, strncat
---

### -{ffj-preex|fj-no-preex}

-ffj-preex specifies to evaluate invariant expressions in advance. -ffj-no-preex option invalidates -ffj-preex option. -ffj-no-preex is set by default.

When -ffj-preex option is set, side effects (calculation error or runtime exceptions) may occur in the execution results because instructions that may not to be executed based on the logic of the program are likely to be executed.

-ffj-preex and -ffj-no-preex options require that -O1 option or higher is set.

**-ffj-prefetch-cache-level={1|2|all}**

Specifies the data cache-level of prefetch instructions. By default, **-ffj-prefetch-cache-level=all** is set.

This option requires that at least one of **-ffj-prefetch-sequential** or **-ffj-prefetch-stride** options is set.

**-ffj-prefetch-cache-level=1**

Data is prefetched in the first level cache. Normal prefetch instructions are used.

**-ffj-prefetch-cache-level=2**

Data is prefetched in the second level cache.

**-ffj-prefetch-cache-level=all**

Both **-ffj-prefetch-cache-level=1** and **-ffj-prefetch-cache-level=2** functions are effective. By using two levels of prefetch instructions, the prefetch function becomes more sophisticated.

**-f{fj-prefetch-conditional|fj-no-prefetch-conditional}**

**-ffj-prefetch-conditional** specifies to generate prefetch instructions for array data in the block included in an **"if"** statement and a **"switch"** statement. **-ffj-no-prefetch-conditional** option invalidates **-ffj-prefetch-conditional** option. **-ffj-no-prefetch-conditional** is set by default.

**-ffj-prefetch-conditional** and **-ffj-no-prefetch-conditional** options require that at least one of **-ffj-prefetch-sequential** or **-ffj-prefetch-stride** options is set.

**-ffj-prefetch-iteration=N**

Specifies to prefetch data which is referred to or defined in a loop after *N* iteration(s). *N* should be from 1 to 10000.

This option specifies to prefetch data to only the first level cache.

This option requires that at least one of **-ffj-prefetch-sequential** or **-ffj-prefetch-stride** options is set, and **-ffj-prefetch-cache-level=1** or **-ffj-prefetch-cache-level=all** is set.

This option cannot be specified simultaneously with **-ffj-prefetch-line** option.

**-ffj-prefetch-iteration-L2=N**

Specifies to prefetch data which is referred to or defined in a loop after *N* iteration(s). *N* should be from 1 to 10000.

This option specifies to prefetch data to only the second level cache.

This option requires that at least one of **-ffj-prefetch-sequential** or **-ffj-prefetch-stride** options is set, and **-ffj-prefetch-cache-level=2** or **-ffj-prefetch-cache-level=all** is set.

This option cannot be specified simultaneously with **-ffj-prefetch-line-L2** option.

**-ffj-prefetch-line=N**

Specifies to prefetch data which is referred to or defined in a line after *N* line(s). *N* should be from 1 to 100. This option specifies to prefetch data to only the first level cache.

This option requires that **-ffj-prefetch-sequential** option is set, and **-ffj-prefetch-cache-level=1** or **-ffj-prefetch-cache-level=all** is set.

This option cannot be specified simultaneously with **-ffj-prefetch-iteration** option.

**-ffj-prefetch-line-L2=N**

Specifies to prefetch data which is referred to or defined in a line after *N* line(s). *N* should be from 1 to 100. This option specifies to prefetch data to only the second level cache.

This option requires that **-ffj-prefetch-sequential** option is set, and **-ffj-prefetch-cache-level=2** or **-ffj-prefetch-cache-level=all** is set.

This option cannot be specified simultaneously with **-ffj-prefetch-iteration-L2** option.

**-f{fj-prefetch-sequential[={auto|soft}]|fj-no-prefetch-sequential}**

**-ffj-prefetch-sequential** specifies to generate prefetch instructions for prefetch array data accessed sequentially within a loop. **-ffj-no-prefetch-sequential** option invalidates **-ffj-prefetch-sequential** option. When **-O1** option is set, **-ffj-prefetch-nosequential** is set by default. When **-O2** or higher option is set, **-ffj-prefetch-sequential** is set by default.

If the argument **={auto|soft}** is omitted, **-ffj-prefetch-sequential=auto** is assumed.

-ffj-prefetch-sequential and -ffj-no-prefetch-sequential options require that -O1 option or higher is set.

-ffj-prefetch-sequential=auto

The compiler automatically selects whether to use the hardware prefetch function or to output the prefetch instruction for array data accessed sequentially within a loop.

-ffj-prefetch-sequential=soft

The compiler specifies to output the prefetch instruction for array data accessed sequentially within a loop without using the hardware prefetch function.

-f{fj-prefetch-stride|fj-no-prefetch-stride}

-ffj-prefetch-stride specifies to generate prefetch instructions for prefetch array data with a stride larger than the cache line size in its loop. In addition, the loops include prefetch instructions to prefetch addresses not determined clearly during compilation. -ffj-no-prefetch-stride option invalidates -ffj-prefetch-stride option. -ffj-no-prefetch-stride is set by default.

-ffj-prefetch-stride and -ffj-no-prefetch-stride options require that -O1 option or higher is set.

-f{fj-prefetch-strong|fj-no-prefetch-strong}

-ffj-prefetch-strong specifies to generate strong prefetch instructions. -ffj-no-prefetch-strong option invalidates -ffj-prefetch-strong option. -ffj-prefetch-strong is set by default.

-ffj-prefetch-strong option specifies to prefetch data to only the first level cache.

-ffj-prefetch-strong option requires that at least one of -ffj-prefetch-sequential or -ffj-prefetch-stride options is set, and -ffj-prefetch-cache-level=1 or -ffj-prefetch-cache-level=all option is set.

-ffj-no-prefetch-strong is effective if -ffj-hpctag option is set.

-f{fj-prefetch-strong-L2|fj-no-prefetch-strong-L2}

This option specifies to generate strong prefetch instructions. -ffj-no-prefetch-strong-L2 option invalidates -ffj-prefetch-strong-L2 option. -ffj-prefetch-strong-L2 is set by default.

-ffj-prefetch-strong-L2 option specifies to prefetch data to only the second level cache.

-ffj-prefetch-strong-L2 option requires that at least one of -ffj-prefetch-sequential or -ffj-prefetch-stride options is set, and -ffj-prefetch-cache-level=2 or -ffj-prefetch-cache-level=all option is set.

-ffj-no-prefetch-strong-L2 is effective if -ffj-hpctag option is set.

-f{fj-promote-licm-addressing|fj-no-promote-licm-addressing}

-ffj-promote-licm-addressing specifies to perform before SIMDization the optimization which extract loop invariant terms from address calculations to promote movement of invariant expressions.

-ffj-no-promote-licm-addressing option invalidates -ffj-promote-licm-addressing option.

-ffj-promote-licm-addressing option is set by default when -O1 option or higher is set.

-ffj-promote-licm-addressing option requires -O1 option or higher is set.

-f{fj-regalloc-using-latency|fj-no-regalloc-using-latency}

-ffj-regalloc-using-latency specifies that at the register allocation for instructions in a loop, the registers for spill are determined in consideration of the instruction latency.

-ffj-no-regalloc-using-latency option invalidates -ffj-regalloc-using-latency option.

-ffj-regalloc-using-latency option is set by default when -O1 option or higher is set.

-ffj-regalloc-using-latency option requires -O1 option or higher is set.

-f{fj-sched-insn-contiguous|fj-no-sched-insn-contiguous}

-ffj-sched-insn-contiguous specifies that the instruction scheduler before and after the register allocation order each of floating-point arithmetic instructions, load instructions, and store instructions sequentially if possible.

-ffj-no-sched-insn-contiguous option invalidates -ffj-sched-insn-contiguous option.

-ffj-sched-insn-contiguous option is set by default when -O1 option or higher is set.

-ffj-sched-insn-contiguous option requires -O1 option or higher is set.

#### -f{fj-swp|fj-no-swp}

The -ffj-swp option performs software pipelining. The -ffj-no-swp option invalidates the -ffj-swp option. The -ffj-no-swp option is set by default. These options are effective only if the -O1 option or higher and -mcpu=a64fx option are specified.

No optimization is performed for the following loops.

- Pipelining is ineffective.
- Vectorized by the -msve-vector-bits=scalable option.

The -ffj-swp option induces the -msve-vector-bits=512 option. The induction is disabled when the -flto option is enabled. Also, the induction is disabled when the -msve-vector-bits=scalable option is specified after the -ffj-swp option. The induction of the -msve-vector-bits=512 option does not depend on the existence of the `loop swp` specifier.

Pipelining may not be performed if the -g option is set.

When specifying the -Rpass=.\* , -Rpass-analysis=.\* , and -Rpass-missed=.\* options, the optimization messages for software pipelining are for one or more loops generated by the compiler from a single loop explicitly written in the source program. Therefore, if multiple loops are generated by the compiler, different multiple messages may be emitted for a single loop on the source file.

Specifying the -Rpass=.\* option, without the -Rpass-analysis=.\* and -Rpass-missed=.\* options, can avoid confusion in the messages resulting from the loops generated by the compiler, and in this case the message "remark: loop pipelined" means that software pipelining is applied to one of these loops.

#### -f{fj-zfill[=N]|fj-no-zfill}

The -ffj-zfill option specifies to perform zfill optimization. The zfill optimization speeds up write operations for array data that is only written in a loop, by using an instruction that allocates space on the cache for writing (DC ZVA) without loading data from the memory. The zfill optimization works on the data *N* blocks ahead of the address pointed to by the target store instruction where one block is 256 byte-long and *N* is an integer value between 1 and 100. If a value is not specified for *N*, the compiler will automatically determine a value. The -ffj-no-zfill option invalidates the -ffj-zfill option. The -ffj-no-zfill option is always applied when the -O1 option is effective. The -ffj-no-zfill option is set by default when the -O2 option or higher is set.

The -ffj-zfill option requires that the -mcpu=a64fx option, -msve-vector-bits=512 option, and -O2 option or higher are enabled. The -ffj-zfill option induces the -msve-vector-bits=512 option. The induction is disabled when the -flto option is enabled. Also, the induction is disabled when the -msve-vector-bits=scalable option is specified after the -ffj-zfill option. The induction of the -msve-vector-bits=512 option does not depend on the existence of the `loop zfill` specifier.

Note that if an object program compiled with the -ffj-zfill option is executed on a CPU other than the one on which a single cache write operation of a DC ZVA instruction is 256 bytes, the execution may be terminated abnormally or an incorrect result may occur. The amount of the cache write operation for a DC ZVA instruction on an A64FX processor is 256 bytes.

Performance may also be reduced under the following conditions:

- The program is not affected by memory bandwidth bottleneck.
- When the loop iteration count is too few.
- When the block size is explicitly specified with the -ffj-zfill=*N* option and the memory size where the data is written in the loop is smaller than *N* blocks.

This optimization is not expected to be effective for all the loops in the program, so it is recommended to use the `loop zfill` specifier for each loop rather than the -ffj-zfill[=*N*] option for the entire program.

For details about zfill, see Section "3.3.5 zfill".

#### -f{builtin|no-builtin}

-fbuiltin promotes optimization by recognizing the operation of standard library functions. The -fno-builtin option invalidates -fbuiltin option. -fbuiltin is set by default.

If the user-defined function with the same name as a standard library function is used, unexpected results may occur.

The recognized standard library functions are as follows:

```
abort, abs, acos, acosf, acoshf, asin, asinf, asinhf, atan, atan2, atan2f, atanf, atanhf,
calloc, cbrt, cbrtf, ceil, ceilf, clearerr, copysign, copysignf, cos, cosf, cosh, coshf, csqrt,
```

```
csqrtf, erf, erfc, erfcf, erff, exit, exp, exp2, exp2f, expf, expmlf, fabs, fabsf, fclose, feof,
ferror, fflush, fgetc, fgetpos, fgets, floor, floorf, fma, fmaf, fmax, fmaxf, fmin, fminf, fmod,
fmodf, fputc, fputs, free, frexp, fseek, fwrite, getenv, hypotf, ilogb, ilogbf, isalnum,
isalpha, iscntrl, isdigit, isgraph, islower, ispunct, isspace, isupper, isxdigit, ldexp,
lgammaf, log, log10, log10f, loglpf, log2, log2f, logbf, logf, malloc, memchr, memcmp, memcpy,
memmove, memset, modf, nextafterf, perror, pow, powf, printf, putchar, rand, realloc,
remainderf, remove, rename, rint, rintf, scalbnf, scanf, setvbuf, sin, sinf, sinh, sinhf,
sprintf, sqrt, sqrtf, srand, sscanf, strcat, strchr, strcmp, strcpy, strcspn, strerror, strlen,
strncat, strncmp, strncpy, strpbrk, strrchr, strspn, strstr, strtod, strtok, strtol, strtoul,
system, tan, tanf, tanh, tanhf, tolower, toupper, vprintf, vsprintf
```

#### `-f{fast-math|no-fast-math}`

`-ffast-math` performs the optimization that changes the method of operator evaluation. `-fno-fast-math` option invalidates `-ffast-math` option. `-fno-fast-math` is set by default.

When the `-ffast-math` option is specified, calculation errors or the runtime exceptions may occur in the execution results and produce unexpected results due to the effects of optimization or flush-to-zero mode which replaces the denormalized numbers with zero with the same sign.

See "[9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

The `-ffast-math` and `-fno-fast-math` options require that `-O1` option or higher is set.

The `-ffast-math` and `-fno-fast-math` options must be set at compiling and linking.

#### `-f{finite-math-only|no-finite-math-only}`

`-ffinite-math-only` option promote the optimization of floating point arithmetic based on the assumption that the argument or operation result is only a finite numerical value.

`-ffast-math` option induces the `-ffinite-math-only` option.

#### `-ffp-contract={fast|on|off}`

`-ffp-contract=fast` specifies to perform optimization using Floating-Point Multiply-Add/Subtract instructions. `-ffp-contract=on` specifies to perform optimization using Floating-Point Multiply-Add/Subtract instructions as if `#pragma STDC FP_CONTRACT ON` (language standard) was specified in the source code. `-ffp-contract=off` option invalidates `-ffp-contract={fast|on}`. `-ffp-contract=off` is set by default.

When `-ffp-contract={fast|on}` option is effective, side effects (calculation errors in rounding error extent) may occur in the execution results and produce unexpected results. See "[9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation](#)" for the side effect of optimization.

`-ffp-contract={fast|off}` options require that `-O1` option or higher is set.

`-ffp-contract=on` option requires that `-O0` option or higher is set.

#### `-f{inline-functions|no-inline-functions}`

`-finline-functions` specifies to perform Inline expansion of function calls defined in source code.

`-finline-functions` is set by default when `-O2` option or higher is set.

This option is effective only if `-O2` option or higher is set.

Therefore, the inline function does not become the target of the inline expansion when the compiler judges that inline expansion should not be performed.

#### `-f{lto|no-lto}`

`-flto` specifies to perform the link time optimization. `-fno-lto` option invalidates `-flto` option. `-flto` option should be specified when programs are compiled and linked. `-fno-lto` option is assumed by default.

If the `-ffj-line` option is specified with `-flto` option, information about loops required by the Profiler's execution time sampling feature will not be generated. See the "Profiler User's Guide" about the Profiler.

If `-flto` option is set, the only options for optimization that can be specified at the same time are the following.

- `-O[N]` option
- The options described to be set at linking in Section "[9.1.2.2.3 Options for Optimization](#)"

The `-flt` option cannot be specified simultaneously with the `-ffj-stdlib-string` option when the `-shared` option is specified.

If both the `-flt` and `-msve-vector-bits=512` options are specified, the `-msve-vector-bits=512` option is disabled and the `-msve-vector-bits=scalable` option is enabled.

`-f{omit-frame-pointer|no-omit-frame-pointer}`

Indicates whether or not to perform the optimization that the frame pointer register is not kept for a function call.

`-fomit-frame-pointer` option does the optimization. Trace back information is not kept.

If `-fomit-frame-pointer` option is specified, `-fomit-leaf-frame-pointer` option is induced.

If `-fno-omit-frame-pointer` option is specified, `-fno-omit-leaf-frame-pointer` option is induced.

By default, `-fno-omit-frame-pointer` option is set, but in this case `-fno-omit-leaf-frame-pointer` option is not induced.

`-f{openmp|no-openmp}`

`-fopenmp` specifies to accept the Specification of OpenMP Application Program Interface. `-fno-openmp` option invalidates `-fopenmp` option. `-fno-openmp` is set by default.

`-fopenmp` option is needed if an object program compiled with it exists in the command line as an input file.

`-f{openmp-simd|no-openmp-simd}`

The `-fopenmp-simd` specifies that only the OpenMP `simd` construct, the `declare reduction` construct, and the `ordered` construct with the `simd` clause are enabled. `-fno-openmp-simd` is set by default.

When the `-fopenmp-simd` option is specified, only the SIMD Extensions based on the OpenMP specifications are applied, and the parallelization is not applied.

When the `-fno-openmp-simd` option is specified, the `-fopenmp_simd` option is disabled.

When the `-O1` option or less is enabled, SIMD Extensions are not used even if `simd` construct is enabled.

`-f{reroll-loops|no-reroll-loops}`

`-freroll-loops` specifies to apply loop rerolling. `-fno-reroll-loops` is set by default.

`-fsigned-char`

Specifies to treat `char` type variable as signed `char` type.

`-funsigned-char`

Specifies to treat `char` type variable as unsigned `char` type.

`-f{slp-vectorize|no-slp-vectorize}`

`-fslp-vectorize` specifies to apply SLP(Superword Level Parallelism) using SIMD instructions.

When the `-msve-vector-bits=512` option is effective, the `-fslp-vectorize` option is invalidated.

When `-O1` option is set, `-fno-slp-vectorize` is set by default. When `-O2` option or higher is set, `-fslp-vectorize` is set by default.

`-f{strict-aliasing|no-strict-aliasing}`

`-fstrict-aliasing` option specifies to perform optimizations with considering overlaps of memory regions according to the strict aliasing rules in the language standards. `-fno-strict-aliasing` option invalidates `-fstrict-aliasing` option. When `-O0` option is set, `-fno-strict-aliasing` is always set. When `-O1` or higher option is set, `-fstrict-aliasing` is set by default.

`-fstrict-aliasing` and `-fno-strict-aliasing` options require that `-O1` option or higher is set.

`-f{unroll-loops|no-unroll-loops}`

`-funroll-loops` specifies to apply loop unrolling. The suitable multiplicity of loop unrolling is determined by the compiler.

`-fno-unroll-loops` option invalidates `-funroll-loops` option. `-fno-unroll-loops` is set by default when `-O1` option is set. `-funroll-loops` is set by default when `-O2` option or higher is set.

`-f{vectorize|no-vectorize}`

`-fvectorize` specifies to perform optimization using the SIMD Extensions for loop. `-fno-vectorize` option invalidates `-fvectorize` option.

When `-O1` option is set, `-fno-vectorize` is set by default. When `-O2` option or higher is set, `-fvectorize` is set by default.



`-Kopt`

`-Kopt` listed in "[Table 9.6 Options replaced in Clang Mode](#)" is replaced with the corresponding option of Clang Mode.

#### 9.1.2.2.4 Options for Language Specifications

This section explains the options relating to language specifications.

`-std=name`

This option specifies the level of language specifications to be interpreted by the compiler (including the preprocessor).

For *name*, specify one of the following:

```
{ c++03 | c++11 | c++14 | c++17 | gnu++03 | gnu++11 | gnu++14 | gnu++17 }
```

`-std=gnu++14` is set by default.

`-std=c++03`

The C++ language specification used at compilation time is based on the C++03 specifications.

`-std=c++11`

The C++ language specification used at compilation time is based on the C++11 specifications.

`-std=c++14`

The C++ language specification used at compilation time is based on the C++14 specifications.

`-std=c++17`

The C++ language specification used at compilation time is based on the C++17 specifications.

`-std=gnu++03`

In addition to the C++03 specifications, GNU C++ Extensions is used at compilation time.

`-std=gnu++11`

In addition to the C++11 specifications, GNU C++ Extensions is used at compilation time.

`-std=gnu++14`

In addition to the C++14 specifications, GNU C++ Extensions is used at compilation time.

`-std=gnu++17`

In addition to the C++17 specifications, GNU C++ Extensions is used at compilation time.

`-stdlib=name`

Specifies the kind of C++ Standard Library.

For *name*, specify one of the following:

```
{ libc++ | libstdc++ }
```

`-stdlib=libstdc++` is set by default.

`-stdlib=libc++`

The libc++ Version 7.0 is used.

`-stdlib=libstdc++`

The libstdc++ is used.

It is recommended to use `-stdlib=libc++` when processing based on the language specification of the latest C++ language standard.

`--linkcoarray`

Directs to search for a library required for linkage with Fortran that does use the COARRAY specification.

If this option is not specified, search processing is not performed, enabling a reduction in compilation time.

--linkfortran

Directs to search for a library required for linkage with Fortran that does not use the COARRAY specification.

If this option is not specified, search processing is not performed, enabling a reduction in compilation time.

### 9.1.2.2.5 Options for CPU/Architecture

This section describes about the optimization options of the processor and the architecture.

**-march=arch[+features]...**

Specifies to generate the object file that is suitable for the given name of architecture and extension.

For *arch*, specify one of the following:

```
{ armv8-a | armv8.1-a | armv8.2-a | armv8.3-a }
```

**-march=armv8-a**

Specifies to generate object files using instructions in Armv8-A.

**-march=armv8.1-a**

Specifies to generate object files using instructions in Armv8-A and Armv8.1-A.

**-march=armv8.2-a**

Specifies to generate object files using instructions in Armv8-A, Armv8.1-A, and Armv8.2-A.

**-march=armv8.3-a**

Specifies to generate object files using instructions in Armv8-A, Armv8.1-A, Armv8.2-A, and Armv8.3-A.

For *features*, specify one or more of the following by separating them with + :

```
{ sve | nosve | fp16 }
```

**sve**

Specifies to output object files using SVE, which is an Armv8-A extension.

**nosve**

Specifies to output object files without the use of SVE, which is an Armv8-A extension.

Specify the argument +nosve to generate object files for a processor that does not support SVE.

**fp16**

Specifies to generate the object file using the extension of the half-precision floating-point. About the language specifications, see Section "9.5.2 Half-Precision (16 bit) Floating-Point Type".

If this option is omitted, the values corresponding to the -mcpu option are set in *arch* and *features* of -march option.

Table 9.4 Default Values of -march Option

Effective Value of -mcpu Option	Default Value of -march Option
-mcpu=a64fx[+sve]	-march=armv8.2-a+sve+fp16
-mcpu=a64fx+nosve	-march=armv8.2-a+nosve+fp16
-mcpu=generic[+nosve]	-march=armv8-a+nosve
-mcpu=thunderx[+nosve]	
-mcpu=generic+sve	-march=armv8-a+sve
-mcpu=thunderx+sve	
-mcpu=thunderx2t99[+nosve]	-march=armv8.1-a+nosve
-mcpu=thunderx2t99+sve	-march=armv8.1-a+sve

## Note

If the `-mcpu` and `-march` options are specified simultaneously, the architecture and *features* specified by the `-march` option take precedence.

### `-Karch`

`-Karch` option listed in "Table 9.6 Options replaced in Clang Mode" is replaced with the corresponding `-march` option of Clang Mode.

### `-mcpu=cputype[+features]...`

Specifies to generate the object file that is suitable for the given name of processor.

For *cputype*, specify one of the following:

```
{ a64fx | generic | thunderx | thunderx2t99 }
```

#### `-mcpu=a64fx`

`-mcpu=a64fx` specifies to output the object file for the A64FX processor.

#### `-mcpu=generic`

`-mcpu=generic` specifies to output the object file for the Arm processor.

#### `-mcpu=thunderx`

`-mcpu=thunderx` specifies to output the object file for the ThunderX processor.

#### `-mcpu=thunderx2t99`

`-mcpu=thunderx2t99` specifies to output the object file for the ThunderX2 processor.

For *features*, specify one or more of the following by separating them with `+`:

```
{ sve | nosve | fp16 }
```

#### `sve`

Specifies to output object files using SVE, which is an Armv8-A extension.

#### `nosve`

Specifies to output object files without the use of SVE, which is an Armv8-A extension.

Specify the argument `+nosve` to generate object files for a processor that does not support SVE.

#### `fp16`

Specifies to generate the object file using the extension of the half-precision floating-point. About the language specifications, see Section "9.5.2 Half-Precision (16 bit) Floating-Point Type".

By default, `-mcpu=a64fx` is set.

## Note

If the `-mcpu` and `-march` options are specified simultaneously, the architecture and *features* specified by the `-march` option take precedence.

### `-Kcputype`

`-Kcputype` option listed in "Table 9.6 Options replaced in Clang Mode" is replaced with the corresponding `-mcpu` option of Clang Mode.

## 9.1.2.2.6 Options for Code Generation

This section describes about the options of code generation.

## **-fopt**

For *opt*, specify one of the following:

```
{ {PIC|pic} | {ffj-largepage|ffj-no-largepage} }
```

### **-f{PIC|pic}**

Specifies to generate position-independent code (PIC).

### **-f{ffj-largepage|ffj-no-largepage}**

-ffj-largepage creates an executable program which uses the large page function. -ffj-no-largepage option invalidates -ffj-largepage option. -ffj-largepage is set by default.

When -fsanitize option is specified, -ffj-largepage option is invalidated.

-ffj-largepage and -ffj-no-largepage option must be set at linking.

## **-mcmode=*name***

This option specifies the possible largest size of the text area and the static data area in an executable program or a shared object.

For *name*, specify one of the following:

```
{small|large}
```

-mcmode=small is set by default.

### **-mcmode=small**

The total size of the text area and the static data area is limited to 4GB at linking. This option creates an efficient object program.

### **-mcmode=large**

Only the size of the text area is limited to 4GB at linking. This option is used when the static data area is large and an error occurs at linking.

## **-mfj-tls-size={12|24|32|48}**

Specifies the size of an offset necessary for the access to Thread-Local Storage. Units are bits.

As the size of Thread-Local Storage, -mfj-tls-size=12 (4K bytes), -mfj-tls-size=24 (16M bytes), -mfj-tls-size=32 (4G bytes) or -mfj-tls-size=48 (256T bytes) can be specified.

When the size of the Thread-Local Storage exceeds the range of the offset, an error occurs at link time.

This option cannot be specified simultaneously with the -flto option.

## **-Kopt**

-Kopt option listed in "[Table 9.6 Options replaced in Clang Mode](#)" is replaced with the corresponding option of Clang Mode.

## 9.1.2.3 Notes on Compile Options

This section describes notes on the compile options.

### 9.1.2.3.1 Correspondences of Compile Options in Clang Mode and Trad Mode

The compile option system of Clang Mode is different from that of Trad Mode. You should pay attention to change the compiler mode when it used for compiling.

The compile options whose option name and function overview are the same in Clang Mode and Trad Mode are described in "[Table 9.5 Common compile options in Clang/Trad Mode](#)". The compile options whose function overview is same but option name is different is described in "[Table 9.6 Options replaced in Clang Mode](#)".

- A compile option of Trad Mode which is not described in "[Table 9.6 Options replaced in Clang Mode](#)" can not be specified in Clang Mode. If it is specified, an error message is output and the compilation is canceled.
- The -K{SVE|NOSVE} options of Trad Mode cannot be specified in Clang Mode. If specified, it outputs an error message and aborts the compilation. To indicate whether to use SVE, add +{sve|nosve} to the -mcpu or -march option.

- If a compile option of Trad Mode which is described in "[Table 9.6 Options replaced in Clang Mode](#)", a corresponding compile option of Clang Mode is assumed to be specified.
- Even if a compile option whose function overview is same is specified, an optimization result in Clang Mode may be different from that in Trad Mode.
  - If the -Kfast option is specified, the induced options and the architecture set by default are different from the Trad Mode.
    - For included options, see the -Ofast option in "[9.1.2.2.3 Options for Optimization](#)" for Clang Mode and the -Kfast option in "[2.2.2.6 -K Option](#)" for Trad Mode.
    - For architectures that are set by default, see the -march option in "[9.1.2.2.5 Options for CPU/Architecture](#)" for Clang Mode and the -Karchi option in "[2.2.2.6 -K Option](#)" for Trad Mode.
- The following options have different default behavior in Trad Mode and Clang Mode.
  - K{strict\_aliasing|nostrict\_aliasing}
 

Trad Mode : -Knostrict\_aliasing

Clang Mode : (if the -O0 option is set) -fnostrict-aliasing, (if the -O1 option or higher is set) -fstrict-aliasing
  - K{omitfp|noomitfp}
 

Trad Mode : -Knoomitfp

Clang Mode : -fno-omit-frame-pointer, and (only if -m {omit-leaf-frame-pointer | no-omit-leaf-frame-pointer} is not specified) -momit-leaf-frame-pointer

Table 9.5 Common compile options in Clang/Trad Mode

Option Name
-C
-Dname [=tokens]
-E
-Idir
-Ldir
-M
-MD
-MFfilename
-MM
-MMD
-MP
-MTtarget
-P
-S
-SSL2
-SSL2BLAMP
-Uname

Option Name
-Wtool, arg1[, arg2]...
-c
-g
-g0
-lname
-o pathname
-shared
-v
-w

Table 9.6 Options replaced in Clang Mode

Option name in Trad Mode	Option name in Clang Mode
-KA64FX	-mcpu=a64fx
-KARMV8_1_A	-march=armv8.1-a
-KARMV8_2_A	-march=armv8.2-a
-KARMV8_3_A	-march=armv8.3-a
-KARMV8_A	-march=armv8-a
-KGENERIC_CPU	-mcpu=generic
-K{PIC pic}	-f{PIC pic}
-Kcmodel={small large}	-mcmmodel={small large}
-K{eval noeval}	-f{fast-math no-fast-math}
-K{eval_concurrent eval_noconcurrent}	-ffj-{eval-concurrent no-eval-concurrent}
-Kfast	-Ofast
-K{fast_matmul nofast_matmul}	-ffj-{fast-matmul no-fast-matmul}
-K{fp_contract nofp_contract}	-ffp-contract={fast off}
-K{fp_relaxed nofp_relaxed}	-ffj-{fp-relaxed no-fp-relaxed}
-K{ilfunc[={loop procedure}] noilfunc}	-ffj-{ilfunc[={loop procedure}] no-ilfunc}
-K{largepage nolargepage}	-ffj-{largepage no-largepage}
-K{lib nolib}	-f{builtin no-builtin}
-K{loop_fission loop_nofission}	-ffj-{loop-fission no-loop-fission}
-Kloop_fission_threshold=N	-ffj-loop-fission-threshold=N
-Knoprefetch	-ffj-no-prefetch

Option name in Trad Mode	Option name in Clang Mode
-K{ocl noocl}	-ffj-{ocl no-ocl}
-K{omitfp noomitfp}	-f{omit-frame-pointer no-omit-frame-pointer}
-K{openmp noopenmp}	-f{openmp no-openmp}
-K{openmp_simd noopenmp_simd}	-f{openmp-simd no-openmp-simd}
-Koptmsg=2	-Rpass=.*
-K{preex nopreex}	-ffj-{preex no-preex}
-Kprefetch_cache_level={1 2 all}	-ffj-prefetch-cache-level={1 2 all}
-K{prefetch_conditional prefetch_noconditional}	-ffj-{prefetch-conditional no-prefetch-conditional}
-Kprefetch_iteration=N	-ffj-prefetch-iteration=N
-Kprefetch_iteration_L2=N	-ffj-prefetch-iteration-L2=N
-Kprefetch_line=N	-ffj-prefetch-line=N
-Kprefetch_line_L2=N	-ffj-prefetch-line-L2=N
-K{prefetch_sequential[={auto soft}] prefetch_nosequential}	-ffj-{prefetch-sequential[={auto soft}] no-prefetch-sequential}
-K{prefetch_stride prefetch_nostride}	-ffj-{prefetch-stride no-prefetch-stride}
-K{prefetch_strong prefetch_nostrong}	-ffj-{prefetch-strong no-prefetch-strong}
-K{prefetch_strong_L2 prefetch_nostrong_L2}	-ffj-{prefetch-strong-L2 no-prefetch-strong-L2}
-K{simd nosimd}	-f{vectorize no-vectorize}
-K{strict_aliasing nostrict_aliasing}	-f{strict-aliasing no-strict-aliasing}
-K{swp noswp}	-ffj-{swp no-swp}
-Kswp_strong	-ffj-swp
-Kswp_weak	
-K{unroll nounroll}	-f{unroll-loops no-unroll-loops}
-K{zfill[=N] nozfill}	-ffj-{zfill[=N] no-zfill}
-N{exceptions noexceptions}	-f{exceptions no-exceptions}
-N{fjcex nofjcex}	-ffj-{fjcex no-fjcex}
-N{fjprof nofjprof}	-ffj-{fjprof no-fjprof}
-N{hook_time nohook_time}	-ffj-{hook-time no-hook-time}
-N{line noline}	-ffj-{line no-line}
-Nlst	-ffj-lst
-Nlst=p	-ffj-lst=p

Option name in Trad Mode	Option name in Clang Mode
-Nlst=t	-ffj-1st=t
-Nlst_out=file	-ffj-1st-out=file
-Nsrc	-ffj-src
-V	--version
{-x- -x0}	-f{inline-functions no-inline-functions}

### 9.1.2.3.2 Side Effect of Optimizations for Floating-Point Operation

Optimizations for floating-point operation might cause the side effect. This section explains the side effect (mainly, computation error).


This system basically creates objects which comply with IEEE 754 arithmetic. Note that the numerical operations may not comply with IEEE 754 arithmetic due to the optimizations with calculation errors in "[Table 9.7 Side effect of optimization for floating-point operation](#)".

See Section "[9.1.2.2.3 Options for Optimization](#)" for compiler options. See Section "[9.2.2.1.2 Optimization Control Specifier](#)" for optimization control specifiers.

Table 9.7 Side effect of optimization for floating-point operation

Compiler Option	Optimization Control Specifier	Side Effect
-ffj-fast-matmul	-	Calculation errors may occur when using high speed library call for the loop of matrix multiplication.
-ffj-ilfunc[={loop procedure}]	-	Side effects similar to the ones caused by -ffj-fp-relaxed may occur when using inline expansion for the math functions, because reciprocal approximation instructions and trigonometric instructions, etc. are used. Plus, use of reciprocal approximation instructions or Floating-Point Multiply-Add/Subtract instructions regardless of set of -ffp-contract=off and/or -ffj-no-fp-relaxed options or their setting order.
-ffj-fp-relaxed	-	Side effects may occur because reciprocal approximation operation instructions are used on single-precision or double-precision floating point division or <code>sqrt</code> functions.  The side effects that may occur are: <ul style="list-style-type: none"> <li>- Rounding errors.</li> <li>- Replacing denormalized numbers found in the arguments or the return value with zero.</li> <li>- Replacing negative zeroes found in the arguments or the return value with positive zeroes.</li> <li>- Behaviors not conforming to IEEE 754 when NaN, positive or negative Inf, numbers which are close to maximal normalized number or numbers which are close to minimal normalized number are found in the arguments or the return value.</li> </ul>
-Ofast	-	-Ofast option induces the -ffast-math option, -ffj-ilfunc option, and so on.
-ffast-math	-	Calculation errors may occur when optimization that changes the method of operator evaluation. Moreover, the behaviors may not conforming to IEEE754.



Compiler Option	Optimization Control Specifier	Side Effect
		 <b>Example</b> <hr/> <div style="border: 1px solid black; padding: 5px; width: fit-content;"> <math>x*y + x*z \rightarrow x*(y+z)</math> </div> <hr/> <p>Moreover, flush-to-zero mode is used by <code>-ffast-math</code> option. If a result or a source operand is a denormalized number, flush-to-zero mode replaces it with zero with the same sign.</p>
<code>-ffp-contract={fast on}</code>	<code>fp contract({fast on})</code>	Rounding errors may occur when optimization using Floating-Point Multiply-Add/Subtract operation instructions is performed on source programs.

-: None

Since `#pragma STDC FENV_ACCESS ON` is not supported in this system, even if `#pragma STDC FENV_ACCESS ON` is specified in the program, a floating-point exception, that is not raised based on the logic of the program, may be raised at execution time. In this case, the following phenomena may occur.

- If a program uses the `fetestexcept` function included in standard library functions and accesses a floating-point status flag by it, the flag which should not be set based on the logic of the program is set.
- If a program uses the `feenableexcept` function included in the GNU C Library and enables a trap of a floating-point exception by it, the SIGFPE signal which should not be raised based on the logic of the program is raised.

This phenomenon may be avoided by setting the compiler option `-O0`, which suppresses optimizations.

### 9.1.2.3.3 Note on Using SVE

Even if the `-ffj-prefetch-sequential` or `-ffj-prefetch-stride` options are set, the prefetch instruction will not be generated for the loops which meet all of the following conditions:

- The induction variable increment is greater than 1.
- The size of the array type is less than 8 bytes.

### 9.1.2.3.4 Notes on Using SIMD Built-in Functions

This section provides notes on using SIMD built-in functions.

- If the shared library contains a function that has the SIMD built-in function with vector type as an argument, take one of the following actions. If the following actions are not taken, errors may occur in the execution results.
  - Specify the `-Wl,-z,now` options when the program linked.
  - Set 1 to the environment variable `LD_BIND_NOW` at runtime,

Refer to "[9.1.6.1 Environment Variable for Execution](#)" for the environment variable `LD_BIND_NOW`.

#### Example

SIMD built-in function with vector type as an argument

```
#include <arm_sve.h>
extern void subl(svint64_t p_val);
int main() {
    svint64_t p_val;
    p_val = svdup_n_s64(123);
    subl(p_val); /* Call function with vector type as argument */
    return 0;
}
```

- Note the SIMD built-in functions cannot be used when the `-msve-vector-bits=512` option is specified.

### 9.1.2.3.5 Notes on Specified SVE Vector Register Size

`-msve-vector-bits=512` option specifies the bitwise SVE vector register size. When this option is specified, the object program is generated considering the SVE vector register size as a fixed value at compilation. Therefore, the generated executable program works normally on the CPU architecture which has the same size of the SVE vector register as the specified size.

When `-msve-vector-bits=512` option is specified, and the program is executed on the CPU whose implemented SVE vector register size is different from 512 bits, an abnormal end occurs. And, the result of executions is not guaranteed.

When `-msve-vector-bits=scalable` option is effective, the executable program does not depend on the SVE vector register size.

See also Section "[B.2.9 Changing SVE Vector Register Size](#)" for notes on changing the SVE vector register size.

## 9.1.3 Environment Variable for Compile Command

---

This section explains environment variables that are valid only in Clang Mode.

For information about common environment variables in Clang/Trad Mode, refer to "[2.3 Environment Variable for Compile Command](#)".

`FCCpx_clang_ENV`  
`FCC_clang_ENV`

Environment variables to set compiler options.

The environment variable `FCCpx_clang_ENV` is for the cross compiler and the environment variable `FCC_clang_ENV` is for the native compiler.

The user can specify Clang Mode specific and Clang/Trad Mode common compiler options for the value of these environment variables. Refer to "[9.1.2.3.1 Correspondences of Compile Options in Clang Mode and Trad Mode](#)" for a description of the common compiler options for Clang/Trad Mode.

Compiler options specified to these environment variables are valid only in Clang Mode.

Refer to "[2.4 Compilation Profile File](#)" for the priority of compiler options.

## 9.1.4 Compilation Profile File

---

The default values of compiler options can be changed by specifying the compilation profile file.

The following compilation profile files are available:

Table 9.8 Compilation Profile Files (Clang Mode)

Kind	Mode	File Name
Cross compiler	Clang/Trad Mode common	<code>/etc/opt/FJSVstclang/FCCpx_PROF</code>
	Clang Mode specific	<code>/etc/opt/FJSVstclang/FCCpx_clang_PROF</code>
Native compiler	Clang/Trad Mode common	<code>/etc/opt/FJSVstclang/FCC_PROF</code>
	Clang Mode specific	<code>/etc/opt/FJSVstclang/FCC_clang_PROF</code>

Common compilation profile files for Clang/Trad Mode can specify common compiler options for Clang/Trad Mode. Refer to "[9.1.2.3.1 Correspondences of Compile Options in Clang Mode and Trad Mode](#)" for a description of the common compiler options for Clang/Trad Mode.

Clang Mode specific compilation profile files can specify Clang Mode specific and Clang/Trad Mode common compiler options.

For information about description formats of the compilation profile file and the priority of compiler options, refer to "[2.4 Compilation Profile File](#)".

## 9.1.5 Predefined Macro Names

---

Some of the predefined macros are changed by the compiler option `-std=name`.

The following tables show the values of some predefined macros. To display all predefined macros, specify the compiler options -E and -dM.

Table 9.9 Predefined Macros whose value is changed by the compiler option -std=name

Identifier	The level of language specifications specified with the compiler option -std=name							
	c++98 c++03	gnu++98 gnu++03	c++11	gnu++11	c++14	gnu++14	c++17	gnu++17
__STRICT_ANSI__	1	-	1	-	1	-	1	-
__cplusplus	199711L		201103L		201402L		201703L	
linux	-	1	-	1	-	1	-	1
unix	-	1	-	1	-	1	-	1

-: Undefined

Table 9.10 Predefined Macros whose value is not changed by the compiler option -std=name

Identifier	Value
__GNU_SOURCE	1
__LP64	1
__OPENMP (*1)	201511
__REENTRANT (*2)	1
__ARM_ARCH	8
__CLANG_FUJITSU	1
__DATE__	The date of compilation (in the format of the asctime function)
__ELF__	1
__EXCEPTIONS (*3)	1
__FCC_major__	The major version number of the compiler
__FCC_minor__	The minor version number of the compiler
__FCC_patchlevel__	The patch level of the compiler
__FCC_version__	The string which represents the version of the compiler
__FILE__	Source file name
__GNUC__	4
__GNUC_MINOR__	2
__GNUC_PATCHLEVEL__	1
__GNUG__	4
__LINE__	Line number of source file
__LP64__	1
__OPTIMIZE__ (*4)	1
__PIC__	1 (*5)
	2 (*6)
__PRAGMA_REDEFINE_EXTNAME	1
__PTRDIFF_TYPE__	long int
__SIZE_TYPE__	long unsigned int

Identifier	Value
__STDC__	1
__STDC_HOSTED__	1
__STDCPP_THREADS__	1 (*7)
__TIME__	The time of compilation (in the format of the <code>asctime</code> function)
__WCHAR_TYPE__	unsigned int (*8)
__aarch64__	1
__clang_major__	The major version number of Clang/LLVM
__clang_minor__	The minor version number of Clang/LLVM
__clang_patchlevel__	The patch level of Clang/LLVM
__clang_version__	The string which represents the version of Clang/LLVM
__linux	1
__linux__	1
__pic__	1 (*5)
	2 (*6)
__unix	1
__unix__	1

-: Undefined

\*1) When the compiler option `-Kopenmp` or `-fopenmp` is set

\*2) When the compiler option `-pthread` is set

\*3) When the compiler option `-Nexceptions` or `-fexceptions` is set

\*4) When the compiler option `-O1` or higher is set

\*5) When the compiler option `-fpic` or `-fPIC` is set

\*6) When the compiler option `-fPIC` or `-fPIE` is set

\*7) When the compiler option `-stdlib=libc++` is set

\*8) This value is different from Trad Mode, and unintended behavior may occur when you execute a program linking an object generated in Trad Mode and an object generated in Clang Mode.

## 9.1.6 Procedure of Execution

This section explains the procedure of executing C++ language programs.

### 9.1.6.1 Environment Variable for Execution

Runtime controls can be changed by setting environment variables. "Table 9.11 Environment Variables for Execution (Clang Mode)" shows the environment variables that can be set at runtime.

For details about OpenMP environment variables, see Section "4.3.2.2 Environment Variable for OpenMP Specifications".

Table 9.11 Environment Variables for Execution (Clang Mode)

Environment Variable Name	Operands	Description
FLIB_HOOK_TIME	<i>time</i>	The user-defined function is called at interval of <i>time</i> millisecond(s). <i>time</i> can be a value between 0 and 2147483647. If 0 is specified for <i>time</i> , calling at regular interval is disabled.

Environment Variable Name	Operands	Description
		See Section " <a href="#">8.3 Hook Function</a> ", for information about the hook function.
LD_BIND_NOW	1	Indicates not to use the delay link for symbol resolution of shared library addresses.

### 9.1.6.2 Notes for Execution

When executing programs created on this system, note the following.

- Variable allocation at execution

Programs created by this system allocate the local and private variables in each function to the stack region. When much more space is required for local and private variables in the function, the stack area needs to be extended to the appropriate size.

The stack area of the process can be set by `ulimit` (bash built-in command), etc.

## 9.2 Optimization

---

### 9.2.1 Overview of Optimization

---

Refer to the Section "[3.1 Overview of Optimization](#)" for overview of optimization.



Note

Result of optimization

The optimization result by Clang Mode may be different from that of Trad Mode.

Compilation message

In Clang Mode, the messages are in English only. And, the message format of Clang Mode is different from that of Trad Mode.

### 9.2.2 Using Optimization Functions

---

This section describes techniques for effectively using the optimization functions.

#### 9.2.2.1 Using Optimization Control Line (Pragma Directives)

It may be possible to increase the optimization level using certain `#pragma` directives in the source code.

In Clang Mode, the following two types of pragma directives can be used.

- Fujitsu-specific pragma supported by Clang Mode
- Clang/LLVM pragma supported by Clang Mode

The `-ffj-ocl` option enables the Fujitsu-specific pragma supported by Clang Mode, and the `-ffj-no-ocl` option suppresses its effect.

The Clang/LLVM pragma supported by Clang Mode is always valid.



Note

The result is not guaranteed when using the pragma not listed in this section.

##### 9.2.2.1.1 Types of Optimization Control Lines

The optimization control line of Fujitsu-specific pragma supported by Clang Mode can be specified in the following format. Enables the specified optimization control specifier for just before the loop.

### Description Format

```
#pragma fj Optimization-control-specifier
```

### Insert Position

This is positioned just before the corresponding loop.

Also, the loop line of the Clang/LLVM pragma or another Fujitsu-specific pragma can be interposed between the loop and this pragma.

For Clang/LLVM pragma supported by Clang Mode, the following format can be used. Enables the specified optimization control specifier for just before the loop or the compound statement.

### Description Format

```
#pragma clang Optimization-control-specifier
```

### Insert Position

- In the case of an optimization control specifier that begins with loop

Put it just before the corresponding loop.

Also, the loop line of the Clang/LLVM pragma or another Fujitsu-specific pragma can be interposed between the loop.

- In the case of an optimization control specifier that begins with fp

Put it before all explicit declarations and statements inside the corresponding compound statement.

## 9.2.2.1.2 Optimization Control Specifier

The optimization control specifier that can be specified in Clang Mode are shown in "[Table 9.12 Optimization Control Specifiers that can be specified to optimization control line](#)".



### Note

The optimization control specifier of Trad Mode is replaced with that of Clang Mode.

The list of the optimization control specifier to be replaced is shown in "[Table 9.13 Optimization control specifiers replaced in Clang Mode](#)".

If the optimization control specifier that is not listed in "[Table 9.12 Optimization Control Specifiers that can be specified to optimization control line](#)" is specified, it is ignored.

Table 9.12 Optimization Control Specifiers that can be specified to optimization control line

Optimization Control Specifiers	Format of pragma	Explanation
loop clone var==n(*1)	#pragma fj loop clone var==n	Directs to generate conditional branches on the expression of the specified arguments and to generate loop copies in the conditional blocks.
loop eval_concurrent	#pragma fj loop eval_concurrent	Performs the optimization that improves the parallelism of floating-point arithmetic instructions using the optimization of tree-height-reduction. The optimization performed by -ffj-eval-concurrent option can be specified at the loop unit.
loop eval_noconcurrent	#pragma fj loop eval_noconcurrent	Suppresses the optimization that is performed by -ffj-eval-concurrent option.
loop loop_fission_target	#pragma fj loop loop_fission_target [c1]	Performs automatic loop fission optimization. This specifier can be specified only for innermost loops. The compiler behavior when c1 is written is the same as when not written.

Optimization Control Specifiers	Format of pragma	Explanation
loop loop_fission_threshold <i>n</i>	#pragma fj loop loop_fission_threshold <i>n</i>	Specifies the threshold <i>n</i> to decide the granularity of loop after automatic loop fission. <i>n</i> is an integer value from 1 to 100. This specifier can be specified only for innermost loops.
loop preex	#pragma fj loop preex	Optimizes by evaluating an invariant first. The optimization performed by -ffj-preex option can be specified at the loop unit.
loop nopreex	#pragma fj loop nopreex	Suppresses the optimization of evaluating an invariant first.
loop prefetch	#pragma fj loop prefetch	Performs the automatic prefetch function of the compiler. The optimizations performed by -ffj-prefetch-sequential or -ffj-prefetch-stride option can be specified at the loop unit.
loop noprefetch	#pragma fj loop noprefetch	Suppresses the automatic prefetch function of the compiler.
loop prefetch_sequential [auto soft]	#pragma fj loop prefetch_sequential [auto soft]	Directs that prefetch instructions are created for array data that is accessed sequentially. The optimizations performed by -ffj-prefetch-sequential option can be specified at the loop unit.
loop prefetch_nosequential	#pragma fj loop prefetch_nosequential	Directs that prefetch instructions should not be generated for array data that is accessed sequentially.
loop prefetch_stride	#pragma fj loop prefetch_stride	Directs that prefetch instructions are created for array data accessed with a stride larger than the cache line size used in the loop.
loop prefetch_nostride	#pragma fj loop prefetch_nostride	Directs that prefetch instructions for stride access should not be generated.
loop prefetch_strong	#pragma fj loop prefetch_strong	Directs that the prefetch instructions for the first level cache are to be the strong prefetch. The optimizations performed by -ffj-prefetch-strong option can be specified at the loop unit.
loop prefetch_nostrong	#pragma fj loop prefetch_nostrong	Directs that the prefetch instructions generated for the first level cache will not be strong prefetch.
loop prefetch_strong_L2	#pragma fj loop prefetch_strong_L2	Directs that the prefetch instructions generated for the second level cache are to be strong prefetch. The optimizations performed by -ffj-prefetch-strong-L2 option can be specified at the loop unit.
loop prefetch_nostrong_L2	#pragma fj loop prefetch_nostrong_L2	Directs that the prefetch instructions generated for the second level cache are not to be strong prefetch.
loop swp	#pragma fj loop swp	Performs software pipelining.
loop noswp	#pragma fj loop noswp	Suppresses software pipelining.
loop zfill [ <i>N</i> ]	#pragma fj loop zfill [ <i>N</i> ]	Directs that the zfill optimization to be performed. The optional parameter <i>N</i> is an integer between 1 and 100 that specifies the number of blocks the DC ZVA instruction writes.

Optimization Control Specifiers	Format of pragma	Explanation
<code>loop nozfill</code>	<code>#pragma fj loop nozfill</code>	Directs that the zfill optimization not to be performed.
<code>fp contract(fast)</code>	<code>#pragma clang fp contract(fast)</code>	Performs optimizations using the Floating-Point Multiply-Add/Subtract instructions. The optimizations performed by <code>-ffp-contract=fast</code> option can be specified at the section unit.
<code>fp contract(on)</code>	<code>#pragma clang fp contract(on)</code>	Performs optimizations using the Floating-Point Multiply-Add/Subtract instructions. This pragma is identical to using <code>#pragma STDC FP_CONTRACT(ON)</code> .
<code>fp contract(off)</code>	<code>#pragma clang fp contract(off)</code>	Suppresses optimizations using the Floating-Point Multiply-Add/Subtract instructions.
<code>loop unroll(enable)</code>	<code>#pragma clang loop unroll(enable)</code>	Performs unrolling for the corresponding loop.
<code>loop unroll_count(n)</code>	<code>#pragma clang loop unroll_count(n)</code>	The optimizations performed by <code>-funroll-loops</code> option can be specified at the loop unit. When <code>unroll(full)</code> is set, all statements in the target loop is unrolled.
<code>loop unroll(full)</code>	<code>#pragma clang loop unroll(full)</code>	
<code>loop unroll(disable)</code>	<code>#pragma clang loop unroll(disable)</code>	Suppresses unrolling for the corresponding loop.
<code>loop vectorize(assume_safety)</code>	<code>#pragma clang loop vectorize(assume_safety)</code>	Indicates that the loop has no dependencies on array elements or pointer variables.
<code>loop vectorize(enable)</code>	<code>#pragma clang loop vectorize(enable)</code>	Performs the optimization that uses SIMD Extensions.
<code>loop vectorize_width(n, scalable)</code>	<code>#pragma clang loop vectorize_width(n, scalable)</code>	The optimizations performed by <code>-fvectorize</code> option can be specified at the loop unit. When <code>vectorize_width(n, scalable)</code> is set to use SIMD Extensions through SVE, the SIMD width becomes a multiple of <i>n</i> .
<code>loop vectorize(disable)</code>	<code>#pragma clang loop vectorize(disable)</code>	Suppresses the optimization that uses SIMD Extensions.

\*1) *var* must be declared before it is used.

### loop clone specifier

The `loop clone` specifier instructs to generate conditional branches on the expression of the specified arguments and to generate loop copies in the conditional blocks. If multiple `loop clone` specifiers are specified for the same loop, the order of the generated branches is the order of the corresponding arguments. For the nested loops, the `loop clone` specifier cannot be specified for loops of different nesting levels. This specifier promotes other optimizations, such as full unrolling. Note that the size of the object program and compile time may increase because the clone optimization makes copies of loops. This specifier is effective when the `-O3` option is set.

Variable *var* is a variable of type `int`. There are some cases that cannot be specified like follows.

- Structure member variable
- Union member variable
- Class member variable
- Reference argument
- Array elements
- Threadprivate variable



- Variable with scope resolution operator

An integer value from -2147483647 to 2147483647 can be specified as constant values *n*.



## Example

### Specification of the optimization control line

- Example 1: Specifying only one loop clone specifier

```
#pragma fj loop clone N==10
for(i = 0; i < N; ++i) {
    a[i] = i;
}
```

[Optimized pseudo-code]

```
if(N == 10) {
    for(i = 0; i < N; ++i) {
        a[i] = i;
    }
} else {
    for(i = 0; i < N; ++i) {
        a[i] = i;
    }
}
```

Loop is copied with "if" statements in the order of the specification.

- Example 2: Specifying multiple loop clone specifiers for the same loop

```
#pragma fj loop clone N==10
#pragma fj loop clone M==20
for(i = 0; i < N; i++) {
    a[i] = M;
}
```

[Optimized pseudo-code]

```
if(N == 10) {
    for(i = 0; i < N; ++i) {
        a[i] = M;
    }
} else if (M == 20) {
    for(i = 0; i < N; ++i) {
        a[i] = M;
    }
} else {
    for(i = 0; i < N; ++i) {
        a[i] = M;
    }
}
```

Loop is copied with "if" statements in the order of the multiple specifications.

- Example 3: Using a flag to specify loop clone specifier

```
int flag = (start == 0 && end == 10) ? 1 : 0;
#pragma fj loop clone flag==1
for(i = start; i < end; ++i) {
    a[i] = i;
}
```

[Optimized pseudo-code]

```

int flag = (start == 0 && end == 10) ? 1 : 0;
if(flag == 1) {
    for(i = start; i < end; ++i) {
        a[i] = i;
    }
} else {
    for(i = start; i < end; ++i) {
        a[i] = i;
    }
}

```

Loop is copied with "if" statements in the order of the specifications.

- Example 4: Incorrect specification of the loop clone specifier

```

#pragma fj loop clone M==10
for(i = 0; i < M; ++i) {
#pragma fj loop clone N==20
    for(j = 0; j < N; ++j) {
        a[i][j] = 0;
    }
}

```

For the nested loops, the loop clone specifier cannot be specified for loops of different nesting levels.

### loop eval\_concurrent specifier

The loop eval\_concurrent specifier instructs to perform the optimization that improves the parallelism of floating-point arithmetic instructions using the optimization of tree-height-reduction.

#### Example

Specification of the optimization control line

```

#pragma fj loop eval_concurrent
for(i = 0; i < n; i++) {
    x[i] = a[i] * b[i] + c[i] * d[i] + e[i] * f[i] + g[i] * h[i];
}

```

### loop eval\_noconcurrent specifier

The loop eval\_noconcurrent specifier instructs to suppress the optimization that improves the parallelism of floating-point arithmetic instructions using the optimization of tree-height-reduction.

#### Example

Specification of the optimization control line

```

#pragma fj loop eval_noconcurrent
for(i = 0; i < n; i++) {
    x[i] = a[i] * b[i] + c[i] * d[i] + e[i] * f[i] + g[i] * h[i];
}

```

### loop loop\_fission\_target specifier

The loop loop\_fission\_target specifier instructs to perform automatic loop fission for target loop.

This specifier can be specified only for innermost loops.

The compiler behavior when `cl` is written after the specifier is the same as when not written.

## Example

Specification of the optimization control line

```
#pragma fj loop loop_fission_target
for (i = 0; i < n; i++) {
    ...
}
```

### **loop loop\_fission\_threshold specifier**

The `loop loop_fission_threshold` specifier specifies the granularity of loop after automatic loop fission by the threshold `n` written after the specifier. `n` is an integer value from 1 to 100.

This specifier can be specified only for innermost loops.

## Example

Specification of the optimization control line

```
#pragma fj loop loop_fission_target
#pragma fj loop loop_fission_threshold 10
for (i = 0; i < n; i++) {
    ...
}
```

### **loop preex specifier**

The `loop preex` specifier instructs to optimize of evaluating the invariant first.

## Example

Specification of the optimization control line

```
#pragma fj loop preex
for(i = 0; i < n; i++) {
    if(m[i] != 0) {
        a[i] = a[i] / b[k];
    }
}
```

[Optimized pseudo-code]

```
t = 1 / b[k];
for(i = 0; i < n; i++) {
    if(m[i] != 0) {
        a[i] = a[i] * t;
    }
}
```

The optimization of evaluating the invariant first is performed in the target loop.

### **loop nopreex specifier**

The `loop nopreex` specifier instructs to suppress the optimization of evaluating the invariant first.

## Example

Specification of the optimization control line

```
#pragma fj loop nopreex
for(i = 0; i < n; i++) {
    if(m[i] != 0) {
        a[i] = a[i] / b[k];
    }
}
```

The optimization of evaluating the invariant first is suppressed in the target loop.

## loop prefetch specifier

The `loop prefetch` specifier instructs to perform the automatic prefetch function. This function inserts the prefetch instruction in the optimal position to execute, as determined by the compiler.

## Example

Specification of the optimization control line

```
#pragma fj loop prefetch
for(i = 0; i < 10; i=i+m) {
    a[i] = b[i];
}
```

## loop noprefetch specifier

The `loop noprefetch` specifier suppresses the automatic prefetch function.

## Example

Specification of the optimization control line

```
#pragma fj loop noprefetch
for(i = 0; i < 10; i=i+m) {
    a[i] = b[i];
}
```

The automatic prefetch function is suppressed in the loop.

## loop prefetch\_sequential specifier

The `loop prefetch_sequential` specifier directs that prefetch instructions are created for array data that is accessed sequentially.

When `prefetch_sequential auto` specifier is used, the compiler automatically selects whether to use hardware-prefetch or to create prefetch instructions for array data that is accessed sequentially in the loop.

When `prefetch_sequential soft` specifier is used, the compiler does not use hardware-prefetch, but rather creates prefetch instructions for array data that is accessed sequentially in the loop.

If neither `auto` nor `soft` is specified, `auto` is set by default.

## Example

Specification of the optimization control line

- Example 1:

```
#pragma fj loop prefetch_sequential auto
for(i = 0; i < n; i++) {
    a1[i] = a2[i] + a3[i] + a4[i] + a5[i]
        + a6[i] + a7[i] + a8[i] + a9[i] + a10[i]
        + a11[i] + a12[i] + a13[i] + a14[i] + a15[i]
        + a16[i] + a17[i];
}
```

Hardware prefetch is used for the target loop rather than generating prefetch instructions.

- Example 2:

```
#pragma fj loop prefetch_sequential soft
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i];
}
```

Prefetch instructions are generated for the target loop rather than using hardware-prefetch.

---

### loop prefetch\_nosequential specifier

The loop `prefetch_nosequential` specifier directs that prefetch instructions should not be generated for array data that is accessed sequentially.

#### Example

---

Specification of the optimization control line

```
#pragma fj loop prefetch_nosequential
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i];
}
```

Prefetch instructions are not generated for array data that is accessed sequentially in the target loop.

---

### loop prefetch\_stride specifier

The loop `prefetch_stride` specifier directs that prefetch instructions are created for array data accessed with a stride larger than the cache line size used in the loop.

#### Example

---

Specification of the optimization control line

```
#pragma fj loop prefetch_stride
for(i = 0; i < n; i=i+m) {
    a[i] = b[i];
}
```

Prefetch instructions are generated for the target loop.

---

### loop prefetch\_nostride specifier

The loop `prefetch_nostride` specifier directs that prefetch instructions are not created for array data accessed with a stride larger than the cache line size used in the loop.

## Example

Specification of the optimization control line

```
#pragma fj loop prefetch_nostride
for(i = 0; i < n; i=i+m) {
    a[i] = b[i];
}
```

Prefetch instructions are not generated for the target loop.

## **loop prefetch\_strong specifier**

The loop prefetch\_strong specifier directs that prefetch instructions for the first level cache are to be the strong prefetch.

## Example

Specification of the optimization control line

```
#pragma fj loop prefetch_strong
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i];
}
```

Prefetch instructions for the first level cache generated for the target loop will be strong prefetch.

## **loop prefetch\_nostrong specifier**

The loop prefetch\_nostrong specifier directs that prefetch instructions generated for the first level cache will not be strong prefetch.

## Example

Specification of the optimization control line

```
#pragma fj loop prefetch_nostrong
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i];
}
```

Prefetch instructions for the first level cache generated for the target loop will not be strong prefetch.

## **loop prefetch\_strong\_L2 specifier**

The loop prefetch\_strong\_L2 specifier directs that prefetch instructions for the second level cache are to be the strong prefetch.

## Example

Specification of the optimization control line

```
#pragma fj loop prefetch_strong_L2
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i];
}
```

Prefetch instructions for the second level cache generated for the target loop will be strong prefetch.

### loop prefetch\_nostrong\_L2 specifier

The `loop prefetch_nostrong_L2` specifier directs that prefetch instructions generated for the first level cache will not be strong prefetch.

#### Example

Specification of the optimization control line

```
#pragma fj loop prefetch_nostrong_L2
for(i = 0; i < n; i++) {
    a[i] = a[i] + b[i];
}
```

Prefetch instructions for the second level cache generated for the target loop will not be strong prefetch.

### loop swp specifier

The `loop swp` specifier performs software pipelining.

#### Example

Specification of the optimization control line

```
#pragma fj loop swp
for (i = 0; i < n; i++) {
    a[i] = b[i]/c[i];
}
```

### loop noswp specifier

The `loop noswp` specifier suppresses software pipelining.

#### Example

Specification of the optimization control line

```
#pragma fj loop noswp
for (i = 0; i < n; i++) {
    a[i] = b[i]/c[i];
}
```

### loop zfill specifier

The `loop zfill` specifier directs to apply the `zfill` optimization. The `zfill` optimization speeds up write operations for array data that is only written in a loop, by using an instruction that allocates space on the cache for writing without loading data from the memory. The `zfill` optimization works on the data  $N$  blocks ahead of the address pointed to by the target store instruction where one block is 256 byte-long and  $N$  is an integer value between 1 and 100. If a value is not specified for  $N$ , the compiler will automatically determine a value.

For details about `zfill`, see "[3.3.5 zfill](#)".

#### Example

Specification of the optimization control line

- Example 1

```
#pragma fj loop zfill
for(i = 0; i < n; i++) {
    ...
}
```

This specifies to improve the cache utilization for the data on certain blocks ahead, and the number is determined by the compiler.

- Example 2

```
#pragma fj loop zfill 1
for(i = 0; i < n; i++) {
    ...
}
```

This specifies to improve the cache utilization for the data on 1 block ahead.

---

### loop nozfill specifier

The `loop nozfill` specifier directs not to apply the `zfill` optimization.

#### Example

---

Specification of the optimization control line

```
#pragma fj loop nozfill
for(i = 0; i < n; i++) {
    ...
}
```

This specifies that the `zfill` optimization will not be performed.

---

### fp contract(fast) specifier

The `fp contract(fast)` specifier instructs to apply optimizations using the Floating-Point Multiply-Add/Subtract instructions. However, using these instructions may cause calculation errors in the result due to rounding errors.

#### Example

---

Specification of the optimization control line

```
for(i = 0; i < n; i++) {
    #pragma clang fp contract(fast)
    a[i] = a[i] + b[i] * c[i];
}
```

The optimization that uses the Floating-Point Multiply-Add/Subtract instructions is performed for the target compound statement.

---

### fp contract(on) specifier

The `fp contract(on)` specifier instructs to apply optimizations using the Floating-Point Multiply-Add/Subtract instructions. It is performed as if the `#pragma STDC FP_CONTRACT ON` (language standard) was specified. However, using these instructions may cause calculation errors in the result due to rounding errors.



## Example

Specification of the optimization control line

```
for(i = 0; i < n; i++) {
    #pragma clang fp contract(on)
    a[i] = a[i] + b[i] * c[i];
}
```

The optimization that uses the Floating-Point Multiply-Add/Subtract instructions is performed for the target compound statement.

### **fp contract(off) specifier**

The `fp contract(off)` specifier instructs to suppress optimizations using the Floating-Point Multiply-Add/Subtract instructions.

## Example

Specification of the optimization control line

```
for(i = 0; i < n; i++) {
    #pragma clang fp contract(off)
    a[i] = a[i] + b[i] * c[i];
}
```

The optimization that uses the Floating-Point Multiply-Add/Subtract instructions is not performed for the target compound statement.

### **loop unroll(enable) specifier**

The `loop unroll(enable)` specifier instructs the compiler to perform the optimization of loop unrolling for the corresponding loop. The number to be unrolled is determined by the compiler automatically.

Note that the `loop unroll(enable)` specifier targets only the loop specified immediately after.

## Example

Specification of the optimization control line

```
#pragma clang loop unroll(enable)
for(i = 0; i < n; i++) {
    ...
}
```

### **loop unroll(full) specifier**

The `loop unroll(full)` specifier instructs the compiler to perform the optimization of loop unrolling for the corresponding loop.

When the `loop unroll(full)` specifier is used, the statement is unrolled up to the number of the iteration in the specified loop.

If the number of the iteration is unknown at compilation, no optimization of loop unrolling is performed.

Note that the `loop unroll(full)` specifier targets only the loop specified immediately after.

## Example

Specification of the optimization control line

```
#pragma clang loop unroll(full)
for(i = 0; i < 8; i++) {
```

```
} ...
```

All statements in the target loop is unrolled.

### loop unroll\_count specifier

The number from 2 to 100 following the specifier instructs the upper bound on the number of loops to be unrolled.

Cannot omit that upper bound.

Note that the loop unroll\_count specifier targets only the loop specified immediately after.

#### Example

Specification of the optimization control line

```
#pragma clang loop unroll_count(8)
for(i = 0; i < n; i++) {
    ...
}
```

The "statement" is unrolled 8 times in the target loop.

### loop unroll(disable) specifier

The loop unroll(disable) specifier instructs the compiler to suppress unrolling for the corresponding loop.

Note that the loop unroll(disable) specifier targets only the loop specified immediately after.

#### Example

Specification of the optimization control line

```
#pragma clang loop unroll(disable)
for(i = 0; i < n; i++) {
    ...
}
```

Loop unrolling is suppressed in the target loop.

### loop vectorize(assume\_safety) specifier

The loop vectorize(assume\_safety) specifier indicates that there is no dependency on an element of the array or the pointer variable in the loop. This specification causes SIMD Extensions to be performed on loops that cannot be optimized because the data dependency is unknown. Depending on the operation type and the loop structure, it may not be subject to SIMD.

#### Note

If the loop vectorize(assume\_safety) specifier is incorrectly specified in a loop for the element of the dependent array or the pointer variable, the result will be unpredictable.

## Example

Specification of the optimization control line

```
#pragma clang loop vectorize(assume_safety)
for (int i = 0; i < count; i++) {
    a[index[i]] = a[index[i]] + 1;
}
```

### loop vectorize(enable) specifier

The `loop vectorize(enable)` specifier instructs to perform the optimization that uses SIMD Extensions. However, SIMD Extensions may not be used depending on the type of operation and the loop structure.

The compiler automatically determines a suitable SIMD width.

## Example

Specification of the optimization control line

```
double a[10], b[10];
...
#pragma clang loop vectorize(enable)
for(i = 0; i < 10; i++) {
    a[i] = a[i] + b[i];
}
```

The `simd` specifier is in effect for the data in the loop.

### loop vectorize\_width(n, scalable) specifier

The `loop vectorize(n, scalable)` specifier instructs to perform the optimization that uses SIMD Extensions through SVE. The SVE vector register is not considered to be a specific size. However, SIMD Extensions may not be used depending on the type of operation and the loop structure.

This specifier is effective only when the `-msve-vector-bits=scalable` option is effective.

The number of elements of data which one SIMD instruction processes if the size of the SVE vector register is 128-bit (the minimum) is specified by  $n$ . The actual number of elements of data which one SIMD instruction processes, which is called a SIMD width, is decided at execution time depending on the SVE vector register size on the CPU architecture. For example, if the SVE vector register size is 512-bit, the SIMD width is  $(512/128)*n = 4*n$ . If one SIMD instruction cannot process the SIMD width at once, the instruction is used more than once. You can specify 2 or 4 for  $n$ .

## Example

Specification of the optimization control line

```
float a[100];
double b[100];
...
#pragma clang loop vectorize_width(4, scalable)
for(i = 0; i < 100; i++) {
    a[i] += 1.0;
    b[i] += 1.0;
}
```

If the SVE register size is 512-bit for example, data from `a[0]` to `a[15]`, from `b[0]` to `b[7]`, and from `b[8]` to `b[15]` are each processed in one SIMD instruction.

## loop vectorize(disable) specifier

The `loop vectorize(disable)` specifier instructs to suppress the optimization that uses SIMD Extensions.



### Example

Specification of the optimization control line

```
double a[10], b[10];
...
#pragma clang loop vectorize(disable)
for(i = 0; i < 10; i++) {
    a[i] = a[i] + b[i];
}
```

## 9.2.2.1.3 Notes on Optimization Control Lines and Optimization Control Specifiers


This section describes notes on the optimization control lines and the optimization control specifiers.

Differences in behaviors depending on types of optimization control lines


Depending on types of optimization control lines described in "9.2.2.1.1 Types of Optimization Control Lines", the compiler behaves differently in some cases.

- When multiple optimization control lines of the same type are specified at the same location

Specified optimization control specifiers		Compiler behavior	
		Clang/LLVM optimization control lines	Fujitsu-specific optimization control lines
Same optimization control specifiers	loop clone specifier	(No corresponding specifier)	It can be compiled. All loop clone specifiers are enabled.
	Others	It causes a compilation error.	It can be compiled.
Conflicting optimization control specifiers, such as applying and suppressing a certain optimization	Following optimization control specifiers - fp contract(fast) - fp contract(on) - fp contract(off)	It can be compiled. The last optimization control line is enabled.	The last optimization control line is enabled.
	Others Example ..... loop unroll(enable) specifier and loop unroll(disable) specifier .....	It causes a compilation error.	
Different optimization control specifiers for an optimization Example ..... loop unroll(full) specifier and loop unroll_count(n) specifier .....		It causes a compilation error.	

Specified optimization control specifiers	Compiler behavior	
	Clang/LLVM optimization control lines	Fujitsu-specific optimization control lines
Optimization control specifiers with different parameter values  <b>Example</b> ..... <code>loop unroll_count(n) specifier</code> .....	It causes a compilation error.	

- When there is an error in an optimization control line

Type of the error	Compiler behavior	
	Clang/LLVM optimization control lines	Fujitsu-specific optimization control lines
Error in a specifier in an optimization control line or an element which follows a specifier  <b>Example</b> ..... - Example 1 <pre>#pragma loop unrecognized_specifier</pre> - Example 2 <pre>#pragma loop prefetch_sequential unrecognized_parameter</pre> .....	It causes a compilation error.	It can be compiled.  A warning message is output and the optimization control line is ignored.

#### Correspondences of optimization control specifiers in Clang Mode and Trad Mode

In Clang Mode and Trad Mode, the optimization control specifiers whose type of optimization is equivalent but name is different are described by the loop line format in "Table 9.13 Optimization control specifiers replaced in Clang Mode".

- When the optimization control specifiers of Trad Mode described in the table is used, that of Clang Mode is assumed to be used.
- An optimization result may be different even if the type of optimization is equivalent.
- If an optimization control specifier of Trad Mode which is not described in the table is used, it is ignored and a warning message is output.

Table 9.13 Optimization control specifiers replaced in Clang Mode

Trad Mode	Clang Mode
<code>#pragma loop clone var==n</code>	<code>#pragma fj loop clone var==n</code> <sup>[a]</sup>
<code>#pragma loop eval_concurrent</code>	<code>#pragma fj loop eval_concurrent</code>
<code>#pragma loop eval_noconcurrent</code>	<code>#pragma fj loop eval_noconcurrent</code>
<code>#pragma loop loop_fission_target [cl]</code>	<code>#pragma fj loop loop_fission_target [cl]</code>
<code>#pragma loop loop_fission_threshold n</code>	<code>#pragma fj loop loop_fission_threshold n</code>
<code>#pragma loop preex</code>	<code>#pragma fj loop preex</code>
<code>#pragma loop nopreex</code>	<code>#pragma fj loop nopreex</code>

Trad Mode	Clang Mode
<code>#pragma loop prefetch</code>	<code>#pragma fj loop prefetch</code>
<code>#pragma loop noprefetch</code>	<code>#pragma fj loop noprefetch</code>
<code>#pragma loop prefetch_sequential [auto soft]</code>	<code>#pragma fj loop prefetch_sequential [auto soft]</code>
<code>#pragma loop prefetch_nosequential</code>	<code>#pragma fj loop prefetch_nosequential</code>
<code>#pragma loop prefetch_stride</code>	<code>#pragma fj loop prefetch_stride</code>
<code>#pragma loop prefetch_nostride</code>	<code>#pragma fj loop prefetch_nostride</code>
<code>#pragma loop prefetch_strong</code>	<code>#pragma fj loop prefetch_strong</code>
<code>#pragma loop prefetch_nostrong</code>	<code>#pragma fj loop prefetch_nostrong</code>
<code>#pragma loop prefetch_strong_L2</code>	<code>#pragma fj loop prefetch_strong_L2</code>
<code>#pramga loop prefetch_nostrong_L2</code>	<code>#pragma fj loop prefetch_nostrong_L2</code>
<code>#pragma loop swp</code>	<code>#pragma fj loop swp</code>
<code>#pragma loop noswp</code>	<code>#pragma fj loop noswp</code>
<code>#pragma loop zfill [N]</code>	<code>#pragma fj loop zfill [N]</code>
<code>#pragma loop nozfill</code>	<code>#pragma fj loop nozfill</code>
<code>#pragma[ fj] loop simd <sup>[b]</sup></code>	<code>#pragma clang loop vectorize(enable)</code>
<code>#pragma[ fj] loop nosimd</code>	<code>#pragma clang loop vectorize(disable)</code>
<code>#pragma[ fj] loop unroll</code>	<code>#pragma clang loop unroll(enable)</code>
<code>#pragma[ fj] loop unroll n</code>	<code>#pragma clang loop unroll_count(n)</code>
<code>#pragma[ fj] loop unroll "full"</code>	<code>#pragma clang loop unroll(full)</code>
<code>#pragma[ fj] loop nounroll</code>	<code>#pragma clang loop unroll(disable)</code>

[a] Trad Mode and Clang Mode have different notations, types, and values for their arguments.

[b] No argument

## 9.2.3 SIMD

Refer to the Section "[3.2.7.1 Normal SIMD](#)" for normal SIMD.

### 9.2.3.1 Math Functions that SIMD Extensions can be applied to

The following shows the math functions that SIMD extensions can be applied to when `-fvectorize` option is effective.

```
ceil, ceilf, copysign, copysignf, fabs, fabsf, floor, floorf, fmax, fmaxf, fmin, fminf, round, roundf,
trunc, truncf
```

The following shows the math functions that SIMD extensions can be applied to when `-fvectorize` option and `-fbuiltin` option are effective.

```
abs, cimag, cimagf, creal, crealf, fma, fmaf, nearbyint, nearbyintf, rint, rintf
```

The following shows the math functions that SIMD extensions can be applied to when `-fvectorize` option, `-fbuiltin` option, and `-ffast-math` option are effective.

```
cabs, cabsf, sqrt, sqrtf
```

When `-fvectorize` option, `-fbuiltin` option, and `-ffj-ilfunc` option are effective, SIMD extensions can be applied to the math functions that are subject to the inline expansion.

See also the description of the `-ffj-ilfunc` options in the section "[9.1.2.2.3 Options for Optimization](#)" for information on math functions that are subject to the inline expansion.

## 9.3 Multiprocessing

---

### 9.3.1 Overview of Multiprocessing

---

What is Multiprocessing?

In this document, multiprocessing means that one program is executed on two or more CPUs that can work independently at the same time. As used here, it does not mean executing two or more programs simultaneously.

Effect of Multiprocessing

The effect of multiprocessing is to reduce elapsed execution time by using two or more CPUs simultaneously. For instance, if a loop can be executed in parallel by dividing it, the execution time of this loop may be cut in half.

Although the elapsed time will usually decrease with multiprocessing, the total CPU time required to execute the program may increase. This is because the total CPU time is at least as large as the CPU time when the program is executed on a single processor, and the overhead time for multiprocessing may increase the total CPU time.

The cause of the overhead of parallel execution is the parallel region that is one executed by a team consisting of more than one thread. One parallelization loop is generated for one parallel region. The cost of generating the parallel region is the overhead of parallel execution.

The performance of a multiprocessing program is usually evaluated by the reduction of the elapsed time. The total CPU time of a program using multiprocessing is larger than that required by serial processing. This is because of the overhead time required by multiprocessing.

Requirements for Effective Multiprocessing

A computer environment that can use two or more CPUs simultaneously is necessary to reduce elapsed time by multiprocessing. A multiprocessing program can be executed on hardware with only a single CPU; however, the elapsed time will not be less than the execution time for a comparable program written without multiprocessing features. Moreover, even if the program is executed on hardware with two or more CPUs, shortening the elapsed time is difficult when other jobs are executing on the same CPUs. The reason for this is that the probability that two or more CPUs can be allocated at the same time decreases.

To achieve effective multiprocessing, it is necessary to run a multiprocessing program on hardware with multiple CPUs and within a system environment that has room for multiprocessing.

To reduce the relative overhead of multiprocessing, it is necessary that a loop have many iterations and many statements in its body.

### 9.3.2 Parallelization by OpenMP Specification

---

See Section "[4.3 Parallelization by OpenMP Specification](#)" for the multiprocessing.

## 9.4 Emitting Information

---

### 9.4.1 Emitting Information at Compilation

---

This section explains the information emitted by this system at compilation.

### 9.4.1.1 Header

When either the `-ffj-lst`, `-ffj-lst-out=file` or `-ffj-src` compilation option is specified, a header is emitted for each type of information.

Header Format

Fujitsu C/C++ Version <i>version</i> <i>date</i>	
<i>version</i>	Language processing system version.
<i>date</i>	Compilation date and time in the format of the <code>asctime</code> function.

### 9.4.1.2 Source List

When the program is compiled, the source list is put out according to the following options.

- `-ffj-lst[={plt}`

Specifies to output source list to file(s).

- `-ffj-lst-out=file`

Specifies the filename to output source list.

When this option is specified, the `-ffj-lst=p` option is also effective.

- `-ffj-src`

Specifies to output source list to the standard output.

Note that when the `-ffj-lst` or `-ffj-lst-out=file` option is specified, the source list is output to the *file*.

The source list includes symbols indicating optimization if they are performed.

#### 9.4.1.2.1 Output Format

The source list is emitted as shown in the following format.

Source List Format

```

Compilation information
  Current directory : directory-name
  Source file      : source-file-name
(line-no.)(optimize)
nnnnnnnn i mmmmv source
...
      <<< Loop-information Start >>>
      <<< details optimization information
      <<< Loop-information End >>>
nnnnnnnn i mmmmv source
...

```

<i>directory-name</i>	The name of directory where source file is stored
<i>source-file-name</i>	Source file name
<i>nnnnnnnn</i>	Line number of source (variable length)
<i>i</i>	Symbols for inline expansion
<i>mmmm</i>	The number of loop unrolling (variable length)
<i>v</i>	Symbols for using SIMD Extensions
<i>source</i>	Source line
<<< Loop-information Start >>> <sup>[a]</sup>	Details of optimization and parallelization information header
<i>details optimization information</i> <sup>[a]</sup>	Details of optimization and parallelization information which has been performed to the next statement.



<<< Loop-information End >>> <sup>[a]</sup>	Details of optimization and parallelization information footer
---	--

[a] Emitted only if the -ffj-lst=t option is effective.

### 9.4.1.2.2 Information Included in Source List

#### Line number of source

Line number of source is emitted.

#### Symbols for inline expansion

Symbols for inline expansion are emitted at the line including function call.

Symbol	Description
i	Indicates inline expansion is performed
Blank	Indicates inline expansion is not performed

#### The number of loop unrolling

If loop unrolling is performed, multiplicity of loop unrolling is emitted.

If the loop full unrolled, "f" is emitted instead of multiplicity of loop unrolling.

Otherwise, blank is emitted.

#### Symbols for using SIMD Extensions

Symbols for the optimization that uses SIMD Extensions are emitted at the line including loop control statements.

The loop control statement is "for" statement, "while" statement, "do-while" statement, and "if-goto" statement.

Symbol	Description
v	Indicates all statements in this loop have been optimized using SIMD Extensions.
m	Indicates that only some of the statements in this loop have been optimized using SIMD Extensions.
s	Indicates none of the statements in this loop have been optimized using SIMD Extensions.
Blank	Indicates none of the statements in this loop are targeted for optimization using SIMD Extensions.

#### Details Optimization Information

When the -ffj-lst=t option is specified, the source list containing details of optimization and parallelization information which has been performed are written to file(s).

The following optimization information are put just before each loop.

##### Optimization Information In Loop Units

- FISSION(num: *N*)

It means that loop was split.

- *N* is the number of loops after fission.

- SOFTWARE PIPELINING

It means that loop was pipelined.

- SIMD

It means that SIMD instructions were generated for the loop. Displayed as one of the followings:

- SIMD(VL: *length*[,*length*]... Interleave: *num*[,*num*])

A SIMD instruction treats *length* elements of array. When loop fission is applied to the loop and the values of *length* for the each loop are different, two or more lengths are displayed.

*num* displays the number of instructions expanded in the loop after SIMD to make memory access after SIMD more efficient.

- SIMD(VL: AGNOSTIC; VL: *length*[,*length*]... in 128-bit Interleave: *num*[,*num*])

It means that the vector register of SVE is SIMD without regard to a specific size. The *length* represents the number of elements of the array processed in 1 SIMD instruction, assuming that the vector register size of the SVE is 128 bits. If the *length* is different for each loop that is loop divided, multiple *lengths* are displayed.

*num* displays the number of instructions expanded in the loop after SIMD to make memory access after SIMD more efficient.

- PATTERN MATCHING(matmul)

It means that loop was changed to library function call (matmul).

- FULL UNROLLING

It means that loop was fully unrolled.

- CLONE

It means that loop was loop cloned.

#### Information That Relates To The Prefetch

##### Prefetch instruction

- PREFETCH(SOFT) : *N*

It indicates the number of all prefetch instructions that exist in the loop.

In addition, the number and the array name of the prefetch instruction of each access type of the prefetch are shown in the form of the following.

```
access type: N
array name: N, ...
```

##### access type

- SEQUENTIAL : *N*

It indicates the number of prefetch instructions for array data accessed sequentially within a loop.

- STRIDE : *N*

It indicates the number of prefetch instructions for array data that is accessed with a stride larger than the cache line size used in the loop.

##### array name

- array name: *N*,...

It indicates the number of prefetch instructions for the array name. The array which was generated by compiler is shown as "(unknown)".

#### Information That Relates To The Register

- SPILLS :

It indicates the number of register save/restore instructions in the innermost loop for each register type.

- GENERAL : SPILL *N* FILL *N*

It indicates the number of save/restore instructions to the memory of the general-purpose register.

- SIMD&FP : SPILL *N* FILL *N*

It indicates the number of instructions for saving/restoring SIMD and floating-point registers to memory.

- SCALABLE : SPILL N FILL N

It indicates the number of instructions for saving/restoring extended register to memory.

- PREDICATE : SPILL N FILL N

It indicates the number of instructions for saving/restoring the predicate register to memory.

### 9.4.1.2.3 Example of Source Listing

Example of source listing is shown in following list.

#### Example

##### Example 1: Source listing

```
Compilation information
Current directory : directory-name
Source file : source-file-name
(line-no.)(optimize)
 1          #include <stdio.h>
 2
 3          float sub(int i);
 4
 5          int main() {
 6              int i;
 7              float a[1000];
 8
 9          v   for(i = 0; i < 1000; i++) {
10              a[i] = i;
11          }
12
13              printf("%lf\n", a[0]);
14              i = 0;
15              _loop:
16              if(i < 1000) {
17                  goto _next;
18              }
19              a[i] = i;
20              i++;
21              goto _loop;
22              _next:
23
24              printf("%lf\n", a[0]);
25
26          v   for(i = 0; i < 1000; i++) {
27      i       a[i] = sub(i);
28          }
29          f   for(i = 0; i < 2; i++) {
30              a[i] = i;
31          }
32              printf("%lf\n", a[0]);
33              return 0;
34          }
35
36          float sub(int j) {
37              return (float)j;
38          }
```

In this sample, following information can be read.

- The statements included in the "for"-statement at line-9 are optimized using SIMD Extensions.
- The statements included in the "for"-statement at line-26 are optimized using SIMD Extensions.

- The function called at line-27 is optimized with inline expansion.
- The statements included in the "for"-statement at line-29 are full unrolled.

## Example

### Example 2: Source listing

```

Compilation information
Current directory : directory-name
Source file : source-file-name
(line-no.)(optimize)
 1      #include <stdio.h>
 2
 3      float sub(int i);
 4
 5      int main() {
 6          int i;
 7          float a[10000];
 8
 9          <<< Loop-information Start >>>
10          <<< [OPTIMIZATION]
11          <<< SIMD(VL: 4 Interleave: 1)
12          <<< Loop-information End >>>
13      v   for(i = 0; i < 10000; i++) {
14          a[i] = i;
15      }
16
17      printf("%lf\n", a[0]);
18
19      i = 0;
20      _loop:
21      if(i < 10000) {
22          goto _next;
23      }
24      a[i] = i;
25      i++;
26      goto _loop;
27      _next:
28      printf("%lf\n", a[0]);
29      <<< Loop-information Start >>>
30      <<< [OPTIMIZATION]
31      <<< SIMD(VL: 4 Interleave: 1)
32      <<< Loop-information End >>>
33      v   for(i = 0; i < 10000; i++) {
34      i   a[i] = sub(i);
35      }
36      <<< Loop-information Start >>>
37      <<< [OPTIMIZATION]
38      <<< FULL UNROLLING
39      <<< Loop-information End >>>
40      f   for(i = 0; i < 8; i++) {
41          a[i] = i;
42      }
43      printf("%lf\n", a[0]);
44      return 0;
45      }
46
47      float sub(int j) {
48          return(float)j;
49      }

```

In this sample, following information can be read.

- The statements included in the "for"-statement at line-9 are optimized using SIMD Extensions.
- The loop by "for"-statement at line-25 is optimized using SIMD Extensions.
- The function called at line-26 is optimized with inline expansion.
- The statements included in the "for"-statement at line-28 are full unrolled.

---

#### 9.4.1.2.4 Notes on Information at Compilation (Source List)

The -ffj-lst=p option, -ffj-lst=t option and -ffj-src option could output wrong compilation information (optimization information) in the following cases.

- When a function is inline expanded, different optimizations could be applied to each line. In this case, the compiler could output information incorrectly as follows.
  - Output messages redundantly.
  - Output optimization messages inconsistent with compilation information (optimization information).
  - Output no compilation information (optimization information).

When a function in a loop is expanded inline, the prefetch number of the function is recorded in the compilation time information (optimization information) of the loop.

- The compiler applies multiple optimizations such as SIMD Extensions, loop fusion, and loop distribution to one loop. In this case, the compiler could output information incorrectly as follows.
  - Output optimization information on line number to an incorrect line.
  - Output compilation information (optimization information) inconsistent with optimization messages.
- Loop unswitching optimization creates a loop in each "TRUE" and "FALSE" blocks of the "if" statement, but the compiler outputs optimization message for only one of the loops. In addition, optimization information on line number may not be output.
- When multiple loops are written in the same line in a program, the compiler outputs optimization information for only one of the loops. Please make sure to write one loop in one line in order to output optimization information.
- Detailed optimization information for a loop which consists of "if-goto" statements is not output. Optimization information on line number could be output to an incorrect line.
- When link time optimization is applied, the following phenomena may occur.
  - The optimization which is different from the optimization displayed in compilation information may be applied at runtime.
  - Some compilation information (Optimization Information) will not be output.
- Since the interleave function uses a part of the SIMD function, the SIMD Extensions symbol "v" is displayed even if only the interleave function is operated without SIMD Extensions. In this case, "VL: length" is not displayed as detailed optimization information, so it can be judged that SIMD Extensions are not operated.

#### Example

- When the SIMD Extensions and interleave function are operated

```
          <<< Loop-information Start >>>
          <<< [OPTIMIZATION]
          <<<   SIMD(VL: 4 Interleave: 1)
          <<< Loop-information End >>>
9         v   for(i = 0; i < 10000; i++) {
10             a[i] = i;
11         }
```

- When the interleave function is operated without SIMD Extensions

```

          <<< Loop-information Start >>>
          <<< [OPTIMIZATION]
          <<< SIMD(Interleave: 1)
          <<< Loop-information End >>>
9         v   for(i = 0; i < 10000; i++) {
10                a[i] = i;
11                }

```

## 9.5 Language Specifications

---

### 9.5.1 Supported Language Standard

---

This section describes support status of the language standard.

Clang Mode supports the following language standard.

- C++03
- C++11
- C++14
- C++17

#### C++17 Standard

The following features are not available in the C++ standard library libc++.

- Safe conversions in `unique_ptr<T[]>`.
- Removing `auto_ptr`, `random_shuffle()`, And Old `<functional>` Stuff.
- Cleaning-up `noexcept` in the Library.
- The Parallelism TS Should be Standardized
- Mathematical Special Functions for C++17
- Adopt Library Fundamentals V1 TS Components for C++17
- `constexpr std::hardware_{constructive,destructive}_interference_size`
- Splicing Maps and Sets
- Deprecating Vestigial Library Parts in C++17
- Better Names for Parallel Execution Policies in C++17
- Hotel Parallelifornia: `terminate()` for Parallel Algorithms Exception Handling
- Elementary string conversions
- Merging `shared_ptr` changes from Library Fundamentals to C++17
- Fail or succeed: there is no atomic lattice
- Throwing out of a parallel algorithm terminates - but how?
- Wording for GB 58 - structured bindings for `node_handles`
- Removing C dependencies from signal handler wording
- Toward a resolution of US7 and US14: Integrating template deduction for class templates into the standard library
- Unifying `<numeric>` Parallel Algorithms
- Iterator Concerns for Parallel Algorithms

- Allowing copies as arguments to function objects given to parallel algorithms in response to CH11
- Wording for CH 10: Complexity of parallel algorithms
- Resolving atomic<T> named base class inconsistencies
- Algorithm Complexity Constraints and Parallel Overloads
- Inline Variables for the Standard Library
- Deprecating <codecv>
- Final C++17 Parallel Algorithms Fixes
- Repairing elementary string conversions

For details about the features of the C++ standard library libstdc++, refer to the GCC website.

## 9.5.2 Half-Precision (16 bit) Floating-Point Type

---

The following two types of half precision (16bit) floating-point type are supported.

`__fp16` data type: Defined in Arm C language extension (IEEE 754-2008)

`__fp16` data type is not an arithmetic data type. `__fp16` data type is used for storing a value and converting a type. Therefore, when a value declared with `__fp16` data type is used at an arithmetic operation, it is converted to single-precision floating-point data type automatically. And, after the arithmetic operation, single-precision floating-point data type is converted to `__fp16` data type.

`_Float16` data type: Defined in extension of C11 (ISO/IEC TS 18661-3:2015)

`_Float16` data type is an arithmetic data type. When an operation is performed by using a value declared with `_Float16` data type, a half-precision arithmetic operation is used. Using `_Float16` data type may increase the execution performance because a type conversion to single-precision floating-point data type is unnecessary.

However, there is no implicit data conversion between `_Float16` data type and single-precision floating-point data type. Therefore, an explicit type conversion is necessary when a value declared with `_Float16` data type is passed as an argument with single-precision floating-point data type.



### Example

Explicit type conversion between `_Float16` and single-precision floating-point data type

```
void half_function(void)
{
    _Float16 value = 1.0f16;
    printf("%f", (float)value);
}
```

## 9.5.3 C++ Standard Library Header <filesystem>

---

When using C++ Standard Library header <filesystem>, specify the specific option at link time as follows.

- When `-stdlib=libstdc++` option is enabled at compile time, specify `-lstdc++fs` option at link time.
- When `-stdlib=libc++` option is enabled at compile time, specify `-lc++fs` option at link time.

## 9.6 Notes on Linking with Different Languages and Trad/Clang Modes

---

This section describes notes about linking with object programs in different programming languages and Trad/Clang Modes.

## 9.6.1 Compile commands and required options when linking

---

Refer to "7.1 Compile Commands and Required Options when Linking".

## 9.6.2 Notes on Linking with the Object of Trad Mode

---

Note the following points when the object file generated in Clang Mode is linked with the object file generated in Trad Mode.

- Specify `libc++` instead of the default `libstdc++` for the C++ Standard Library used in Clang Mode.
- In the following case, the behavior may be undefined, and compiler operation is not guaranteed.
  - An object generated in Clang Mode and an object generated in Trad Mode have an interface between functions. And,
  - An argument or return value of the function is a class type.

## 9.6.3 Linking with C

---

Note the following:

- To call a C++ function from a C function, the C++ function must be declared with linkage to C specified (via `extern "C"`) in the C++ program.
- To call a C function from a C++ function, the C function must be declared with linkage to C specified (via `extern "C"`) in the C++ program.
- The user can specify any type permitted under C for the arguments and return values of a function to be called in a C specified function specified with `extern "C"`.

Example 1: Passing control to the C++ program first

C++ source: `cplusmain.cc`

```
#include <iostream>

extern "C" {
    void cfunc(short int);
    int cfunc(int);
}

int cfunc(int i) {
    std::cout << "C++: cfunc called, i=" << i << std::endl;
    return i;
}

int main() {
    std::cout << "C++ main()" << std::endl;
    cfunc(10);
    return 0;
}
```

C source: `csub.c`

```
#include <stdio.h>

int cfunc(int);

void cfunc(short int si) {
    printf("C: cfunc called, si=%d\n", si);
    cfunc(si);
}
```



Compiling and linking:

```
$ fccpx -Nclang -c csub.c
$ FCCpx -Nclang csub.o cplusplusmain.cc
```

Execution:

```
$ ./a.out
```

Result:

```
C++ main()
C:  cfunc called, si=10
C++: cplusplus called, i=10
```

Example 2: Passing control to the C program first

C++ source: cplusplussub.cc

```
#include <iostream>

extern "C" {
    void cfunc(short int);
    int cplusplus(int);
}

int cplusplus(int i) {
    std::cout << "C++: cplusplus called, i=" << i << std::endl;
    cfunc(i);
    return i;
}
```

C source: cmain.c

```
#include <stdio.h>

int cplusplus(int);

int main() {
    printf("C main()\n");
    cplusplus(10);
    return 0;
}

void cfunc(short int si) {
    printf("C:  cfunc called, si=%d\n", si);
}
```

Compiling and linking:

```
$ fccpx -Nclang -c cmain.c
$ FCCpx -Nclang cmain.o cplusplussub.cc
```

Execution:

```
$ ./a.out
```

Result:

```
C main()
C++: cplusplus called, i=10
C:  cfunc called, si=10
```

## 9.6.4 Linking with Fortran

Note the following:

- To call a Fortran function from C++, the Fortran function must be declared with linkage to C specified in the C++ program.
- To call C++ function from Fortran, the function must be declared with linkage to C specified (extern "C") in C++ program.
- The user can specify any type permitted under C for the arguments and return values of a function to be called in a C specified function.
- If the program to which control is first passed is C++, the function main must be "MAIN\_\_" or "main".
- If the program to which control is first passed is Fortran, the return value of Fortran is set the return value of executable program. For detail of the return value of Fortran, refer to "Fortran User's Guide".
- Other rules are the same as those for interlanguage linkage between C and Fortran. Please refer to "Fortran User's Guide".

When using the COARRAY specification, specify the --linkcoarray option instead of the --linkfortran option. For more information on the COARRAY specification, see "Fortran User's Guide Additional Volume COARRAY".



### Example

Example 1: Passing control to the C++ program first

C++ source: cplusplusmain.cc

```
#include <iostream>
extern "C" {
    int cpfunc_(int *);
    int func_(int *);
}
int cpfunc_(int *p) {
    std::cout << " C++: cpfunc_ called, *p=" << *p << std::endl;
    return *p;
}
int main() {
    std::cout << " C++ MAIN__()" << std::endl;
    int i=10;
    func_(&i);
    return 0;
}
```

Fortran source: fortransub.f95

```
integer function func(x)
integer*4 x
print *, "Fortran: func() called"
call cpfunc(x)
func=x
end
```

Compiling and linking

- When using the libcpp for C++ Standard Library, and using compile command for the Fortran language at linking

```
$ FCCpx -Nclang -c cplusplusmain.cc -stdlib=libcpp
$ frtpx --linkstl=libcpp fortransub.f95 cplusplusmain.o
```

- When using the libstdc++ for C++ Standard Library, and using compile command for the Fortran language at linking

```
$ FCCpx -Nclang -c cplusplusmain.cc -stdlib=libstdc++
$ frtpx --linkstl=libstdc++ fortransub.f95 cplusplusmain.o
```

- When using the compile command for the C++ language at linking

```
$ frtpx -c fortransub.f95
$ FCCpx -Nclang --linkfortran fortransub.o cplusplusmain.cc
```

#### Execution

```
$ ./a.out
```

#### Result

```
C++ MAIN__()
Fortran: func() called
C++: cpfunc_ called. *p=10
```

### Example 2: Passing control to the Fortran program first

C++ source: cplusplussub.cc

```
#include <iostream>
extern "C" {
    int cpfunc_(int *);
}
int cpfunc_(int *p) {
    std::cout << " C++: cpfunc_ called, *p=" << *p << std::endl;
    return *p;
}
```

Fortran source: fortranmain.f95

```
print *, "Fortran: main program"
call cpfunc(10)
end
```

#### Compiling and linking

- When using the libc++ for C++ Standard Library, and using compile command for the Fortran language at linking

```
$ FCCpx -Nclang -c cplusplussub.cc -stdlib=libc++
$ frtpx --linkstl=libc++ fortranmain.f95 cplusplussub.o
```

- When using the libstdc++ for C++ Standard Library, and using compile command for the Fortran language at linking

```
$ FCCpx -Nclang -c cplusplussub.cc -stdlib=libstdc++
$ frtpx --linkstl=libstdc++ fortranmain.f95 cplusplussub.o
```

- When using the compile command for the C++ language at linking

```
$ frtpx -c fortranmain.f95
$ FCCpx -Nclang --linkfortran fortranmain.o cplusplussub.cc
```

#### Execution

```
$ ./a.out
```

#### Result

```
Fortran: main program
C++: cpfunc_ called, *p=10
```

---

## 9.7 SIMD Built-in Functions

Clang Mode supports SIMD (Single Instruction Multi Data) built-in functions.

For details of SIMD built-in functions, refer to the document "ARM C Language Extensions for SVE" published by Arm for developers.



## Note

The Clang Mode supports the SIMD built-in functions documented in the document version 00bet1 (First public release). Note that the SIMD built-in functions added in document version 00bet2 and later are not supported. For more information, see "ARM C Language Extensions for SVE - 1.1.2. Change history".

## 9.8 Compatibility with GNU C++ Specifications

---

This chapter describes the GNU C++ language specifications (GNU C++ Extensions) and compile options (GNU C++ compatible options) in Clang Mode.

### 9.8.1 GNU C++ Extensions

---

For the GNU C++ extensions in Clang Mode, refer to the clang and GCC website.

### 9.8.2 GNU C++ Compatible Options

---

Clang Mode supports the following compatible options (GNU C++ compatible options). For details of GNU C++ compatible options, refer to the GCC website.

`{--print-file-name|-print-file-name)=include`

This option specifies to print the include directory.

`{--print-prog-name|-print-prog-name)={as|ld|objdump|ranlib|ar}`

This option specifies to print the name of programs which are called by compile command.

`--shared`

This option makes the linker create shared objects and not dynamic link executable files.

This option is passed to the linker.

`--version`

This option specifies to emit the version and copyright information of the compiler to the standard output.

`-Wp,-MD,filename`

This option is equivalent to the following options:

`-MD -MF filename`

`-Xlinker option`

This option directs that the *option* is to be passed as arguments to the linker.

`-dM`

This option specifies to output all macro definitions when the `-E` option is effective.

`-f{exceptions|no-exceptions}`

The `-fexceptions` option defines the `__EXCEPTIONS` macro. The `-fno-exceptions` option invalidates the `-fexceptions` option. `-fexceptions` is set by default.

`-fno-common`

This option specifies that the global variable without the initial value is allocated to the data section in the object file.

This option is useful for variables declared with the same name without the "extern" specifier in two or more sources, as this can be detected as an error at linking time.

`-f{optimize-sibling-calls|no-optimize-sibling-calls}`

The compiler option `-foptimize-sibling-calls` specifies to perform optimization of the sibling call. The compiler option `-fno-optimize-sibling-calls` invalidates the compiler option `-foptimize-sibling-calls`. The compiler option `-fno-optimize-sibling-calls` is set by default. When the compiler option `-O1` or higher is set, the compiler option `-foptimize-sibling-calls` is set by default.

The compiler options `-foptimize-sibling-calls` and `-fno-optimize-sibling-calls` require that the compiler option `-O1` or higher is set.

`-f{pie|PIE}`

This option specifies to generate position-independent code.

If the `-fPIE` option is specified, a slower object is generated with a long instruction sequence. However a number of unique global symbols are able to be used in the final shared library that the object will be a part of. If the `-fpie` option is specified, a faster object is generated with a short instruction sequence. However few of unique global symbols are able to be used in the final shared library that the object will be a part of. The number of the unique external symbols that can be referred for these cases is the total about all libraries united at the same time.

If `-fPIE` and `-fpie` options are specified, only the later one is effective.

This option is effective at compilation time only.

The difference between these options and the `-f{pic|PIC}` options is that the objects can be linked with the `-pie` linker option to generate a position independent executable.

`-fvisibility={default|internal|hidden|protected}`

This option specifies the visibility attribute of global symbols in components (executable or shared libraries). `-fvisibility=default` is set by default.

`-fvisibility=default`

The symbol is able to be referred from other components.

`-fvisibility=hidden`

The symbol is not able to be referred from other components. However, it is able to be referred from the other components if the symbol address can be handled in the other components.

`-fvisibility=internal`

This option is equivalent to the `-fvisibility=hidden` option.

`-fvisibility=protected`

The symbol is able to be referred from other components. However, the symbol is not overridden by other symbol with the same name in other components.

`-g{dwarf|dwarf-4}`

This option specifies to add DWARF4 debugging information to the object file.

This option is equivalent to the `-g` option.

`-idirafter dir`

This option adds the *dir* directory to the end of the standard search paths to search the headers.

If multiple directories are specified in multiple `-idirafter` options, the directories are searched in the specified order.

For details of the search order of the headers, refer to the description of `-I` option in Section "[9.1.2.2.1 General Options for Compiler](#)".

If a header is specified with an absolute path name, only the specified absolute path name is searched. If the specified directory does not exist, this option is invalidated.

`-include file`

This option specifies to include the *file* at the top of the source file.

If multiple files are specified in multiple `-include` options, the files are included in the specified order.

`-isystem dir`

This option adds the *dir* directory to the search paths to search the standard headers.

If multiple directories are specified in multiple `-isystem` options, the directories are searched in the specified order.

If a header is specified with an absolute path name, only the specified absolute path name is searched. If the specified directory does not exist, this option is invalidated.

**-nostartfiles**

This option specifies not to use standard system startup files at linking.

**-nostdinc**

This option specifies not to use the standard directory to search headers.

**-nostdlib**

This option specifies not to use standard system startup files and standard libraries at linking.

**-pthread**

This option specifies to create thread safe objects by using the POSIX thread library.

**-rdynamic**

This option specifies to the linker to add all symbols to the dynamic symbol table. This option passes the `-export-dynamic` option to the linker.

**-undef**

This option specifies to suppress the system specific and GNU C++ specific predefined macros.

**-w**

This option suppresses the output of warning messages.

## 9.9 Code Coverage

---

This section explains the code coverage.

The code coverage is the function to measure the execution frequency of the executable statement at execution, and to examine the code coverage rate of the program.

The code coverage uses the coverage tool based on LLVM compiler infrastructure (hereafter called `llvm-cov`), which is an open source software.

Refer to an online manual of the `llvm-cov` command by the `man` command for the use of the `llvm-cov` command.

### 9.9.1 How to Use Code Coverage


---

The code coverage is used the following procedures.

1. Compilation
2. Execution
3. Execution of `llvm-cov` command using the `gcov` mode

A following figure shows the procedure of the code coverage.

Figure 9.1 Procedure of the code coverage



## 9.9.2 Necessary Files to Use Code Coverage

This section explains the necessary files to use the code coverage.

### 9.9.2.1 The ".gcno" File

The .gcno file is a binary file which contains the information obtained at compilation.

File generation step	The file is generated at compilation
File name	The file name is the same as the source program with file extension .gcno.
File generated location	The .gcno file is generated in the directory where the compilation is performed.

An example of generating .gcno file is shown as follows.

## Example

```
$ FCC -Nclang --coverage ../simp.cc
$ ls
a.out simp.gcno
```

When the assembler source file and the object file are generated with specifying the `-o` option, the `-o` option is also applied to the `.gcno` file. An example of generating files with specifying the `-o` option is shown as follows.

## Example

```
$ FCC -Nclang --coverage ../simp.cc -c -o www/yyy.o
$ ls www/
yyy.gcno yyy.o
```

### 9.9.2.2 The ".gcda" File

The `.gcda` file is a binary file which contains the information obtained at execution.

File generation step	The file is generated at execution. The execution frequency is accumulated when the <code>.gcda</code> file with the same name already exists.
File name	The file name is the same as the source program with file extension <code>.gcda</code> .
File generated location	The <code>.gcda</code> file is generated in the directory where the compilation is performed.

An example of generating `.gcda` file is shown as follows.

## Example

```
$ FCC -Nclang --coverage ../simp.cc
$ ls
a.out simp.gcno
$ ./a.out
$ ls
a.out simp.gcda simp.gcno
```

When the assembler source file and the object file are generated with specifying the `-o` option, the `-o` option is also applied to the `.gcda` file. An example of generating files with specifying the `-o` option is shown as follows.

## Example

```
$ FCC -Nclang --coverage ../simp.cc -c -o www/yyy.o
$ ls www/
yyy.gcno yyy.o
$ FCC -Nclang --coverage www/yyy.o
$ ./a.out
$ ls www/
yyy.gcda yyy.gcno yyy.o
```

The storage location directory of the `.gcda` file can be changed by specifying the `-fprofile-dir=dir_name` option.

An example of generating files with specifying the `-fprofile-dir=dir_name` option is shown as follows.





## Example

```
$ FCC -Nclang --coverage ../simp.cc -c -o www/yyy.o -fprofile-dir=/tmp
$ ls www/
YYY.gcno yyy.o
$ FCC -Nclang --coverage www/yyy.o
$ ./a.out
$ ls www/
YYY.gcno yyy.o
$ ls /tmp/www/
YYY.gcda
```

### 9.9.3 Notes on Code Coverage

This section explains notes of the code coverage.

- The execution performance may decrease because the instructions that measure the execution frequency are added in the object program.
- The storage directory of the .gcda file, which is generated at execution, is determined at compilation.
  - The location of the .gcda file storage directory does not change even if the executable program is moved after compilation.
  - Use the `-fprofile-dir=dir_name` option at compilation to change the location of the .gcda file storage directory.
- The code coverage function is not available for programs that need the `-mcmmodel=large` option.
- The execution frequency may be measured incorrectly in the following cases:
  - Two or more executable statements are included in one line.
  - The `exit(3)` function is called.
  - The `setjmp` or `longjmp` function is called.
  - An exception is caught.
  - The `#line` directive is included in the source program.
  - Optimizations of the compiler are applied.

## 9.10 Restrictions and Notes

This section describes the restrictions and notes of Clang Mode.

### 9.10.1 Restrictions of Macros `CMPLX`, `CMPLXF`, and `CMPLXL` in `complex.h`

The macros `CMPLX`, `CMPLXF`, and `CMPLXL` in `complex.h` are not defined.

If a complex type variable is initialized using the macros `CMPLX`, `CMPLXF`, or `CMPLXL` in `complex.h`, the following occurs.

- C program : compilation error or link error
- C++ program : compilation error

Please initialize complex type variables without using the macros `CMPLX`, `CMPLXF`, and `CMPLXL` in `complex.h`.



## Example

Initialization without using macro `CMPLX`

C program

```
double complex a = 1.0 + 2.0 * I;
```

C++ program

```
std::complex<double> a(1.0, 2.0);
```

---

# Appendix A Precompiled Headers

It is often desirable to avoid recompiling a set of header files, especially when they introduce many lines of code and the primary source files that `#include` are relatively small. The compile command front end provides a mechanism for, in effect, taking a snapshot of the state of the compilation at a particular point and writing it to a disk file before completing the compilation; then, when recompiling the same source file or compiling another file with the same set of header files, it can recognize the "snapshot point" and verify that the corresponding precompiled header (PCH) file is read and reusable. Under the right circumstances, this can produce a dramatic improvement in compilation time. However, PCH file uses a large amount of disk space.

## A.1 Automatic Precompiled Header Processing

When `--pch` appears on the command line, automatic precompiled header processing is specified. This means the front end will automatically look for a qualifying precompiled header file to read and/or will create it to use on a subsequent compilation.

The PCH file will contain a snapshot of all the code preceding the "header stop" point. The header stop point is typically the first token in the primary source file that does not belong to a preprocessing directive, but it can also be specified directly by `#pragma hdrstop` (see below) if that comes first. For example:

```
#include "xxx.h"
#include "yyy.h"
int i;
```

The header stop point is `int` (the first non-preprocessor token) and the PCH file will contain a snapshot reflecting the inclusion of `xxx.h` and `yyy.h`. If the first non-preprocessor token or the `#pragma hdrstop` appears within a `#if` block, the header stop point is the outermost enclosing `#if`. To illustrate this, a more complex example is provided:

```
#include "xxx.h"
#ifndef YY_H
#define YY_H 1
#include "yyy.h"
#endif
#if TEST
int i;
#endif
```

Here, the first token that does not belong to a preprocessing directive is again `int`, but the header stop point is the start of the `#if` block containing it. The PCH file will reflect the inclusion of `xxx.h` and conditionally the definition of `YY_H` and inclusion of `yyy.h`; it will not contain the state produced by `#if TEST`.

A PCH file will be produced only if the header stop point and the code preceding it (mainly, the header files themselves) meet certain requirements:

- The header stop point must appear at file scope -- it may not be within an unclosed scope established by a header file. For example, a PCH file will not be created in this case:

```
// xxx.h
class A {
    // xxx.C
    #include "xxx.h"
    int i;
};
```

- The header stop point may not be inside a declaration started within a header file, or part of a declaration list of a linkage specification. For example, in the following case the header stop point is `int`, but since it is not the start of a new declaration, no PCH file will be created:

```
// yyy.h
static

// yyy.C
#include "yyy.h"
int i;
```

- Similarly, the header stop point may not be inside a `#if` block or a `#define` started within a header file.
- The processing preceding the header stop must not have produced any errors. (Note: warnings and other diagnostics will not be reproduced when the PCH file is reused.)
- No references to predefined macros `__DATE__` or `__TIME__` may have appeared.
- No use of the `#line` preprocessing directive may have appeared.
- `#pragma no_pch` (see below) must not have appeared.
- The code preceding the header stop point must have introduced a sufficient number of declarations to justify the overhead associated with precompiled headers.

When a precompiled header file is produced, it contains, in addition to the snapshot of the compiler state, some information that can be checked to determine under what circumstances it can be reused. This includes:

- The compiler version, including the date and time the compiler was built.
- The current directory (i.e., the directory in which the compilation is occurring).
- The command line options.
- The initial sequence of preprocessing directives from the primary source file, including `#include` directives.
- The date and time of the header files specified in `#include` directives.

This information comprises the PCH "prefix". The prefix information of a given source file can be compared to the prefix information of a PCH file to determine whether the latter is applicable to the current compilation.

As an illustration, consider two source files:

```
// a.C
#include "xxx.h"
... // Start of code

// b.C
#include "xxx.h"
... // Start of code
```

`a.C` is compiled with `--pch`, a precompiled header file named `a.pch` is created. Then, when `b.C` is compiled (or when `a.C` is recompiled), the prefix section of `a.pch` is read and compared with the current source file. If the command line options are identical, if `xxx.h` has not been modified, and so forth, then, instead of opening `xxx.h` and processing it line by line, the front end reads in the rest of `a.pch` and thereby establishes the state for the rest of the compilation.

It may be that more than one PCH file is applicable to a given compilation. If so, the largest (i.e., the one representing the most preprocessing directives from the primary source file) is used. For instance, consider a primary source file that begins with

```
#include "xxx.h"
#include "yyy.h"
#include "zzz.h"
```

If there is one PCH file for `xxx.h` and a second for `xxx.h` and `yyy.h`, the latter will be selected (assuming both are applicable to the current compilation).

Moreover, after the PCH file for the first two headers is read in and the third is compiled, a new PCH file for all three headers may be created.

When a precompiled header file is created, it takes the name of the primary source file, with the suffix replaced by an implementation-specified suffix (see `PCH_FILE_SUFFIX`, which is set to `pch` by default). Unless `--pch_dir` is specified (see below), it is created in the directory of the primary source file.

When a precompiled header file is created or used, a message such as

```
"sample.C": creating precompiled header file "sample.pch"
```

is issued. The user may suppress the message by using the command-line option `--no_pch_messages`.

In automatic mode (i.e., when `--pch` is used) the front end will deem a precompiled header file obsolete and delete it under the following circumstances:

- if the precompiled header file is based on at least one out-of-date header file but is otherwise applicable for the current compilation; or
- if the precompiled header file has the same base name as the source file being compiled (e.g., `xxx.pch` and `xxx.C`) but is not applicable for the current compilation (e.g., because of different command-line options).

This handles some common cases; other PCH file clean-up must be dealt with by other means (e.g., by the user).

Support for precompiled header processing is not available when multiple source files are specified in a single compilation: an error will be issued and the compilation aborted if the command line includes a request for precompiled header processing and specifies more than one primary source file.

## A.2 Manual Precompiled Header Processing

---

Command-line option `--create_pch file-name` specifies that a precompiled header file of the specified name should be created.

Command-line option `--use_pch file-name` specifies that the indicated precompiled header file should be used for this compilation; if it is invalid (i.e., if its prefix does not match the prefix for the current primary source file), a warning will be issued and the PCH file will not be used.

When either of these options is used in conjunction with `--pch_dir`, the indicated file name (which may be a path name) is appended to the directory name, unless the file name is an absolute path name.

The `--create_pch`, `--use_pch`, and `--pch` options may not be used together. If more than one of these options is specified, only the last one will apply. Nevertheless, most of the description of automatic PCH processing applies to one or the other of these modes. The header stop points are determined the same way, PCH file applicability is determined the same way, and so forth.

## A.3 Other Ways to Control Precompiled Headers

---

There are several ways in which the user can control and/or tune how precompiled headers are created and used.

- `#pragma hdrstop` may be inserted in the primary source file at a point prior to the first token that does not belong to a preprocessing directive. It enables the user to specify where the set of header files subject to precompilation ends. For example,

```
#include "xxx.h"
#include "yyy.h"
#pragma hdrstop
#include "zzz.h"
```

Here, the precompiled header file will include processing state for `xxx.h` and `yyy.h` but not `zzz.h`. (This is useful if the user decides that the information added by what follows the `#pragma hdrstop` does not justify the creation of another PCH file.)

- `#pragma no_pch` may be used to suppress precompiled header processing for a given source file.

# Appendix B Restrictions and Notes

## B.1 Restriction

This section describes the restrictions of this compilation system.

### B.1.1 OpenMP Program

`threadprivate` Directive

Class array type variables cannot be specified in the `threadprivate` directive.

Example:

```
class A {
    public:
        int mem;
        A(){}
        ~A(){}
}object;

#pragma omp threadprivate(object)

int main(){}
```

OpenMP Directive in Constructor or Destructor

When the OpenMP directive is specified in a constructor or destructor, execution may be incorrect.

Example:

```
class A {
    public:
        int mem;
        A() {
            #pragma omp parallel
            mem++;
        }
        ~A() {
            #pragma omp parallel
            mem--;
        }
} object;

int main(){}
```

Exception Handler in the Static Part of a Parallel Region

When exception handler is specified in the static part of a parallel region, execution may be incorrect.

Example:

```
void func(){
    throw 10;
}
int main(){
    int i =100;
#pragma omp parallel
    {
        try {
            func();
        }
        catch (int){
            i++;
        }
    }
}
```

```
}  
}  
}
```

## B.1.2 C++ Standard Library

---

The C++ Standard Library supports only the "C" locale.

## B.1.3 Restriction on Inheritance of a Constructor with Default Parameters Using "using" Inheritance between Inner Classes

---

In Trad Mode, a compilation error may occur when two inner classes in a class definition have an inheritance relationship and a constructor with default parameters using "using" inheritance between the inner classes. In this case, compile in Clang Mode with the `-Nclang` option, or rewrite the program without the "using" inheritance.

Structures also have the same restriction.



### Example

---

Inheritance of a constructor with default parameters using "using" inheritance between inner classes

```
class A {  
    public:  
    class B {  
        B(int x = 0) { }  
    };  
    class C : B {  
        using B::B;  
        C(void *);  
    };  
};  
A::C c;
```

## B.2 Notes

---

### B.2.1 Variable Arguments of Function

---

In this system, the data with types `class`, `struct` and `union` cannot be specified directly at the variable arguments of a function. If you want to specify the data of these types, please specify the pointer of the data in the arguments.

### B.2.2 Result of the Signed Integer Type Expression

---

When the evaluation result of the signed integer type expression exceeds the range of its type, the result is not guaranteed.

### B.2.3 Undefined Behavior

---

The "undefined behavior" in the C++ Specification is not guaranteed.

### B.2.4 Arrays of Variable Length

---

Programs created by this system allocate arrays of variable length in functions to the stack region.

When the arrays of variable length is huge or inline expansion of function calls is applied to the functions which have arrays of variable length, the stack area needs to be extended to the appropriate size.

The stack area of the process can be set by `ulimit` (bash built-in command), etc.

## B.2.5 Restriction of Register Definition for Floating Point Types with asm Keyword

---

In this system, the data with a long double type or complex type cannot be defined in certain registers with asm keyword.

## B.2.6 Restriction of Floating Point Types Expression as a Register Constraint for asm Operands

---

In this system, the data with a long double type or complex type cannot be specified as a register constraint for asm operands.

## B.2.7 Restriction of Random Access Iterator Type

---

In this system, the variable of random access iterator type cannot be specified for the loop control variable of "for" construct or "parallel for" construct in OpenMP.

## B.2.8 Debug Using Debugger

---

This section explains the notes when the executable file generated in Trad Mode is debugged by using the debugger.

### Notes on using the debugger

In the following cases, debugging is restricted.

- Variables, functions, and types that are declared under the C++11/C++14/C++17 specifications.
- Variables, functions, and types that are declared under the GNU C++ Extensions.

### Debugging of optimized program

- Debugging of the SIMD type variables is restricted.
- Debugging of the parameters and the global variables is guaranteed at the entry of each function.
- Debugging of the local variables is not guaranteed.
- Debugging of the variables of which type is the following is restricted.
  - long double
  - array
  - class
  - structure
  - union
- Debugging of the parameters which are not referred in its function is restricted.
- Debugging of the inlined functions is restricted.
- The displaying of line number is not guaranteed.

## B.2.9 Changing SVE Vector Register Size

---

This system assumes that the SVE vector register size is not changed during executing a program.

If the the effective size of the SVE vector register is changed during executing a program using the system call `prctl(2)` or in other ways, the behavior is undefined.

If you want to change the SVE vector register size, you should set it before executing a program that uses SVE.

## B.2.10 Integer Division Exception when the Divisor Is Zero

---

Note that "Integer divide by zero" is not detected in the Fujitsu CPU A64FX with Arm architecture.



If the compiler option `-NRtrap` is specified and the divisor is 0, the program that expects "Integer divide by zero" to be detected may have different results.

It is recommended to check the value of the divisor just before the integer division.



## Example

Example of checking the value of a divisor just before integer division

```
#include <stdio.h>
int main()
{
    int a,b,c;
    a = 1;
    b = 0;
    if (b == 0) {
        fprintf(stderr, "Integer divide by zero\n");
        exit(1);
    }
    return 0;
}
```

## B.2.11 About Compile Time for Template Functions

When compiling many template functions in Trad Mode, the compile time may increase. In this case, compile in Clang Mode with the `-Nclang` option.

## B.2.12 Link Error (undefined reference to)

When compiling a program, if the following operation is performed by configure etc., a link error (undefined reference to) may occur because the combining of this system original objects that normally performed at linking is suppressed.

- The `-L` option which is passed to the linker by this system is directly specified on the compile command to pass the linker by user.



## Example

Example of a link error

```
test.o: In function `__fjc_check_hpctag':
(.text+0x1620): undefined reference to `__jwe_check_hpctag'
```

This can be avoided by setting the environment variable `FCOMP_LINK_FJOB`. Refer to "[2.3 Environment Variable for Compile Command](#)" for the environment variable `FCOMP_LINK_FJOB`.

## B.2.13 Restriction of registers with asm keyword and on an inline asm clobber lists

In this system, the following registers reserved by the compiler cannot be specified in the register name of asm keyword and on an inline asm clobber lists.

Compiler reserved registers: x19, w19, r19, x29, w29, r29, sp, xzr, xsp, wzr, wsp

# Appendix C Compatibility with GNU C++ Specifications

This system partially supports the GNU C++ Compiler specification (version 6.1).

This chapter describes the GNU C++ language specifications (GNU C++ Extensions) and compile options (GNU C++ compatible options).

## C.1 GNU C++ Extensions

The GNU C++ Extensions supported by this system are as follows. For more detail about GNU C++ Extensions, refer to the GCC website.

- Statements and Declarations in Expressions
- Locally Declared Labels
- Labels as Values
- Referring to a Type with `typeof`
- Conditionals with Omitted Operands
- Double-Word Integers
- Complex Numbers  
Complex integral types are not supported.
- Hex Floats
- Arrays of Length Zero
- Arrays of Variable Length
- Macros with a Variable Number of Arguments
- Slightly Looser Rules for Escaped Newlines
- Non-Lvalue Arrays May Have Subscripts
- Arithmetic on void- and Function-Pointers
- Non-Constant Initializers
- Compound Literals
- Designated Initializers
- Case Ranges
- Cast to a Union Type
- Declaring Attributes of Functions/Specifying Attributes of Variables/Specifying Attributes of Types  
For the list of attributes that can be used, see "[C.1.1 Attributes](#)".
- The Character ESC in Constants
- Inquiring on Alignment of Types or Variables
- Alternate Keywords
- Function Names as Strings
- Builtins
- Unnamed struct/union fields within structs/unions
- Assembler Instructions with C Expression Operands
- Dollar Signs in Identifier Names
- Pragmas Accepted by GCC
  - `#pragma redefine_extname`

- #pragma pack
- #pragma weak
- Thread-Local Storage

## C.1.1 Attributes

---

The attributes supported by this system are as follows. For more detail about attributes, refer to the GCC website.

For the attributes supported by Clang Mode, see "<https://releases.lvm.org/7.0.0/tools/clang/docs/AttributeReference.html>".

Note that the attributes at explicit instantiation are not available in this system.

- `aarch64_vector_pcs` (available for Functions)
- `alias`
- `aligned` (available for Variables, Types)
- `always_inline`
- `const`
- `constructor` (available for Functions)
- `destructor` (available for Functions)
- `deprecated` (available for Functions, Types)
- `format`
- `format_arg` (available for Functions)
- `init_priority` (available for Functions, Types)
- `malloc`
- `may_alias` (available for Types)
- `mode`
- `no_instrument_function` (available for Functions)
- `noinline`
- `nonnull`
- `nopl` (available for Functions)
- `noreturn`
- `nothrow`
- `packed` (available for Variables, Types)
- `pure`
- `section` (available for Functions, Variables)
- `sentinel` (available for Functions)
- `strong` (available for using-directive)
- `unused` (available for Functions, Variables, Types)
- `used` (available for Functions, Variables)
- `transparent_union`
- `visibility` (available for Functions, Types)

- warn\_unused\_result
- weak (available for Functions, Variables)
- weakref (available for Functions)

## C.1.2 Built-in Functions

---

The Built-in Functions supported by this system are as follows. For more detail about built-in Functions, refer to the GCC website.

- \_\_atomic\_load\_n
- \_\_atomic\_load
- \_\_atomic\_store\_n
- \_\_atomic\_store
- \_\_atomic\_exchange\_n
- \_\_atomic\_exchange
- \_\_atomic\_compare\_exchange\_n
- \_\_atomic\_compare\_exchange
- \_\_atomic\_add\_fetch
- \_\_atomic\_sub\_fetch
- \_\_atomic\_and\_fetch
- \_\_atomic\_xor\_fetch
- \_\_atomic\_or\_fetch
- \_\_atomic\_nand\_fetch
- \_\_atomic\_fetch\_add
- \_\_atomic\_fetch\_sub
- \_\_atomic\_fetch\_and
- \_\_atomic\_fetch\_xor
- \_\_atomic\_fetch\_or
- \_\_atomic\_fetch\_nand
- \_\_atomic\_test\_and\_set
- \_\_atomic\_clear
- \_\_atomic\_thread\_fence
- \_\_atomic\_signal\_fence
- \_\_atomic\_always\_lock\_free
- \_\_atomic\_is\_lock\_free
- \_\_builtin\_abort
- \_\_builtin\_alloca
- \_\_builtin\_bswap16
- \_\_builtin\_bswap32
- \_\_builtin\_bswap64

- \_\_builtin\_calloc
- \_\_builtin\_choose\_expr
- \_\_builtin\_clz
- \_\_builtin\_clzl
- \_\_builtin\_clzll
- \_\_builtin\_constant\_p
- \_\_builtin\_ctz
- \_\_builtin\_ctzl
- \_\_builtin\_ctzll
- \_\_builtin\_exit
- \_\_builtin\_expect
- \_\_builtin\_extract\_return\_addr
- \_\_builtin\_ffs
- \_\_builtin\_ffsl
- \_\_builtin\_ffsll
- \_\_builtin\_fpclassify
- \_\_builtin\_fprintf
- \_\_builtin\_fputc
- \_\_builtin\_fputs
- \_\_builtin\_fscanf
- \_\_builtin\_fwrite
- \_\_builtin\_huge\_val
- \_\_builtin\_huge\_valf
- \_\_builtin\_huge\_vall
- \_\_builtin\_index
- \_\_builtin\_inf
- \_\_builtin\_inff
- \_\_builtin\_infl
- \_\_builtin\_isfinite
- \_\_builtin\_isinf
- \_\_builtin\_isinf\_sign
- \_\_builtin\_isinff
- \_\_builtin\_isinfl
- \_\_builtin\_isnan
- \_\_builtin\_isnanf
- \_\_builtin\_isnanl
- \_\_builtin\_isnormal

- \_\_builtin\_malloc
- \_\_builtin\_memchr
- \_\_builtin\_memcmp
- \_\_builtin\_memcpy
- \_\_builtin\_memmove
- \_\_builtin\_memcpy
- \_\_builtin\_memset
- \_\_builtin\_nan
- \_\_builtin\_nanf
- \_\_builtin\_nanl
- \_\_builtin\_nans
- \_\_builtin\_nansf
- \_\_builtin\_nansl
- \_\_builtin\_offsetof
- \_\_builtin\_popcount
- \_\_builtin\_popcountl
- \_\_builtin\_popcountll
- \_\_builtin\_prefetch
- \_\_builtin\_printf
- \_\_builtin\_putchar
- \_\_builtin\_puts
- \_\_builtin\_return\_address
- \_\_builtin\_rindex
- \_\_builtin\_scanf
- \_\_builtin\_signbit
- \_\_builtin\_signbitf
- \_\_builtin\_signbitl
- \_\_builtin\_sprintf
- \_\_builtin\_sscanf
- \_\_builtin\_stpcpy
- \_\_builtin\_strcat
- \_\_builtin strchr
- \_\_builtin\_strcmp
- \_\_builtin\_strcpy
- \_\_builtin\_strerror
- \_\_builtin\_strlen
- \_\_builtin\_strncat

- \_\_builtin\_strncmp
- \_\_builtin\_strncpy
- \_\_builtin\_strpbrk
- \_\_builtin\_strrchr
- \_\_builtin\_strspn
- \_\_builtin\_strstr
- \_\_builtin\_trap
- \_\_builtin\_types\_compatible\_p
- \_\_builtin\_vfprintf
- \_\_builtin\_vprintf
- \_\_builtin\_vsprintf
- \_\_sync\_fetch\_and\_add
- \_\_sync\_fetch\_and\_sub
- \_\_sync\_fetch\_and\_or
- \_\_sync\_fetch\_and\_and
- \_\_sync\_fetch\_and\_xor
- \_\_sync\_fetch\_and\_nand
- \_\_sync\_add\_and\_fetch
- \_\_sync\_sub\_and\_fetch
- \_\_sync\_or\_and\_fetch
- \_\_sync\_and\_and\_fetch
- \_\_sync\_xor\_and\_fetch
- \_\_sync\_nand\_and\_fetch
- \_\_sync\_bool\_compare\_and\_swap
- \_\_sync\_val\_compare\_and\_swap
- \_\_sync\_synchronize
- \_\_sync\_lock\_test\_and\_set
- \_\_sync\_lock\_release

## C.2 GNU C++ Compatible Options

---

This system partially supports the compatible options (GNU C++ compatible options).

For more detail about these GNU C++ compatible options, refer to the GCC website.

`{--print-file-name|-print-file-name)=include`

This option specifies to print the include directory.

`{--print-prog-name|-print-prog-name)={as|ld|objdump|ranlib|ar}`

This option specifies to print the name of programs which are called by compile command.

`--shared`

This option makes the linker create shared objects and not dynamic link executable files.

This option is passed to the linker.

#### --version

This option specifies to emit the version and copyright information of the compiler to the standard output.

#### -Wp,-MD,*filename*

This option is equivalent to the following options:

-MD -MF *filename*

#### -Xlinker *option*

This option directs that the *option* is to be passed as arguments to the linker.

#### -dM

This option specifies to output all macro definitions when the -E option is effective.

#### -f{align-loops[=*N*]|no-align-loops}

The -falign-loops option specifies to align the loop to the boundary of power-of-two byte.

The -fno-align-loops option invalidates the -falign-loops option. -fno-align-loops is set by default.

#### -f{exceptions|no-exceptions}

The -fexceptions option specifies to enable exception handling. Also, -fexceptions defines the \_\_EXCEPTIONS macro. The -fno-exceptions option invalidates the -fexceptions option. -fexceptions is set by default.

The -fexceptions option is necessary when a program handles exception.

If -fno-exceptions is specified and the program uses exception handling syntax, compilation is terminated.

#### -ffast-math

This option specifies to perform the math functions optimization that do not conform strictly to the specifications/rules of IEEE 754 or ISO. When this option is effective, it may give rise to side effects (calculation error or runtime exceptions) in the execution results.

#### -ffp-contract=fast

This option specifies to perform optimization using Floating-Point Multiply-Add/Subtract instructions.

When this option is set, side effects (calculation errors in rounding error extent) may occur in the execution results and produce unexpected results.

#### -f{inline-functions|no-inline-functions}

-finline-functions option specifies to perform the inline expansion for simple functions defined in the source code instead of function calls. The function which is the subject of inline expansion is determined by the compiler automatically. The -fno-inline-functions option is set by default.

#### -floop-parallelize-all

This option specifies to perform automatic parallelization. However, if the effect of parallel execution is not expected, automatic parallelization is not performed.

If the object program compiled with this option is included in the filename list, this option must also be specified when linking.

#### -fno-common

This option specifies that the global variable without the initial value specification is allocated to the data section in the object file.

This option is useful for variables declared with the same name without the "extern" specifier in two or more sources, as this can be detected as an error at linking time.

#### -f{omit-frame-pointer|no-omit-frame-pointer}

The -fomit-frame-pointer option specifies to suppress the storing of the frame pointer in the function that need not store the frame pointer in the register. The -fno-omit-frame-pointer option is set by default.

#### -fopenmp

The -fopenmp option specifies to enable Specification of OpenMP Application Program Interface.



When the `-fopenmp` option is specified, `-mt` is set.

The `-fopenmp` option is needed if an object program compiled with the `-fopenmp` option exists in the command line as input files.

`-f{openmp-simd|no-openmp-simd}`

The `-fopenmp-simd` specifies that only the OpenMP `simd` construct and the `declare simd` construct are enabled.

`-fno-openmp-simd` is set by default.

`-f{optimize-sibling-calls|no-optimize-sibling-calls}`

The `-foptimize-sibling-calls` option specifies to perform optimization of the sibling call. The `-fno-optimize-sibling-calls` option invalidates the `-foptimize-sibling-calls` option. `-fno-optimize-sibling-calls` is set by default.

`-f{pic|PIC}`

This option specifies to generate position-independent code (PIC).

If the `-fPIC` option is specified, a slower object is generated with a long instruction sequence. However a number of unique global symbols are able to be used in the final shared library that the object will be a part of. If the `-fpic` option is specified, a faster object is generated with a short instruction sequence. However few of unique global symbols are able to be used in the final shared library that the object will be a part of. The number of the unique external symbols that can be referred for these cases is the total about all libraries united at the same time.

If the `-fPIC` and `-fpic` options are specified, only the later one is effective.

This option is effective at compilation time only.

`-f{pie|PIE}`

This option specifies to generate position-independent code.

If the `-fPIE` option is specified, a slower object is generated with a long instruction sequence. However a number of unique global symbols are able to be used in the final shared library that the object will be a part of. If the `-fpie` option is specified, a faster object is generated with a short instruction sequence. However few of unique global symbols are able to be used in the final shared library that the object will be a part of. The number of the unique external symbols that can be referred for these cases is the total about all libraries united at the same time.

If `-fPIE` and `-fpie` options are specified, only the later one is effective.

This option is effective at compilation time only.

The difference between these options and the `-f{pic|PIC}` options is that the objects can be linked with the `-pie` linker option to generate a position independent executable.

`-f{plt|no-plt}`

The `-fplt` option specifies to use Procedure Linkage Table (PLT) for accessing a global symbol in the position-independent code (PIC). The `-fno-plt` option invalidates the `-fplt` option. `-fplt` is set by default.

`-fprofile-dir=path`

This option specifies storage location directory of the `.gdcda` file which is necessary for the use of the code coverage. For *path*, the storage location directory name is specified by the relative path or the absolute path.

`-f{schedule-insns|no-schedule-insns}`

The `-fschedule-insns` option specifies to perform the optimization of an instruction scheduling before the register is allocated. The `-fno-schedule-insns` option invalidates the `-fschedule-insns` option. `-fno-schedule-insns` is set by default.

`-f{schedule-insns2|no-schedule-insns2}`

The `-fschedule-insns2` option specifies to perform the optimization of an instruction scheduling after the register is allocated. The `-fno-schedule-insns2` option invalidates the `-fschedule-insns2` option. `-fno-schedule-insns2` is set by default.

`-f{signed-char|unsigned-char}`

The `-fsigned-char` option specifies that `char` declaration is assumed to be `signed char` declaration. The `-funsigned-char` option invalidates the `-fsigned-char` option. The `-funsigned-char` option is applied by default.

### -f{strict-aliasing|no-strict-aliasing}

The `-fstrict-aliasing` option specifies that in case of indirect memory access through a pointer, when the accessing types are different, no memory alias is assumed. The `-fno-strict-aliasing` option invalidates the `-fstrict-aliasing` option. The `-fno-strict-aliasing` is set by default.

### -funroll-loops

This option specifies to apply loop unrolling.

### -f{unsafe-math-optimizations|no-unsafe-math-optimizations}

The `-funsafe-math-optimizations` option specifies to use flush-to-zero mode. The default is `-fno-unsafe-math-optimizations`.

When `-funsafe-math-optimizations` option is specified, side effects may occur in the execution results.

These options must be set at linking.

### -g{dwarf|dwarf-4}

This option specifies to add DWARF4 debugging information to the object file.

This option is equivalent to the `-g` option.

### -idirafter *dir*

This option adds the *dir* directory to the end of the standard search paths to search the headers.

If multiple directories are specified in multiple `-idirafter` options, the directories are searched in the specified order.

For details about the order of the search of the header, see the `-I` option in Section "2.2.2.1 General Options for Compiler".

If a header is specified with an absolute path name, only the specified absolute path name is searched.

If the specified directory does not exist, this option is invalidated.

### -include *file*

This option specifies to include the *file* at the top of the source file.

If multiple files are specified in multiple `-include` options, the files are included in the specified order.

### -isystem *dir*

This option adds the *dir* directory to the search paths to search the standard headers.

If multiple directories are specified in multiple `-isystem` options, the directories are searched in the specified order.

For details about the order of the search of the header, see the `-I` option in Section "2.2.2.1 General Options for Compiler".

If a header is specified with an absolute path name, only the specified absolute path name is searched.

If the specified directory does not exist, this option is invalidated.

### -march=*arch*[+*features*]

This option specifies the architecture. In *arch*, specifies one of `armv8-a`, `armv8.1-a`, `armv8.2-a`, or `armv8.3-a`.

In *features*, specifies `sve` or `nosve`. The default for *features* is `sve`.

If `-march` option is omitted, `-march=armv8.3-a+sve` is set by default.

### -mcmmodel={small|large}

This option specifies the possible largest size of the text area and the static data area in an executable program or a shared object. `-mcmmodel=small` is set by default.

#### -mcmmodel=small

The total size of the text area and the static data area is limited to 4GB at linking. This option creates an efficient object program.

#### -mcmmodel=large

Only the size of the text area is limited to 4GB at linking. This option is used when the static data area is large and an error occurs at linking.

**-mcpu=*cpu***

This option specifies the processor. In *cpu*, specifies one of a64fx, thunderx2t99, or generic. -mcpu=a64fx is set by default.

**-m{pc-relative-literal-loads|no-pc-relative-literal-loads}**

The -mpc-relative-literal-loads option specifies to access the literal pool by one instruction assuming that the code area in the function is within 1MB.

The -mno-pc-relative-literal-loads option invalidates the -mpc-relative-literal-loads. -mno-pc-relative-literal-loads is set by default.

**-mtls-size={12|24|32|48}**

This option specifies the size of an offset necessary for the access to Thread-Local Storage. Units are bits.

**-mtune=*cpu***

This option specifies the processor which optimizations targets. In *cpu*, specifies one of a64fx, thunderx2t99, or generic. -mtune=a64fx is set by default.

**-nostartfiles**

This option specifies not to use standard system startup files at linking.

**-nostdinc**

This option specifies not to use the standard directory to search headers.

**-nostdlib**

This option specifies not to use standard system startup files and standard libraries at linking.

**-pthread**

This option specifies to create thread safe objects by using the POSIX thread library.

**-rdynamic**

This option specifies to the linker to add all symbols to the dynamic symbol table.

This option passes the -export-dynamic option to the linker.

**-undef**

This option specifies to suppress the system specific and GNU C++ specific predefined macros.

**-v**

This option specifies to display the executing command which compile command called as compiler, assembler and linker.

**-w**

This option suppresses the output of warning messages.

# Appendix D Data and Memory Regions

This appendix describes attribute of the data and memory regions in this system.

## D.1 Data Size and Alignment

Data types, sizes, and alignments of basic data system are shown below.

Data type <sup>[a]</sup>	Size(byte)	Alignment (byte)
char	1	1
short int	2	2
int	4	4
long int	8	8
long long int		
float	4	4
double	8	8
long double	16	16
float _Complex	8	4
double _Complex	16	8
long double _Complex	32	16
struct	-	[b]
union		

[a] Common to signed and unsigned.

[b] The largest value among the maximum alignment of the member variables and 8 byte.

## D.2 Memory Regions

The main usages of the memory regions are shown below.

Memory regions	Usage
Text	Memory regions where instructions are arranged.
Static data(.data)	Memory regions where static data with initialization is stored.
Static data(.bss)	Memory regions where static data without initialization is stored.
Process stack	Stack regions for the process and the main thread.
Thread stack	Stack regions for the child process.
Dynamically allocated memory region	Memory regions allocated when requiring memory to be allocated dynamically, or memory regions allocated as thread heap region or thread stack.
Shared memory	Memory regions shared among processes.

## D.3 Data Allocation

Variables in a program are arranged to different memory regions depending on each type of the variables and existence of initial values.

## Example

```
int e[2];
int f[2] = {3,3};

int main(){
    int      a[2];
    int      b[2] = {1,1};
    static float c[2];
    static float d[2] = {2.0,2.0};
    double   *g;

    g = (double *)malloc(sizeof(double)*2);
    ...
}
```

In the above example, each array is allocated as follows.

Array name	Region used
a	Allocated to the process stack region, because these are local arrays.
b	
c	Allocated to the static data (.bss) region, because it is a static array and not initialized.
d	Allocated to the static data (.data) region, because it is a static array and initialized.
e	Allocated to the static data (.bss) region, because it is a global array and not initialized.
f	Allocated to the static data (.data) region, because it is a global array and initialized
g	Allocated to the dynamically allocated memory region.

## D.4 Data Allocation to Stack Region

In this system, the order of data allocation to the stack region can be controlled by the compiler option.

The compiler option that controls the order of data allocation is shown as follows.

`-N{reordered_variable_stack|noreordered_variable_stack}`

The `-Nnoreordered_variable_stack` option directs the compiler the order in which to allocate the automatic variables to the stack area.

If the `-Nnoreordered_variable_stack` option is set, the allocation order of automatic variables is determined in the following order.

1. Descending order of their alignments
2. Descending order of data sizes if the alignments are equal
3. The order of appearance of the declaration statements in the source program if both the alignments and data sizes are equal

The stack area of the entire program can be reduced by allocating the automatic variables in descending order of their alignments.

When the `-Nnoreordered_variable_stack` option is specified, automatic variables are allocated in order of the appearance of the declaration statements in the source program. `-Nnoreordered_variable_stack` is set by default.

The order of allocation is not guaranteed when the `-Nnoline`, `-g0`, or `-Kocl` option is set.

## Example

---


### - Source Program

```
int    a;
double b;
char   c;
double d[2],e[2];
short int f;
a = 1;
b = 2.0;
c = '3';
d[0] = 4.0;
e[0] = 5.0;
f = 6;
```

### - When specifying the order of data allocation for the example source program


Automatic variables are stored in the stack as shown below. The stack region can be decreased.

Figure D.1 When the `-Nreordered_variable_stack` option is valid



- When not specifying the order of data allocation for the example source program Automatic variables are stored in the stack as shown below.

Figure D.2 When the -Nnoreordered\_variable\_stack option is valid



# Appendix E Code Coverage

This appendix explains the code coverage.

The code coverage is the function to measure the execution frequency of the executable statement at execution, and to examine the code coverage rate of the program.

The code coverage uses the coverage tool based on LLVM compiler infrastructure (hereafter called llvm-cov), which is an open source software.

Refer to an online manual of the llvm-cov command by the man command for the use of the llvm-cov command.

## E.1 How to Use Code Coverage

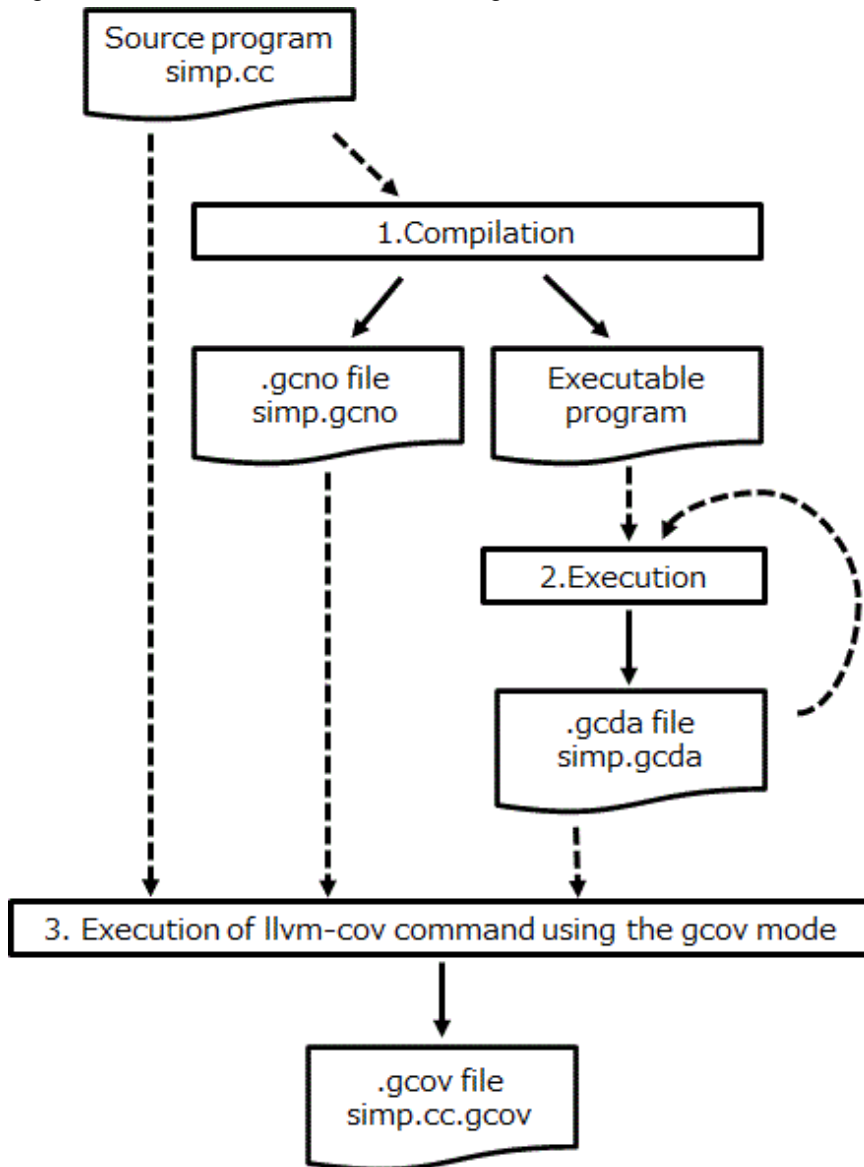
The code coverage is used the following procedures.

1. Compilation
2. Execution
3. Execution of llvm-cov command using the gcov mode

A following figure shows the procedure of the code coverage.



Figure E.1 Procedure of the code coverage



## E.2 Necessary Files to Use Code Coverage

This section explains the necessary files to use the code coverage.

### E.2.1 The ".gcno" File

The .gcno file is a binary file which contains the information obtained at compilation.

File generation step	The file is generated at compilation
File name	The file name is the same as the source program with file extension .gcno.
File generated location	The .gcno file is generated in the directory where the compilation is performed.

An example of generating .gcno file is shown as follows.

## Example

```
$ FCC -Ncoverage ../simp.cc
$ ls
a.out simp.gcno
```

When the assembler source file and the object file are generated with specifying the `-o` option, the `-o` option is also applied to the `.gcno` file. An example of generating files with specifying the `-o` option is shown as follows.

## Example

```
$ FCC -Ncoverage ../simp.cc -c -o www/yyy.o
$ ls www/
yyy.gcno yyy.o
```

## E.2.2 The ".gcda" File

The `.gcda` file is a binary file which contains the information obtained at execution.

File generation step	The file is generated at execution. The execution frequency is accumulated when the <code>.gcda</code> file with the same name already exists.
File name	The file name is the same as the source program with file extension <code>.gcda</code> .
File generated location	The <code>.gcda</code> file is generated in the directory where the compilation is performed.

An example of generating `.gcda` file is shown as follows.

## Example

```
$ FCC -Ncoverage ../simp.cc
$ ls
a.out simp.gcno
$ ./a.out
$ ls
a.out simp.gcda simp.gcno
```

When the assembler source file and the object file are generated with specifying the `-o` option, the `-o` option is also applied to the `.gcda` file. An example of generating files with specifying the `-o` option is shown as follows.

## Example

```
$ FCC -Ncoverage ../simp.cc -c -o www/yyy.o
$ ls www/
yyy.gcno yyy.o
$ FCC -Ncoverage www/yyy.o
$ ./a.out
$ ls www/
yyy.gcda yyy.gcno yyy.o
```

The storage location directory of the `.gcda` file can be changed by specifying the `-Nprofile_dir=dir_name` option.

An example of generating files with specifying the `-Nprofile_dir=dir_name` option is shown as follows.



## Example

```
$ FCC -Ncoverage ../simp.cc -c -o www/yyy.o -Nprofile_dir=/tmp
$ ls www/
YYY.gcno yyy.o
$ FCC -Ncoverage www/yyy.o
$ ./a.out
$ ls www/
YYY.gcno yyy.o
$ ls /tmp/www/
YYY.gcda
```

## E.3 Notes on Code Coverage

---

This section explains notes of the code coverage.

- The execution performance may decrease because the instructions that measure the execution frequency are added in the object program.
- The storage directory of the .gcda file, which is generated at execution, is determined at compilation.
  - The location of the .gcda file storage directory does not change even if the executable program is moved after compilation.
  - Use the `-Nprofile_dir=dir_name` option at compilation to change the location of the .gcda file storage directory.
- The code coverage function is not available for programs that need the `-Kcmodel=large` or `-mcmmodel=large` option.
- The execution frequency may be measured incorrectly in the following cases:
  - Two or more executable statements are included in one line.
  - The `exit(3)` function is called.
  - The `setjmp` or `longjmp` function is called.
  - An exception is caught.
  - The `#line` directive is included in the source program.
  - Optimizations of the compiler are applied.

# Appendix F Fujitsu Extended Functions

This appendix describes the Fujitsu Extended Functions supported in this system.

## F.1 How to Use Fujitsu Extended Functions

### F.1.1 Header

Fujitsu Extended Functions need a header shown below to be included in the source file.

- fjcx.h

### F.1.2 Supported Functions

The system supports Fujitsu Extended Functions below.

Table F.1 List of Fujitsu Extended Functions

Function Prototype	Description
double __gettod(void)	The __gettod() function returns the current high-resolution real time, expressed as micro-seconds since the some arbitrary time in the past. Typically, the time is expressed since the most recent system booting.

# Appendix G Runtime Information Output Function

This appendix describes the Runtime Information Output Function.

This function outputs information for speedup and performance confirmation of the program, by specifying the compiler option and setting the environment variable, when the program is executed.

## G.1 Usage of Runtime Information Output Function

This section describes the usage of Runtime Information Output Function.

### G.1.1 Range of Obtainable Information

This section describes the range of information that can be fetched using this function.

Table G.1 Range of obtainable information

Compiler Options	Runtime environment variables	Range of obtainable information
-Nrt_tune	FLIB_RTINFO	from the start of the program to the end
-Nrt_tune_func		each user-defined function
-Nrt_tune_loop		each loop

When the -Nrt\_tune compile option and FLIB\_RTINFO runtime environment variable is specified, this function is enabled and information can be fetched from the start of the program to the end.

Also, if the -Nrt\_tune\_func compile option or -Nrt\_tune\_loop compile option is specified, additional information concerning each function, each loop can be fetched.

Refer to "2.2 Compiler Options" for further details.

If the MPI Programs or COARRAY Programs, information for each process is output.

#### Each User-defined Function

When the -Nrt\_tune\_func compile option is specified, information for each user-defined function is output.

When information is fetched for each user-defined function, the information below is fetched in the cases shown below.

- If a function is called from within a function:
- Information for the called function is not included in the calling source function information.
- If the same function is called from multiple locations :
- The total values are used as the information for that function.
- Only user-defined function s are output. Information concerning system calls and library functions is not fetched.

#### Each Loop

When the -Nrt\_tune\_loop compile option is specified, information for each loop is output.

The following information is output, as information for each loop.

- Information for each loop operates in serial processing.
- Information for each loop operates in parallel region by OpenMP.

The cost of the serial loop is output as information for serial loop. If the loop exists in parallel region by OpenMP, it is output as each loop in parallel region by OpenMP. As the information of loop in parallel region by OpenMP, information of the master thread is output.

When information is fetched for each loop, the information below is fetched in the cases shown below.

- If a function is called from within a loop :

- Information for the called function is included in the calling source loop information.
- If a loop is called from within a loop :
- Information for the called loop is included in the calling source loop information.

## G.1.2 Runtime Environment Variables

This section describes the runtime environment variables using this function.

The English small letter and the English capital letter specified for the value of the environment variable at execution time are distinguished. For environment variables that do not require operands, the specified operand value is ignored.

Table G.2 Environment variables that control the information output

Runtime environment variables		Explanation
Variable name	Operand	
FLIB_RTINFO	(None)	Outputs runtime information.
FLIB_RTINFO_NOFUNC	(None)	Even if the compiler option -Nrt_tune_func is effective, the output of information for each function is suppressed.
FLIB_RTINFO_NOLOOP	(None)	Even if the compiler option -Nrt_tune_loop is effective, the output of information for each loop is suppressed.
FLIB_RTINFO_CSV	<i>Filename</i>	Outputs the fetched information to a CSV file.
	-	Operand is optional. Any file name can be specified as operand of the environment variable. If operand is omitted, flib_rtinfo.csv is assumed.

## G.2 Output of Runtime Information Output Function

This section describes the output of Runtime Information Output Function.

### G.2.1 Output Information

This section describes the information output by this function.

Table G.3 Output information

Compiler Options	Runtime environment variables	Output information
-Nrt_tune	FLIB_RTINFO	The following information from the start of the program to the end is output <ul style="list-style-type: none"> <li>- Elapsed time</li> <li>- User CPU time</li> <li>- System CPU time</li> </ul>
-Nrt_tune_func		In addition to the -Nrt_tune output, following Information of each function is output <ul style="list-style-type: none"> <li>- Function cost</li> <li>- Percentage of total cost accounted for by the function cost</li> <li>- Function cost per called once</li> <li>- Function start line number</li> <li>- Function end line number</li> <li>- Call frequency of function</li> </ul>

Compiler Options	Runtime environment variables	Output information
		- Function name
-Nrt_tune_loop		<p>In addition to the -Nrt_tune output, following Information of each loop in serial and each loop that in parallel region by OpenMP is output</p> <ul style="list-style-type: none"> <li>- Loop cost</li> <li>- Percentage of total cost accounted for by the loop cost</li> <li>- Loop cost per called once</li> <li>- Loop start line number</li> <li>- Loop end line number</li> <li>- Call frequency of Loop</li> <li>- Loop nest level</li> <li>- Generated function name</li> </ul>

## G.2.2 Output Format

Runtime information can be output by the text format or CSV format.

### Text format

When FLIB\_RTINFO\_CSV environment variable is not specified, information is output to the standard output by text format.

Refer to "G.2.3 Output Example" for details of output format of information.

### CSV format

When the FLIB\_RTINFO\_CSV environment variable is specified, information is output in CSV format. If a filename is specified in the operands, the fetched information is output to the specified file. If the specified file already exists, the information overwrites the existing file. If the file fails to open, runtime message jwe1653i-w is output and the information is output to the standard output.

When runtime information is output by CSV, data line follows the following compositions. When there is no output information, "-" is output.

Table G.4 Composition of CSV format

Column position	Output information
First	<p>Identifier that shows type</p> <p>[COST] : Information for entire program</p> <p>[COST_ROUTINE] : Information for each function</p> <p>[COST_LOOP_SERIAL] : Information for each loop that operates in serial processing</p> <p>[COST_LOOP_PRL] : Information for each loop that in parallel region by OpenMP</p>
Second and subsequent	It is depending on output information.

Refer to "G.2.3 Output Example" for details of output format of information.

## G.2.3 Output Example

Output example of Runtime Information Output Function is shown below.

Output example of runtime information

```

++=====
++

```

EXECUTION PERFORMANCE INFORMATION																						
++																						
(1)	-----		(2)	-----		(3)	-----															
	Elapsed		User		System																	
	(sec)		(sec)		(sec)																	
+-----+-----+-----+																						
	2.5352		5.0552		0.0060																	
+-----+-----+-----+																						
Routine (4)																						
(5)	-----		(6)	-----		(7)	-----		(8)	-----		(9)	-----		(10)	-----		(11)	-----			
	Cost(sec)		%		Once(sec)		Start		End		Count		Routine name									
+-----+																						
	1.9550		84.38		0.0010		73		88		2000		sub2(int)						(12)			
	0.3620		15.62		0.0362		40		70		10		sub1(int)									
+-----+																						
	2.3169		100.00		-		-		-		-		-						(13)			
+-----+																						
Serial Loop (14)																						
															(15)	-----						
	Cost(sec)		%		Once(sec)		Start		End		Count		Nest		Routine name							
+-----+																						
	2.1647		45.15		2.1647		27		29		1		1		main			(16)				
	2.1419		44.68		0.0011		79		80		2000		2		(sub2(int))							
+-----+																						
	4.7942		100.00		-		-		-		-		-		-			(17)				
+-----+																						
Parallel Loop (18)																						
															(15)	-----						
	Cost(sec)		%		Once(sec)		Start		End		Count		Nest		Routine name							
+-----+																						
	0.2467		50.06		0.0247		60		66		10		2		(sub1(int)._OMP_1)							
	0.2461		49.94		0.0002		61		62		1000		3		(sub1(int)._OMP_1)							
+-----+																						
	0.4928		100.00		-		-		-		-		-		-			(17)				
+-----+																						

- (1) Elapsed time
- (2) User CPU time
- (3) System CPU time
- (4) Information for each function that operates in serial processing (-Nrt\_tune\_func is specified)
- (5) Cost (sec)
- (6) Percentage of the entire cost
- (7) Cost called once
- (8) Start line number
- (9) End line number
- (10) Call frequency
- (11) Function name
- (12) Information for each function
- (13) Total value for each function
- (14) Information for each loop that operates in serial processing (-Nrt\_tune\_loop is specified)



(15) Nest level

(16) Information for each loop

(17) Total value for each loop

(18) Information for each loop that in parallel region by OpenMP (-Nrt\_tune\_loop is specified)

Note: When information is fetched for each loop, the generated function name is output as the function name. The demangle name is output. The function name is enclosed with parentheses at the loop that is nested.

Output example of runtime information (environmental variable FLIB\_RTINFO\_CSV is specified)

```
=====
EXECUTION PERFORMANCE INFORMATION
=====
Type,Elapsed(sec),User(sec),System(sec)
[COST],2.5579,4.1384,0.0030

Routine
Type,Cost(sec),%,Once(sec),Start,End,Count,Routine name
[COST_ROUTINE],1.9128,84.10,0.0010,73,88,2000,"sub2(int)"
[COST_ROUTINE],0.3617,15.90,0.0362,40,70,10,"sub1(int)"
[COST_ROUTINE],2.2745,100.00,-,-,-,-,-

Serial Loop
Type,Cost(sec),%,Once(sec),Start,End,Count,Nest,Routine name
[COST_LOOP_SERIAL],2.1887,45.20,2.1887,27,29,1,1,"main"
[COST_LOOP_SERIAL],2.1652,44.71,0.0011,79,80,2000,2,"(sub2(int))"
[COST_LOOP_SERIAL],0.2805,100.00,-,-,-,-,-

Parallel Loop
Type,Cost(sec),%,Once(sec),Start,End,Count,Nest,Routine name
[COST_LOOP_PRL],0.2463,50.05,0.0246,60,66,10,2,"(sub1(int))._OMP_1"
[COST_LOOP_PRL],0.2458,49.95,0.0002,61,62,1000,3,"(sub1(int))._OMP_1"
[COST_LOOP_PRL],2.2745,100.00,-,-,-,-,-
```

## G.3 Notes on Runtime Information Output Function

This section describes the notes on Runtime Information Output Function.

- If the information is fetched each function or loop by specifying -Nrt\_tune\_func or -Nrt\_tune\_loop, fetched information is increases, and execute time may increase. If the information cannot be obtained, the diagnostic message jwe1655i-i may be output. If this diagnostic message is output during error processing, the order in which the diagnostic messages are output may be incorrect.
- When there are goto statement, setjmp()/longjmp(), throw statement during a program or an exceptions is checked, information may be not correctly fetched.
- There is a possibility of receiving the following influences by optimization of the compiler.
  - When fetched the information for each loop, it may be different which loop is worked as a target of the measurement.
  - Line number may not accord with line number of the source program.
  - The nest relations of the loop may not accord with a nest of the source program. By the following optimization, I may not recognize loop information definitely. By the following optimization, it may not be correctly recognized loop information.
    - inline expansion for functions
    - loop unrolling
    - loop blocking
    - loop interchange
    - loop fusion
    - loop striping

- loop splitting
- loop peeling
- loop unswitching
- change the function to a multi-operation function, use SIMD instructions

# Appendix H Using High-Speed Facility

This appendix provides information on the compilation and execution processes for using the high-speed facility on FX system.

By using the following environment variables, the high-speed facility (the inter-core hardware barrier and the Sector cache) can be used.

## FLIB\_HPCFUNC

The environment variable `FLIB_HPCFUNC` specifies whether to use the high-speed facility (inter-core hardware barrier, Sector cache). If the value `TRUE` is assigned to `FLIB_HPCFUNC`, the high-speed facility can be used. If the value `FALSE` is assigned to that environment variable, this facility cannot be used. The default value is `FALSE`.

If this environment variable is not specified, the high-speed facility cannot be used.

Value	Explanation
TRUE	The high-speed facility (the inter-core hardware barrier and the Sector cache) can be used. This can be used when a NUMA node can be occupied by one process. The upper limit of CPU is the number of CPUs allocated to one node.
FALSE (Default)	The high-speed facility (the inter-core hardware barrier and the Sector cache) cannot be used. If the following environment variables is specified, an error message (jwe1047i-w) is output. - <code>FLIB_SCCR_CNTL=TRUE</code>

## FLIB\_HPCFUNC\_INFO

The environment variable `FLIB_HPCFUNC_INFO` specifies whether to display the CPU affinity of the process at the start of execution. The default value is `FALSE`.

If this environment variable is not specified, the CPU affinity of the process at the start of execution will not be displayed.

Value	Explanation
TRUE	The CPU affinity of the process at the start of execution is displayed.
FALSE (Default)	The CPU affinity of the process at the start of execution is not displayed.

## H.1 Inter-Core Hardware Barrier

The inter-core hardware barrier (hereafter called hardware barrier) is a hardware mechanism which facilitates high speed synchronization among threads in a CPU Chip, and raises the execution performance. This section describes compiling and executing programs to take advantage of hardware barriers using the LLVM OpenMP library.

### H.1.1 Compilation

When generating a program using the hardware barrier as the thread barrier, specify `-Kparallel` or `-Kopenmp` option.

### H.1.2 Execution

#### FLIB\_BARRIER

The barrier-kind of OpenMP and Automatic Parallelization can be controlled using the environment variable `FLIB_BARRIER`.

Allowable values and their meanings are as follows:

#### HARD

The hardware barrier is used.

When the hardware barrier cannot be used, the diagnostic message (jwe1050i-w) is output, and the processing is continued using the software barrier.

When the hardware barrier is used, the part or everything of the following OpenMP functions, cannot be used. When these OpenMP functions are used, the diagnostic message (jwe1051i-w) is output, and these functions are ignored, the processing is continued using the software barrier.

- The controlling number of threads

The number of threads will be determined using the following priority.

1. The value specified in the environment variable `OMP_NUM_THREADS`
2. The limit on the number of CPUs

When the number of threads exceeds the limit on the number of CPUs, the number of threads is the limit on the number of CPUs.

The number of threads specified with the `num_threads` clause or the `omp_set_num_threads` function must be the same as defined by the execution environment.

Irrespective of the environment variable `OMP_DYNAMIC` and the `omp_set_dynamic` function, the number of threads cannot exceed the limit on the number of CPUs.

- The controlling thread affinity

The thread affinity is set using CPUs assigned to the process when the program execution starts.

Irrespective of the environment variable `OMP_PROC_BIND` and the `proc_bind` clause of `parallel` construct, OpenMP thread affinity is ignored.

When the thread affinity is set using other ways, the operation is not guaranteed.

- Nested parallelism.

Irrespective of the environment variable `OMP_NESTED` and the `omp_set_nested` function, nested `parallel` regions are always serialized.

- Task

Always, the undeferred task is generated.

- Cancellation

Irrespective of the environment variable `OMP_CANCELLATION`, the `cancel` construct and cancellation points are effectively ignored.

#### SOFT

The software barrier is used. (default)

#### Note

- When the number of threads is 1, the hardware barrier cannot be used.
- The hardware barrier in FX system is used only inside the NUMA node. When one process is allocated to multiple NUMA nodes, hardware barrier is used inside NUMA node and software barrier is used between NUMA nodes. Therefore, it is recommended to allocate one process to one NUMA node, in case of giving priority to the performance of threads barrier.
- On the following cases, the high speed runtime library cannot determine whether one process can reserve a Barrier resource or not. Therefore behavior of hardware barrier is indeterminate, program may be ended with diagnostic-message (jwe1044i-u or jwe1045i-u), or the program may terminate abnormally.

- For multi-process programs

Example:

- FORTRAN FORK/SYSTEM/SH service function is used.
- C fork system call/system function is used.

## Example

Example: A program that uses the inter-core hardware barrier is executed.

```
$ cat exec.sh
#!/bin/sh
export FLIB_HPCFUNC=TRUE
export FLIB_BARRIER=HARD
./a.out
$ ./exec.sh
```

## H.2 Sector Cache

See Section "[3.5 Software Control of Sector Cache](#)" for details on how to control the Sector cache.

This section explains only notes on execution.

### Notes on execution

The Sector cache for the primary and secondary level caches are available in the following execution environments.

The execution environments	The Sector cache for the first level cache	The Sector cache for the second level cache
Only one process runs on a NUMA node	Available	Available
Multiple processes run on a NUMA node	Available (*1)	Not available

\*1) Not available if the environment variable `FLIB_L1_SCCR_CNTL` is assigned to the value `FALSE`. See Section "[3.5.2.2 Software Control with Environment Variables and Optimization Control Line](#)" for details.

On the following cases, the high speed runtime library cannot determine whether one process can reserve a NUMA node or not. Therefore behavior of Sector cache is indeterminate. Performance of program may be reduced, program may be ended with diagnostic-message (jwe1048i-u), or the program may terminate abnormally.

- For multi-process programs
- Example:
- FORTRAN FORK/SYSTEM/SH service function is used.
  - C fork system call/system function is used.
- Threads are generated by functions other than Fujitsu Automatic Parallelization/OpenMP.

Example:

- Threads are controlled by pthread functions.

To invalidate the Sector cache control function unconditionally regardless of the execution environment, "FALSE" is set for the environment variable `FLIB_SCCR_CNTL`.

See Section "[3.5.2.2 Software Control with Environment Variables and Optimization Control Line](#)" for details.

## Example

Example: A program that uses the Sector cache is executed.

```
$ cat exec.sh
#!/bin/sh
export FLIB_HPCFUNC=TRUE
```

```
./a.out  
$ ./exec.sh
```

---

# Appendix I Multi-Process Execution

This appendix describes how to execute a program in a multi-process and notes.

## I.1 Specifying CPU Affinity

For multi-process execution, the environment variable `FLIB_PROCESS_CPU_AFFINITY` sets the CPU affinity of the process.

The system binds threads to the specified CPU according to the thread affinity policy that is in effect when the program is started.

Environment Variable `FLIB_PROCESS_CPU_AFFINITY`

The CPUs specified in the `cpuid` list are assigned to the process.

The `cpuid` list shall be separated by comma (',') or space (' ').

The `cpuid` list can have the next form that has range with increment.

```
cpuid1[-cpuid2[:inc]]
```

*cpuid1*: cpuid of the beginning of the range. ( $0 \leq \text{cpuid1} < \text{CPU\_SETSIZE}$ )

*cpuid2*: cpuid of the end of the range. ( $0 \leq \text{cpuid2} < \text{CPU\_SETSIZE}$ )

*inc*: increment ( $1 \leq \text{inc} < \text{CPU\_SETSIZE}$ )

In addition, it is necessary to be the following.

```
cpuid1 <= cpuid2
```

It becomes equivalent to the case where all CPUs for every increment value *inc* in the range from *cpuid1* to *cpuid2* are specified.

The `cpuid` can be used the above-mentioned value. However, `cpuid` which can actually be assigned becomes only within the limits of CPU affinity of the process at the start of execution.

See `CPU_SET(3)` about details of `CPU_SETSIZE`.

### Note

- If `cpuid` is outside the CPU affinity of the process at the start of execution, the program may not work properly, for example, an error will be output and the program will terminate, or the program will not end. Correct the setting value. The CPU affinity of the process at the start of execution can be checked by specifying `FLIB_HPCFUNC_INFO=TRUE` or system commands such as `taskset` and `numactl`. See "[Appendix H Using High-Speed Facility](#)" about the environment variable `FLIB_HPCFUNC_INFO`. For system commands, refer to each man manual.
- If the environment variable `FLIB_HPCFUNC=TRUE` is not specified, the environment variable `FLIB_PROCESS_CPU_AFFINITY` is disabled. See "[Appendix H Using High-Speed Facility](#)" about the environment variable `FLIB_HPCFUNC`.
- The inter-core hardware barrier cannot be used when a program in multi-process using Fortran `FORK`, `SYSTEM`, `SH` service functions, C fork system calls, system library functions, and so on. Use the software barrier. See section "[H.1 Inter-Core Hardware Barrier](#)".
- If a value that is not specified to the environment variable `FLIB_PROCESS_CPU_AFFINITY` is specified to the environment variable `GOMP_CPU_AFFINITY`, an error message is output and the program terminates. For the environment variable `GOMP_CPU_AFFINITY`, see Section "[4.2.2.4 CPU Binding for Thread](#)" if using the automatic parallelism feature or "[4.3.2.1 Environment Variable at Execution](#)" if using the parallelism feature based on the OpenMP specification.

### Example

Examples of using environment variable `FLIB_PROCESS_CPU_AFFINITY`

- Example 1:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0,1,2,3,4,5,6,7"
```

The process is bound to CPU 0, 1, 2, 3, 4, 5, 6, and 7.

- Example 2:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0-7"
```

The process is bound to CPU 0, 1, 2, 3, 4, 5, 6, and 7.

- Example 3:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0-7:2"
```

The process is bound to CPU 0, 2, 4, and 6.

- Example 4:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0-5:2,3,8"
```

The process is bound to CPU 0, 2, 3, 4, and 8.

---



# Appendix J Fujitsu OpenMP Library

This appendix describes how parallel processing is performed for a C++ program using Fujitsu OpenMP Library in Trad Mode. See "[Chapter 4 Multiprocessing](#)" when you use LLVM OpenMP Library.

## Note

If all of the following conditions are met, this system uses Fujitsu OpenMP Library.

- -Nfjomplib option is effective at linking
- Trad Mode compiler is used at both compiling and linking

## Information

The table below shows a connectable combination of library for multiprocessing (LLVM OpenMP Library and Fujitsu OpenMP Library) and connectable object file.

	Fortran object file	C/C++ object file	
		Trad Mode (-Nnoclang option)	Clang Mode (-Nclang option)
LLVM OpenMP Library (-Nlibomp option)	Available	Available	Available
Fujitsu OpenMP Library (-Nfjomplib option)	Available	Available	Not available

## J.1 Overview of Multiprocessing

See Section "[4.1 Overview of Multiprocessing](#)".

## J.2 Automatic Parallelization

This section describes automatic parallelization.

### J.2.1 Compilation (Automatic Parallelization)

Following options are used to invoke automatic parallelization.

- Compile time

```
-Kparallel
```

- Link time

```
-Kparallel -Nfjomplib
```

#### Compiler Option for Automatic Parallelization

The compiler option for automatic parallelization is:

```
-K{parallel|parallel_strong|visimpact}  
[,array_private,dynamic_iteration,instance=N,loop_part_parallel,ocl,optmsg=2,parallel_fp_precision,p  
arallel_iteration=N,reduction,region_extension] -Nfjomplib
```

See Section "[2.2 Compiler Options](#)" for details on how to use each option.

## J.2.2 Execution Process (Automatic Parallelization)

---

When a multiprocessing program is executed, the number of threads to be used may be specified by the environment variable `PARALLEL` or `OMP_NUM_THREADS`. The size of the thread stack may be specified by the environment variable `THREAD_STACK_SIZE` or `OMP_STACKSIZE`. The method of synchronizing threads can be changed using the environment variable `FLIB_SPINWAIT` or `OMP_WAIT_POLICY`.

Except these environment variable settings, the procedure for execution is same as for a serial-processing program.

### Environment Variable `PARALLEL`

The number of threads to be used may be specified by the environment variable `PARALLEL`. Specify a number from 1 to 2147483647.

For more information on the number of threads, see "[Number of Threads](#)" in Section "[J.2.2 Execution Process \(Automatic Parallelization\)](#)".

### Number of Threads

The number of threads is determined with the following priority.

1. The value of the environment variable `PARALLEL`
2. The value of the environment variable `OMP_NUM_THREADS`
3. When the environment variable `FLIB_HPCFUNC=TRUE` is specified, the number of CPUs that can be used in the system, when the environment variable `FLIB_HPCFUNC=TRUE` is not specified, 1 thread



---

#### Value of `PARALLEL`

If the `-Kopenmp` option is specified at the linkage phase of the compilation and the environment variable `FLIB_FASTOMP` is set as `TRUE`, the value of the environment variable `OMP_NUM_THREADS` and `PARALLEL` must be the same. Otherwise, a message (`jwe1042i-w`) is output and the smaller value of `OMP_NUM_THREADS` and `PARALLEL` is adopted.

---

The number determined above is compared with the available maximum CPU number, and smaller number is used as the number of threads.

See "[Environment Variable for OpenMP Specifications](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)" for the `OMP_NUM_THREADS` environment variable.

### Environment Variable `THREAD_STACK_SIZE` and `OMP_STACKSIZE`

#### `THREAD_STACK_SIZE`

Specify the stack size (in Kbytes) for each thread using the environment variable `THREAD_STACK_SIZE`. Specify an integer value between from 1 to 18446744073709551615.

When the environment variable `OMP_STACKSIZE` is also specified, the greater value of the two is used as the stack area size for each thread. Refer to "[Stack Size on Execution](#)" in Section "[J.2.2 Execution Process \(Automatic Parallelization\)](#)" for information on stack area.

#### `OMP_STACKSIZE`

The stack area size for each thread can be specified using the environment variable `OMP_STACKSIZE` in byte, Kbytes, Mbytes, or Gbytes.

When the environment variable `THREAD_STACK_SIZE` is also specified, the greater value of the two is used as the stack area size. Refer to "[Stack Size on Execution](#)" in Section "[J.2.2 Execution Process \(Automatic Parallelization\)](#)" for information on stack area.

## Stack Size on Execution

Local variables in the range of a parallelized loop are allocated on the stack region. If there is not enough space for the stack area to allocate local variables, the stack area needs to be extended to an appropriate size.

When the stack area size for each thread needs to be specified, specify the environment variable `THREAD_STACK_SIZE` or `OMP_STACKSIZE`.

When the environment variable is not specified, the stack size of each thread is the same as the stack size of the process.

However, if the maximum stack size for the process is larger than the value calculated by the expression " $\text{min}(\text{the size of available memory in this system, the size of a virtual memory}) / (\text{the number of threads})$ ", the system will determine the stack area size and allocate it as follows:

```
the stack size for each thread (byte) =
(min(the size of available memory in this system, the size of a virtual memory) / (the number of
threads)) / 5
```

However, the assumed value does not guarantee that the system works correctly. It is recommended that an appropriate value for maximum process stack size be specified manually. The maximum process stack area size can be specified using the bash built in command "`ulimit`".

Note that the maximum virtual memory size can also be confirmed using `ulimit` (bash built-in command).

See "[Number of Threads](#)" in Section "[J.2.2 Execution Process \(Automatic Parallelization\)](#)" for details.



### Note

Some types of pthread library do not support setting the size of the stack area because of fixed stack. Static types of pthread library tend to be fixed stack, so ensure that you use a pthread library that is of a shared type.

## Synchronization Process

The synchronization process can be controlled with the environment variable `FLIB_SPINWAIT` or `OMP_WAIT_POLICY`.

`FLIB_SPINWAIT`

<code>unlimited</code>	Uses spin wait until all threads are synchronized. The default is <code>unlimited</code> .
<code>0</code>	Uses no spin wait but suspending wait.
<code>ns</code>	Uses spin wait until after $n$ seconds, and then uses suspending wait. $n$ is an integer value that must be positive. The unit "s" is specified following $n$ .
<code>rms</code>	Uses spin wait until after $n$ milliseconds, and then uses suspending wait. $n$ is an integer value that must be positive. The unit "ms" is specified following $n$ .

Spin wait is a synchronization method that allows all threads to consume CPU time, while suspending wait consumes less CPU time. Since spin wait is more efficient in parallel execution, "`unlimited`" is suitable for reducing elapsed time. On the other hand, "`0`" is suitable for reducing CPU time.

`OMP_WAIT_POLICY`

<code>ACTIVE</code>	Uses spin wait until all threads are synchronized. The default is <code>ACTIVE</code> .
<code>PASSIVE</code>	Uses no spin wait but suspending wait.

Select `ACTIVE` to give priority to the elapsed time. Select `PASSIVE` to give priority to the CPU time.

The settings of environment variable `FLIB_SPINWAIT` will be valid when environment variable `FLIB_SPINWAIT` is specified.

## Restriction of `fork(2)`

Functions that depend on `fork(2)` are restricted.

## CPU binding for thread

When the environment variable `FLIB_HPCFUNC=TRUE` is not set, you can control CPU binding for thread using environment variable `FLIB_CPU_AFFINITY`. For details of `FLIB_HPCFUNC`, see "[FLIB\\_HPCFUNC](#)".

The environment variable `FLIB_CPU_AFFINITY`

Threads are bound to CPUs in order of the specified `cpuid` list.

When the number of specified CPUs is exceeded, it is repeatedly used from the beginning of the list.

The `cpuid` list shall be separated by comma (',' ) or space (' ') .

The `cpuid` list can have the next form that has range with increment.

```
cpuid1[-cpuid2[:inc]]
```

*cpuid1*: `cpuid` of the beginning of the range. ( $0 \leq \text{cpuid1} < \text{CPU\_SETSIZE}$ )

*cpuid2*: `cpuid` of the end of the range. ( $0 \leq \text{cpuid2} < \text{CPU\_SETSIZE}$ )

*inc*: increment ( $1 \leq \text{inc} < \text{CPU\_SETSIZE}$ )

In addition, it is necessary to be the following.

```
cpuid1 <= cpuid2
```

It becomes equivalent to the case where all CPUs for every increment value *inc* in the range from *cpuid1* to *cpuid2* are specified.

The `cpuid` can be used the above-mentioned value.

However, `cpuid` which can actually be assigned becomes only within the limits of CPU affinity of the process at the start of execution.

See `CPU_SET(3)` about details of `CPU_SETSIZE`.

### Note

If `cpuid` is outside the CPU affinity of the process at the start of execution, the program outputs error messages. Correct the setting value. The CPU affinity of the process at the start of execution can be checked by system commands such as `taskset` and `numactl`. For system commands, refer to each man manual.

### Example

Example 1:

```
$ export FLIB_CPU_AFFINITY="12,14,13,15"
```

The thread is bound to CPU in order of 12, 14, 13, and 15.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

Example 2:

```
$ export FLIB_CPU_AFFINITY="12-19"
```

The thread is bound to CPU in order of 12, 13, 14, 15, 16, 17, 18, and 19.

When the number of threads is nine or more, it is repeatedly used from the beginning of the list.

Example 3:

```
$ export FLIB_CPU_AFFINITY="12-19:2"
```

The thread is bound to CPU in order of 12, 14, 16, and 18.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

Example 4:

```
$ export FLIB_CPU_AFFINITY="12-16:2,13,19"
```

The thread is bound to CPU in order of 12, 14, 16, 13, and 19.

When the number of threads is six or more, it is repeatedly used from the beginning of the list.

---

## J.2.3 Example of Compilation and Execution

---

Example 1:

```
$ FCCpx -Kparallel,reduction,ocl -Nfjomplib test1.cc
$ ./a.out
```

In this example, reduction parallelization optimization and the optimization control lines (OCL) are in effect during compilation. The number of threads using at the runtime is defined by the execution environment. For details about the number of threads, see "[Number of Threads](#)" in Section "[J.2.2 Execution Process \(Automatic Parallelization\)](#)".

Example 2:

```
$ FCCpx -Kparallel -Nfjomplib test2.cc
$ PARALLEL=2
$ export PARALLEL
$ ./a.out
$ PARALLEL=4
$ export PARALLEL
$ ./a.out
```

The environment variable PARALLEL is set to 2 and the program executes with two threads.

Next, the environment variable PARALLEL is set to 4 and the program executes with four threads.

## J.2.4 Performance Tuning

---

See Section "[4.2.4 Performance Tuning](#)".

## J.2.5 Automatic Parallelization

---

See Section "[4.2.5 Feature Details on Automatic Parallelization](#)".

## J.2.6 Optimization Control Line

---

See Section "[4.2.6 Optimization Control Line](#)".

## J.2.7 Notes on Automatic Parallelization

---

See Section "[4.2.7 Notes on Automatic Parallelization](#)".

## J.2.8 Linking with Other Multi-Thread Programs

---

This section describes the restrictions of Linking with other multi-thread programs.

OpenMP Program on this System

An object program of an OpenMP program created by this system can be linked with an object program created by this system.

Other Multi-thread Programs

An object program created by this system cannot be linked with an object program created by another multi-thread system.

However, the program may be linked with a pthread program when using the environment variable FLIB\_PTHREAD.

Environment Variable `FLIB_PTHREAD`

You can control the linking with a pthread program, using the environment variable `FLIB_PTHREAD`.

Valid values are as follows. Default value is 0.

Value	Explanation
0 (Default)	<p>The program can be linked with a pthread program which uses the pthread thread as a control thread.</p> <p>However, there are the following restrictions.</p> <ul style="list-style-type: none"> <li>- The pthread thread must use suspending wait.</li> <li>- The pthread program must not be compiled with <code>-Kopenmp</code> option or <code>-Kparallel</code> option. (*)</li> <li>- The routine which is compiled with <code>-Kopenmp</code> option or <code>-Kparallel</code> option must not be used in the pthread program. (*)</li> <li>- The Fortran routine must not be used in the pthread program. (*)</li> </ul>
1	<p>The program can be linked with a pthread program which executes parallel processing. And, the program can be linked with a pthread program which uses the pthread thread as a control thread.</p> <p>However, there are the following restrictions.</p> <ul style="list-style-type: none"> <li>- A pthread parallel processing, and an OpenMP or automatic parallel processing must be executed in order. (*) An OpenMP or automatic parallel processing must not be executed in a pthread parallel processing. And, a pthread parallel processing must not be executed in OpenMP or automatic parallel processing.</li> <li>- The pthread thread must use suspending wait.</li> <li>- The pthread program must not be compiled with <code>-Kopenmp</code> option or <code>-Kparallel</code> option. (*)</li> <li>- The routine which is compiled with <code>-Kopenmp</code> option or <code>-Kparallel</code> option must not be used in the pthread program. (*)</li> <li>- The Fortran routine must not be used in the pthread program. (*)</li> <li>- Fujitsu's math libraries must not be used in the pthread program. (The pthread program must not be compiled with <code>-SSL2BLAMP</code> option.) (*)</li> <li>- You must execute OpenMP parallel processing with two or more threads, before creating pthread threads. And, this number of OpenMP threads cannot be changed later.</li> </ul> <p>And, when this function is used, the following functions are effective.</p> <ul style="list-style-type: none"> <li>- <code>FLIB_SPINWAIT=0</code></li> <li>- <code>FLIB_CPUBIND=off</code></li> </ul> <p>Operation is not guaranteed when a value which is different in the above-mentioned environment variable is specified.</p> <p>See "<a href="#">Synchronization Process</a>" in Section "<a href="#">J.2.2 Execution Process (Automatic Parallelization)</a>" for the environment variable <code>FLIB_SPINWAIT</code>.</p> <p>See Section "<a href="#">J.4.1 CPU Binding</a>" for the environment variable <code>FLIB_CPUBIND</code>.</p>

\*) Operation is not guaranteed when the restriction function is used.

## J.3 Parallelization by OpenMP Specification

This section details parallelization as described in the OpenMP Specification. For more details about the OpenMP specification, refer to the OpenMP Architecture Review Board's website.

This system supports the following specifications.

Supported Specifications
OpenMP 3.1
Part of OpenMP 4.0 <sup>[a]</sup>
Part of OpenMP 4.5 <sup>[a]</sup>

[a] The following functions can be used:

- `simd` construct
- `declare simd` construct

### J.3.1 Compilation (OpenMP Parallelization)

For compiling source programs and linking, specify the following options to the compile command.

The specified option is different by whether to perform parallelization by the OpenMP specification.

Parallelization by the OpenMP specification	Specified option	
	Compile time	Link time
When using the parallelization and the SIMD Extensions	-Kopenmp	-Kopenmp -Nfjomplib
When using the SIMD Extensions only, not using the parallelization	-Kopenmp_simd	-

-: There is no specified option.

#### Compiler Option for OpenMP C++ Program

This section describes the compilation options used to compile OpenMP programs.

##### -Kopenmp

This option enables OpenMP directives and compiles the source programs.

The -Kopenmp option should also be specified at the linkage phase of an object file, when the object file has previously been compiled with this option.



#### Example

```
$ FCCpx a.cc -Kopenmp -c
$ FCCpx a.o -Kopenmp -Nfjomplib
```

##### -Kopenmp\_simd

-Kopenmp\_simd option enables only the `simd` construct and the `declare simd` construct of OpenMP and compiles the source programs.

##### -Nfjomplib

This option specifies to use Fujitsu OpenMP Library for OpenMP.

-Nfjomplib option is required at linking.

## Obtaining of Optimization Information

See Section "[4.3.1.2 Obtaining of Optimization Information](#)".

## Restriction of OpenMP programs

See Section "[4.3.1.3 Restriction of OpenMP programs](#)".

## J.3.2 Execution Process (OpenMP Parallelization)

---

The execution process of an OpenMP C++ program is the same as the procedure for a serial-processing program.

### Environment Variable at Execution

Environment variable `FLIB_FASTOMP`

The user can control the high-speed runtime library using the environment variable `FLIB_FASTOMP`. Allowable values and their meanings are as follows. The default value is `TRUE`, so that the user can use the high-speed runtime library.

- `TRUE`

The high-speed runtime library is adopted instead of the normal runtime library (default).

- `FALSE`

The normal runtime library is adopted.

The high-speed runtime library achieves a fast execution time by partly limiting the OpenMP specification.

Details of the high-speed runtime library are shown below:

- It is well optimized on the assumption that nested parallelism is inhibited and every thread is connected to a CPU respectively.

- If the value `TRUE` assigned to `FLIB_HPCFUNC`, the hardware barrier can be used. For details of `FLIB_HPCFUNC`, see "[FLIB\\_HPCFUNC](#)".

To use the high-speed runtime library, OpenMP programs are restricted as follows:

- The number of threads specified with the `num_threads` clause or the `omp_set_num_threads` function must be the same as defined by the execution environment.

To check the number of threads defined by the execution environment, see "[Notes during Execution](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)".

When the high-speed runtime library is used, control of parallel execution with OpenMP environment variables and the OpenMP API is restricted as follows:

- Irrespective of the environment variable `OMP_NESTED` and the `omp_set_nested` function, nested parallel regions are always serialized.

- Irrespective of the environment variable `OMP_DYNAMIC` and the `omp_set_dynamic` function, the number of threads cannot exceed the limit on the number of CPUs.

The limit on the number of CPUs is described in "[Notes during Execution](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)".

Environment variable `FLIB_SPINWAIT` and `OMP_WAIT_POLICY`

The environment variables `FLIB_SPINWAIT` and `OMP_WAIT_POLICY` control the synchronization process.

`FLIB_SPINWAIT`

<code>unlimited</code>	Uses spin wait until all threads are synchronized. The default is <code>unlimited</code> .
<code>0</code>	Uses no spin wait but suspending wait.
<code>ns</code>	Uses spin wait until after $n$ seconds, and then uses suspending wait. $n$ is an integer value that must be positive. The unit "s" is specified following $n$ .



<code>nms</code>	Uses spin wait until after <i>n</i> milliseconds, and then uses suspending wait.  <i>n</i> is an integer value that must be positive. The unit "ms" is specified following <i>n</i> .
------------------	---

Spin wait is a synchronization mode that allows all threads to consume CPU time, while suspending wait consumes less CPU time. Since spin wait is more efficient in parallel execution, "unlimited" is suitable for reducing elapsed time. On the other hand, "0" is suitable for reducing CPU time.

#### OMP\_WAIT\_POLICY

ACTIVE	Uses spin wait until all threads are synchronized.  The default is ACTIVE.
PASSIVE	Uses no spin wait but suspending wait.

Select ACTIVE to give priority to the elapsed time. Select PASSIVE to give priority to the CPU time.

The settings of environment variable `FLIB_SPINWAIT` will be valid when environment variable `FLIB_SPINWAIT` is specified.

#### Environment variable `OMP_PROC_BIND`

A user can control the CPU binding for thread using the environment variable `OMP_PROC_BIND`.

Valid values and their meanings are as follows:

Value	Meaning
TRUE	The threads are bound to CPUs in order of cpuid list. However, cpuid which can actually be assigned becomes only within the limits of CPU affinity of the process at the start of execution.
FALSE	The thread is not bound to a CPU

The environment variable `FLIB_CPUBIND` or `FLIB_CPU_AFFINITY` is given to priority more than the environment variable `OMP_PROC_BIND`.

However, the environment variable `OMP_PROC_BIND` is effective in the following case.

- The environment variable `FLIB_CPUBIND=off` is specified, and the environment variable `FLIB_CPU_AFFINITY` is not specified.

See "[FLIB\\_CPUBIND](#)" in Section "[J.4.1 CPU Binding](#)" for the environment variable `FLIB_CPUBIND`.

See "[Notes during Execution](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)" for the environment variable `FLIB_CPU_AFFINITY`.

#### Environment variable `THREAD_STACK_SIZE`

The size of the thread stack can be set for the environment variable `THREAD_STACK_SIZE` in Kbytes.

When the environment variable `OMP_STACKSIZE` is specified, the greater value is used as the stack area size for each thread.

See "[Notes during Execution](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)" for details.

#### Environment variable `OMP_STACKSIZE`

The size of the thread stack can be set for the environment variable `OMP_STACKSIZE` in Kbytes, Mbytes, or Gbytes.

When the environment variable `THREAD_STACK_SIZE` is specified, the greater value is used as the stack area size.

See "[Notes during Execution](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)" for details.

## Environment Variable for OpenMP Specifications

The user can set the following environment variables, which are in accordance with the OpenMP Specifications. For the details, refer to the OpenMP specifications.

Environment variable `OMP_SCHEDULE`

Sets the schedule type and chunk size for directives that have the schedule type at runtime.

Environment variable `OMP_NUM_THREADS`

Sets the number of threads to use during execution.

Environment variable `OMP_DYNAMIC`

Enables or disables dynamic thread adjustment.

Environment variable `OMP_PROC_BIND`

Specifies whether threads are bound to CPUs.

Environment variable `OMP_NESTED`

Enables or disables nested parallelism.

Environment variable `OMP_STACKSIZE`

Specifies the stack size for each thread.

Environment variable `OMP_WAIT_POLICY`

Specifies the behavior of the standby thread.

Environment variable `OMP_MAX_ACTIVE_LEVELS`

Specifies the maximum number of nested active parallel regions.

Environment variable `OMP_THREAD_LIMIT`

Specifies the maximum number of threads performed on an Open MP program.

## Notes during Execution

When executing programs created on this system, note the following.

### Variable allocation at execution

Programs created by this system allocate local variables and private variables in a function to the stack region.

If there is not enough space for the stack area to allocate local variables and private variables in the function, the stack area needs to be extended to an appropriate size.

When the stack area size for each thread needs to be specified, specify the environment variable `THREAD_STACK_SIZE` or `OMP_STACKSIZE`.

When the environment variable is not specified, the stack size of each thread is the same as the stack size of the process. However, if the maximum stack size for the process is larger than the value calculated by the expression " $\text{min}(\text{the size of available memory in this system, the size of a virtual memory}) / (\text{the number of threads})$ ", the system will determine the stack area size and allocate it as follows:

```
the stack size for each thread (byte) =
(min(the size of available memory in this system, the size of a virtual memory) / (the number of
threads)) / 5
```

However, the assumed value does not guarantee that the system works correctly. It is recommended that an appropriate value for maximum process stack size be specified manually. The maximum process stack area size can be specified using the bash built in command "`ulimit`".

Note that the maximum virtual memory size can also be confirmed using `ulimit` (bash built-in command).

To make the size of the stack area of each thread the same value, the function to change the number of threads dynamically is not considered. Therefore, the same number of threads for "`FLIB_FASTOMP=TRUE`" is used in this expression.

When the number of threads is dynamically changed, you should set the number of maximum threads to `OMP_NUM_THREADS`.

For more information, see "[Environment Variable for OpenMP Specifications](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)" for the environment variable `THREAD_STACK_SIZE` and `OMP_STACKSIZE`.

## Limits on the Number of CPUs

The maximum number of CPUs is the limit on the number of CPUs on the system.

## Number of Threads

When the high-speed runtime library is used, a common number of threads is used for parallel execution of both the OpenMP program and automatic parallelization, otherwise the numbers of threads are determined independently of each other. Use of the high-speed runtime library is controlled by setting the environment variable `FLIB_FASTOMP` (described in "[Environment Variable for OpenMP Specifications](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)").

When the high-speed runtime library is used (`FLIB_FASTOMP` is `TRUE`):

The number of threads is determined with the following priority:

1. The value of environment variable `OMP_NUM_THREADS`
2. The value of environment variable `PARALLEL`
3. When the environment variable `FLIB_HPCFUNC=TRUE` is specified, the number of CPUs that can be used in the system, when the environment variable `FLIB_HPCFUNC=TRUE` is not specified, 1 thread

The number of threads may not exceed the limit on the number of CPUs. If this occurs, the number of threads becomes equal to the limit on the number of CPUs.

The numbers of threads specified in the program with the `num_threads` clauses of the `parallel` directives and the `omp_set_num_threads` function must adhere to this limitation. If a different number of threads is specified in the program, the execution is stopped with the following message:

```
jwe1041i-s The number of thread cannot be changed when the environment variable FLIB_FASTOMP is TRUE.
```

When the high-speed runtime library is not used (`FLIB_FASTOMP` is `FALSE`):

The number of threads used for OpenMP code is determined with the following priority (while the one for automatic parallelization is described in Section "[J.2 Automatic Parallelization](#)"):

1. The value specified with the `num_threads` clause of the "parallel" directive
2. The value specified with the `omp_set_num_threads` function
3. The value of the environment variable `OMP_NUM_THREADS`
4. The value of the environment variable `PARALLEL`
5. When the environment variable `FLIB_HPCFUNC=TRUE` is specified, the number of CPUs that can be used in the system, when the environment variable `FLIB_HPCFUNC=TRUE` is not specified, 1 thread

When the dynamic thread adjustment feature is enabled, the actual number of threads is the minimum number of threads determined by the rule above and the limits on the number of CPUs.

When the dynamic thread adjustment feature is disabled, multiple threads can be executed on the same processor via time-sharing. Such parallel execution is not effective because the overhead of synchronization among threads is very high. Therefore, it is recommended to set the number of threads to be less than the limit on the number of CPUs.

## Restriction of `fork(2)`

Functions that depend on `fork(2)` are restricted.

## CPU binding for thread

When the environment variable `FLIB_HPCFUNC=TRUE` is not set, you can control CPU binding for thread using environment variable `FLIB_CPU_AFFINITY`. For details of `FLIB_HPCFUNC`, see "[FLIB\\_HPCFUNC](#)".

The environment variable `FLIB_CPU_AFFINITY`

Threads are bound to CPUs in order of the specified cpuid list.

When the number of specified CPUs is exceeded, it is repeatedly used from the beginning of the list.

The cpuid list shall be separated by comma (',' ) or space (' ').

The cpuid list can have the next form that has range with increment.

```
cpuid1[-cpuid2[:inc]]
```

*cpuid1*: cpuid of the beginning of the range. ( $0 \leq \text{cpuid1} < \text{CPU\_SETSIZE}$ )

*cpuid2*: cpuid of the end of the range. ( $0 \leq \text{cpuid2} < \text{CPU\_SETSIZE}$ )

*inc*: increment ( $1 \leq \text{inc} < \text{CPU\_SETSIZE}$ )

In addition, it is necessary to be the following.

```
cpuid1<=cpuid2
```

It becomes equivalent to the case where all CPUs for every increment value *inc* in the range from *cpuid1* to *cpuid2* are specified.

The cpuid can be used the above-mentioned value.

However, cpuid which can actually be assigned becomes only within the limits of CPU affinity of the process at the start of execution.

See CPU\_SET(3) about details of CPU\_SETSIZE.

### Note

If cpuid is outside the CPU affinity of the process at the start of execution, the program outputs error messages. Correct the setting value. The CPU affinity of the process at the start of execution can be checked by system commands such as `taskset` and `numactl`. For system commands, refer to each man manual.

### Example

Example 1:

```
$ export FLIB_CPU_AFFINITY="12,14,13,15"
```

The thread is bound to CPU in order of 12, 14, 13, and 15.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

Example 2:

```
$ export FLIB_CPU_AFFINITY="12-19"
```

The thread is bound to CPU in order of 12, 13, 14, 15, 16, 17, 18, and 19.

When the number of threads is nine or more, it is repeatedly used from the beginning of the list.

Example 3:

```
$ export FLIB_CPU_AFFINITY="12-19:2"
```

The thread is bound to CPU in order of 12, 14, 16, and 18.

When the number of threads is five or more, it is repeatedly used from the beginning of the list.

Example 4:

```
$ export FLIB_CPU_AFFINITY="12-16:2,13,19"
```

The thread is bound to CPU in order of 12, 14, 16, 13, and 19.

When the number of threads is six or more, it is repeatedly used f from the beginning of the list.

## Output from More than One Threads

When errors are detected from more than one thread within a program, traceback maps are output in thread units. Therefore, in the parallel region, the information is displayed at the same number as threads.

### J.3.3 Implementation-Dependent Specifications

The OpenMP specification includes the following processor-dependent specification. For details on each item, refer to the OpenMP Specifications.

#### Memory Model

If a variable is longer than 4 bytes or crosses a 4-byte boundary:

- The value of the variable written from two threads without synchronization may be undefined.
- The value read by a thread from a variable that is written from another thread without synchronization may be undefined.

#### Internal Control Variables

The initial values for each of the internal control variables are as follows:

Internal Control Variable	Initial Value
<code>nthreads-var</code>	1
<code>dyn-var</code>	TRUE
<code>run-sched-var</code>	static without chunk size
<code>def-sched-var</code>	static without chunk size
<code>bind-var</code>	FALSE
<code>stacksize-var</code>	the stack area size for threads in " <a href="#">Notes during Execution</a> " in Section " <a href="#">J.3.2 Execution Process (OpenMP Parallelization)</a> "
<code>wait-policy-var</code>	ACTIVE
<code>thread-limit-var</code>	2147483647
<code>max-active-levels-var</code>	2147483647

#### Dynamic Thread Adjustment Features

This system supports the dynamic thread adjustment described in Section "[Notes during Execution](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)".

By default, the dynamic adjustment features are on.

#### Loop Statement

The type used to calculate the iteration count of a loop that is not nested is `long int`.

When `auto` is set for the internal control variable `run-sched-var`, the effect of the `schedule(runtime)` construct is set to `schedule(static)`.

#### sections Construct

Assignment to a thread of a structured block in a `sections` construct is performed in the same way as a `dynamic schedule`.

Each thread that encounters the `sections` region or that finishes the execution of a structured block in the `sections` region is moved on the next structured block.

#### single Construct

A `single` region is executed by the thread that encounters the region first.

#### simd Construct

The integer type used to compute the iteration count of a collapsed loop is `long int`.

SIMD length is the value to which the compiler decides automatically.

When the `alignment` parameter is not specified in the `aligned` clause, it is considered that the following value was specified for the `alignment` parameter.

When -KSVE option is effective at the compile time	Alignment of the type of the list item
When -KNOSVE option is effective at the compile time	16

#### declare simd Construct

SIMD length when the `simdlen` clause is not specified is as follows.

- When -KSVE option is effective at the compile time

SIMD length is decided at the run time.

- When -KNOSVE option is effective at the compile time

SIMD length is decided by the size of the minimum type among all input parameters and return value.

SIMD length is as follows by the size of the type.

Size of the type (byte)	SIMD length
1	16, 8
2	8, 4
4	4, 2
8	2
16	

When the `alignment` parameter is not specified in the `aligned` clause, it is considered that the following value was specified for the `alignment` parameter.

When -KSVE option is effective at the compile time	Alignment of the type of the list item
When -KNOSVE option is effective at the compile time	16

#### Task scheduling points

In untied task regions, task scheduling points are located at the same positions as in the case of tied task regions: only in `task`, `taskwait`, explicit or implicit `barrier` constructs, and at the completion point of the task.

#### atomic Construct

Two `atomic` regions will be executed independently (not exclusively), when a variable type to be updated is different. When the types match, it may be executed exclusively even if the address is different.

- When updating a logical type, complex type, 1 byte integer type, 2 byte integer type, or quadruple precision (16 byte) real type variable
- When updating array element and index expression is not the same
- The target expression has explicit or implicit type conversion

#### omp\_set\_num\_threads Function

Calling the `omp_set_num_threads` function has no effect when the argument value is equal to or less than 0. A value that exceeds the number of threads supported by the system must not be specified.

#### omp\_set\_schedule Function

There are no schedule types that are implementation-dependent.

#### omp\_set\_max\_active\_levels Function

A call to the `omp_set_max_active_levels` function will be ignored when it is performed from an explicit `parallel` region. It will also be ignored when the argument is an integer that less than 0.

#### omp\_get\_max\_active\_levels Function

The `omp_get_max_active_levels` function can be called from anywhere in the program and it returns the value of the internal control variable `max-active-levels-var`.

#### Environment Variable OMP\_SCHEDULE

When the schedule type specified for OMP\_SCHEDULE is invalid, the schedule type is ignored, and the default value (`static` without chunk size) is used.

When the schedule type specified for OMP\_SCHEDULE is `static`, `dynamic`, or `guided`, and the chunk size is not a positive number, the chunk size will be as follows:

Schedule Type	Chunk Size
<code>static</code>	no chunk size
<code>dynamic</code>	1
<code>guided</code>	

#### Environment Variable OMP\_NUM\_THREADS

When a value equal to or less than 0 is specified for the list of OMP\_NUM\_THREADS, it works in the same ways as when 1 is specified. A value that exceeds the number of threads supported by the system must not be specified.

#### Environment Variable OMP\_PROC\_BIND

When a value other than `TRUE` or `FALSE` is specified for OMP\_PROC\_BIND, the value is ignored, and the default value (`FALSE`) is used.

#### Environment Variable OMP\_DYNAMIC

When a value other than `TRUE` or `FALSE` is specified for OMP\_DYNAMIC, the value is ignored and the default value (`TRUE`) is used.

#### Environment Variable OMP\_NESTED

When a value other than `TRUE` or `FALSE` is specified for OMP\_NESTED, it is ignored and the default value (`FALSE`) is used.

#### Environment Variable OMP\_STACKSIZE

When the value specified for OMP\_STACKSIZE does not meet the defined format, it is ignored and the default value is used.

#### Environment Variable OMP\_WAIT\_POLICY

`ACTIVE` performs spin wait. `PASSIVE` performs suspend wait.

#### Environment Variable OMP\_MAX\_ACTIVE\_LEVELS

When the value specified for OMP\_MAX\_ACTIVE\_LEVELS is an integer that is less than 0, it is ignored and the default value (2147483647) is used.

#### Environment Variable OMP\_THREAD\_LIMIT

When the value specified for OMP\_THREAD\_LIMIT is not a positive integer, it is ignored and the default value (2147483647) is used.

## J.3.4 Notes on OpenMP Programming

---

This section provides notes on OpenMP C++ programming.

### Implementation of `parallel` Region and `explicit task` Region

The structured block within a `parallel` construct or a `task` construct is compiled as an internal function.

The name of the internal function generated from the `parallel` construct or the `task` construct is different depending on the compiler mode. The name of an internal function is shown below.

Type of internal function	Internal function name
Internal function generated from <code>parallel</code> construct	The name " <code>._OMP_id-numberA</code> " is added
Internal function generated from <code>task</code> construct	The name " <code>._TSK_id-numberB</code> " is added

*id-numberA*: Unique number that identifies `parallel` construct in a source file.

*id-numberB*: Unique number that identifies `task` construct in a source file.

## Implementation of `threadprivate` Variable

See Section "[4.3.4.2 Implementation of `threadprivate` Variable](#)".

## Automatic Parallelization for OpenMP Programs

See Section "[4.3.4.3 Automatic Parallelization for OpenMP Programs](#)".

## OpenMP directive

See Section "[4.3.4.4 OpenMP Directive](#)".

## J.3.5 Linking with Other Multi-Thread Programs

---

The following restrictions apply when linking with other multi-thread programs.

### Automatic Parallelization Program on this System

An object program of an automatic parallelization program created by this system can be linked with an object program created by this system.

### Other Multi-thread Programs

An object program created by this system cannot be linked with an object program created by another multi-thread system.

However, the program may be linked with a `pthread` program when using the environment variable `FLIB_PTHREAD`.

### Environment Variable `FLIB_PTHREAD`

You can control the linking with a `pthread` program, using the environment variable `FLIB_PTHREAD`.

Valid values are as follows. Default value is 0.

Value	Explanation
0 (Default)	<p>The program can be linked with a <code>pthread</code> program which uses the <code>pthread</code> thread as a control thread.</p> <p>However, there are the following restrictions.</p> <ul style="list-style-type: none"><li>- The <code>pthread</code> thread must use suspending wait.</li><li>- The <code>pthread</code> program must not be compiled with <code>-Kopenmp</code> option or <code>-Kparallel</code> option. (*)</li><li>- The routine which is compiled with <code>-Kopenmp</code> option or <code>-Kparallel</code> option must not be used in the <code>pthread</code> program. (*)</li><li>- The Fortran routine must not be used in the <code>pthread</code> program. (*)</li></ul>
1	<p>The program can be linked with a <code>pthread</code> program which executes parallel processing.</p> <p>And, the program can be linked with a <code>pthread</code> program which uses the <code>pthread</code> thread as a control thread.</p> <p>However, there are the following restrictions.</p> <ul style="list-style-type: none"><li>- A <code>pthread</code> parallel processing, and an OpenMP or automatic parallel processing must be executed in order. (*)</li></ul> <p>An OpenMP or automatic parallel processing must not be executed in a <code>pthread</code> parallel processing.</p> <p>And, a <code>pthread</code> parallel processing must not be executed in OpenMP or automatic parallel processing.</p> <ul style="list-style-type: none"><li>- The <code>pthread</code> thread must use suspending wait.</li></ul>



Value	Explanation
	<ul style="list-style-type: none"> <li>- The pthread program must not be compiled with -Kopenmp option or -Kparallel option. (*)</li> <li>- The routine which is compiled with -Kopenmp option or -Kparallel option must not be used in the pthread program. (*)</li> <li>- The Fortran routine must not be used in the pthread program. (*)</li> <li>- Fujitsu's math libraries must not be used in the pthread program. (The pthread program must not be compiled with -SSL2BLAMP option.) (*)</li> <li>- You must execute OpenMP parallel processing with two or more threads, before creating pthread threads. And, this number of OpenMP threads cannot be changed later.</li> </ul> <p>And, when this function is used, the following functions are effective.</p> <ul style="list-style-type: none"> <li>- FLIB_SPINWAIT=0</li> <li>- FLIB_CPUBIND=off</li> </ul> <p>Operation is not guaranteed when a value which is different in the above-mentioned environment variable is specified.</p> <p>See "<a href="#">Environment Variable at Execution</a>" in Section "<a href="#">J.3.2 Execution Process (OpenMP Parallelization)</a>" for the environment variable FLIB_SPINWAIT.</p> <p>See Section "<a href="#">J.4.1 CPU Binding</a>" for the environment variable FLIB_CPUBIND.</p>

\*) Operation is not guaranteed when the restriction function is used.

## J.3.6 Debugging OpenMP Programs

See Section "[Chapter 8 Debugging Functions](#)" for the debug feature in this system.

The following restrictions are applied to the debugging of the OpenMP programs:

- A function generated by the `declare simd` construct cannot be debugged using the debugger.

## J.4 Using High-Speed Facility

This section provides information on the compilation and execution processes for using the high-speed facility on FX system.

By using the following environment variables, the high-speed facility (the inter-core hardware barrier and the Sector cache) can be used.

FLIB\_HPCFUNC

The environment variable FLIB\_HPCFUNC specifies whether to use the high-speed facility (inter-core hardware barrier, Sector cache). If the value TRUE is assigned to FLIB\_HPCFUNC, the high-speed facility can be used. If the value FALSE is assigned to that environment variable, this facility cannot be used. The default value is FALSE.

If this environment variable is not specified, the high-speed facility cannot be used.

Value	Explanation
TRUE	<p>The high-speed facility (the inter-core hardware barrier and the Sector cache) can be used.</p> <p>This can be used when a NUMA node can be occupied by one process.</p> <p>The upper limit of CPU is the number of CPUs allocated to one node.</p>
FALSE (Default)	<p>The high-speed facility (the inter-core hardware barrier and the Sector cache) cannot be used.</p> <p>If the following environment variables is specified, an error message (jwe1047i-w) is output.</p> <ul style="list-style-type: none"> <li>- FLIB_SCCR_CNTL=TRUE</li> </ul>

## FLIB\_HPCFUNC\_INFO

The environment variable `FLIB_HPCFUNC_INFO` specifies whether to display the CPU affinity of the process at the start of execution. The default value is `FALSE`.

If this environment variable is not specified, the CPU affinity of the process at the start of execution will not be displayed.

Value	Explanation
<code>TRUE</code>	The CPU affinity of the process at the start of execution is displayed.
<code>FALSE</code> (Default)	The CPU affinity of the process at the start of execution is not displayed.

## J.4.1 CPU Binding

---

If the value `TRUE` is assigned to `FLIB_HPCFUNC`, the thread is bound to one CPU when the program using the automatic parallelization or OpenMP is executed.

See "[J.2 Automatic Parallelization](#)" and "[J.3 Parallelization by OpenMP Specification](#)" for details.

### FLIB\_CPUBIND

The environment variable `FLIB_CPUBIND` controls the CPU binding for threads.

The following value can be set to `FLIB_CPUBIND`. The default value is `"chip_pack"`.

<code>chip_pack</code>	All threads are bound to the same CPU Chip as much as possible.
<code>chip_unpack</code>	Each thread is bound to a different CPU Chip as much as possible.
<code>off</code>	The thread is not bound to the CPU. If you control CPU binding by yourself, this must be set.

When only one CPU Chip can be used, `"chip_pack"` and `"chip_unpack"` have the same effect.

When only one CPU on a CPU Chip can be used, `"chip_pack"` and `"chip_unpack"` have the same effect.

If the value `TRUE` is not assigned to `FLIB_HPCFUNC`, the `FLIB_CPUBIND` setting has no effect.

## J.4.2 Inter-Core Hardware Barrier

---

FX system has the inter-core hardware barrier. If the value `TRUE` is assigned to `FLIB_HPCFUNC`, it can be used as the thread barrier. The inter-core hardware barrier is a hardware mechanism which facilitates high speed synchronization among threads in a CPU Chip, and raises the execution performance.

### Compilation

When generating a program using the inter-core hardware barrier as the thread barrier, specify `-Kparallel` or `-Kopenmp` option.

The hardware barrier is automatically used during execution in an environment in which it can be used. By contrast, the software barrier is used during execution in an environment that cannot use the hardware barrier.

### Execution

The inter-core hardware barrier can be used only when the value `TRUE` is assigned to `FLIB_HPCFUNC`. The inter-core hardware barrier is a hardware mechanism for synchronization among threads in a CPU Chip.

#### FLIB\_CNTL\_BARRIER\_ERR

When the inter-core hardware barrier cannot be used, the following message is output, and the processing is continued using the software barrier.

```
jwe1050i-w The hardware barrier couldn't be used and continues processing using the software barrier.
```

This diagnostic message (jwe1050i-w) error can be controlled using the environment variable `FLIB_CNTL_BARRIER_ERR`.

The allowed values and their meanings are as follows. Default is "TRUE".

- TRUE

The diagnostic message (jwe1050i-w) error is detected.

When the inter-core hardware barrier cannot be used, this message is output, and the processing is continued using the software barrier.

- FALSE

The diagnostic message (jwe1050i-w) error is not detected.

When the inter-core hardware barrier cannot be used, the processing is continued using the software barrier.

#### FLIB\_NOHARDBARRIER

When the environment variable FLIB\_NOHARDBARRIER is set, the inter-core hardware barrier is not used and the software barrier is used. Always outputs the diagnostic message jwe1050i-w.

See "[FLIB\\_CNTL\\_BARRIER\\_ERR](#)" in Section "[J.4.2 Inter-Core Hardware Barrier](#)" for details.

### Note

- When the inter-core hardware barrier is used by a program made using the compile option -Kopenmp, the environment variable FLIB\_FASTOMP shall not be FALSE. See Section "[J.3 Parallelization by OpenMP Specification](#)" for the meaning and use of the environment variable FLIB\_FASTOMP.
  - When the inter-core hardware barrier is used, CPU binding is needed. Therefore, when "off" is set for the environment variable FLIB\_CPUBIND, the inter-core hardware barrier cannot be used. See Section "[J.4.1 CPU Binding](#)" for details of CPU binding.
  - When the number of threads is 1, the inter-core hardware barrier cannot be used. See "[Number of Threads](#)" in Section "[J.2.2 Execution Process \(Automatic Parallelization\)](#)" or "[Notes during Execution](#)" in Section "[J.3.2 Execution Process \(OpenMP Parallelization\)](#)" for the number of threads.
  - When the number of CPUs to a process is three or less, an inter-core hardware barrier may not be able to use. The number of CPUs to a process should be set to four or more.
  - The inter-core hardware barrier in FX system is used only inside the NUMA node. When one process is allocated to multiple NUMA nodes, inter-core hardware barrier is used inside NUMA node and software barrier is used between NUMA nodes. Therefore, it is recommended to allocate one process to one NUMA node, in case of giving priority to the performance of threads barrier.
  - On the following cases, the high speed runtime library cannot determine whether one process can reserve a barrier resource or not. Therefore behavior of hardware barrier is indeterminate, program may be ended with diagnostic-message (jwe1044i-u or jwe1045i-u), or the program may terminate abnormally.
    - For multi-process programs
- Example:
- FORTRAN FORK/SYSTEM/SH service function is used
  - C fork system call/system function is used

### Example

Example: A program that uses the inter-core hardware barrier is executed.

```
$ cat exec.sh
#!/bin/sh
export FLIB_HPCFUNC=TRUE
./a.out
$ ./exec.sh
```

## J.4.3 Sector Cache

---

FX system has Sector cache. If the value TRUE is assigned to FLIB\_HPCFUNC, the Sector cache can be controlled.

See Section "[3.5 Software Control of Sector Cache](#)" for details on how to control the Sector cache.

This section explains only notes on execution.

### Notes on execution

The Sector cache for the primary and secondary level caches are available in the following execution environments.

The execution environments	The Sector cache for the first level cache	The Sector cache for the second level cache
Only one process runs on a NUMA node	Available	Available
Multiple processes run on a NUMA node	Available (*1)	Not available

\*1) Not available if the environment variable FLIB\_L1\_SCCR\_CNTL is assigned to the value FALSE. See Section "[3.5.2.2 Software Control with Environment Variables and Optimization Control Line](#)" for details.

On the following cases, the high speed runtime library cannot determine whether one process can reserve a NUMA node or not. Therefore behavior of Sector cache is indeterminate. Performance of program may be reduced, program may be ended with diagnostic-message (jwe1048i-u) , or the program may terminate abnormally.

- For multi-process programs

Example:

- FORTRAN FORK/SYSTEM/SH service function is used.
- C fork system call/system function is used.
- Threads are generated by functions other than Fujitsu Automatic Parallelization/OpenMP.

Example:

- Threads are controlled by pthread functions.

To invalidate the Sector cache control function unconditionally regardless of the execution environment, "FALSE" is set for the environment variable FLIB\_SCCR\_CNTL.

See Section "[3.5.2.2 Software Control with Environment Variables and Optimization Control Line](#)" for details.



### Example

Example: A program that uses the Sector cache is executed.

```
$ cat exec.sh
#!/bin/sh
export FLIB_HPCFUNC=TRUE
./a.out
$ ./exec.sh
```

## J.5 Multi-Process Execution

---

This section describes how to execute a program in a multi-process and notes.

### J.5.1 Specifying CPU Affinity

---

For multi-process execution, the environment variable FLIB\_PROCESS\_CPU\_AFFINITY sets the CPU affinity of the process.

The system binds threads to the specified CPU according to the thread affinity policy that is in effect when the program is started.

#### Environment Variable `FLIB_PROCESS_CPU_AFFINITY`

The CPUs specified in the cpuid list are assigned to the process.

The cpuid list shall be separated by comma (',') or space (' ').

The cpuid list can have the next form that has range with increment.

```
cpuid1[-cpuid2[:inc]]
```

*cpuid1*: cpuid of the beginning of the range. ( $0 \leq \text{cpuid1} < \text{CPU\_SETSIZE}$ )

*cpuid2*: cpuid of the end of the range. ( $0 \leq \text{cpuid2} < \text{CPU\_SETSIZE}$ )

*inc*: increment ( $1 \leq \text{inc} < \text{CPU\_SETSIZE}$ )

In addition, it is necessary to be the following.

```
cpuid1 <= cpuid2
```

It becomes equivalent to the case where all CPUs for every increment value *inc* in the range from *cpuid1* to *cpuid2* are specified.

The cpuid can be used the above-mentioned value. However, cpuid which can actually be assigned becomes only within the limits of CPU affinity of the process at the start of execution.

See `CPU_SET(3)` about details of `CPU_SETSIZE`.

#### Note

- If cpuid is outside the CPU affinity of the process at the start of execution, the program may not work properly, for example, an error will be output and the program will terminate, or the program will not end. Correct the setting value. The CPU affinity of the process at the start of execution can be checked by specifying `FLIB_HPCFUNC_INFO=TRUE` or system commands such as `taskset` and `numactl`. See section "[J.4 Using High-Speed Facility](#)" about the environment variable `FLIB_HPCFUNC_INFO`. For system commands, refer to each man manual.
- If the environment variable `FLIB_HPCFUNC=TRUE` is not specified, the environment variable `FLIB_PROCESS_CPU_AFFINITY` is disabled. See section "[J.4 Using High-Speed Facility](#)" about the environment variable `FLIB_HPCFUNC`.
- The inter-core hardware barrier cannot be used when a program in multi-process using Fortran `FORK`, `SYSTEM`, `SH` service functions, C fork system calls, system library functions, and so on. Use the software barrier. See section "[J.4.2 Inter-Core Hardware Barrier](#)".

#### Example

##### Examples of using environment variable `FLIB_PROCESS_CPU_AFFINITY`

- Example 1:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0,1,2,3,4,5,6,7"
```

The process is bound to CPU 0, 1, 2, 3, 4, 5, 6, and 7.

- Example 2:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0-7"
```

The process is bound to CPU 0, 1, 2, 3, 4, 5, 6, and 7.

- Example 3:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0-7:2"
```

The process is bound to CPU 0, 2, 4, and 6.

- Example 4:

```
$ export FLIB_PROCESS_CPU_AFFINITY="0-5:2,3,8"
```

The process is bound to CPU 0, 2, 3, 4, and 8.

---

# Appendix K Large Page Library

This appendix describes the large page library that is part of the HPC extension function.

Some application programs use a large amount of memory in the HPC area. The cost of memory management by the OS affects the execution performance of these application programs. To reduce this cost, the HPC extension function provides an original extended function for large pages.

The appendix also describes the basic memory allocation ideas needed to understand the large page library, and more.

## K.1 Overview of the Memory Allocation Function

This section describes the following function and processing as an overview of the memory allocation function, which is an extension function of the FX system:

- Large page function
- Paging methods

### K.1.1 Large Page Function

The large page function extends HugeTLBfs, a standard function in Linux, and allocates memory (large page) with a larger page size than a normal page to application programs that handle large amounts of data. By doing so, it reduces the cost of address translation processing by the OS and increases memory access performance.

#### Memory Address Translation and TLB

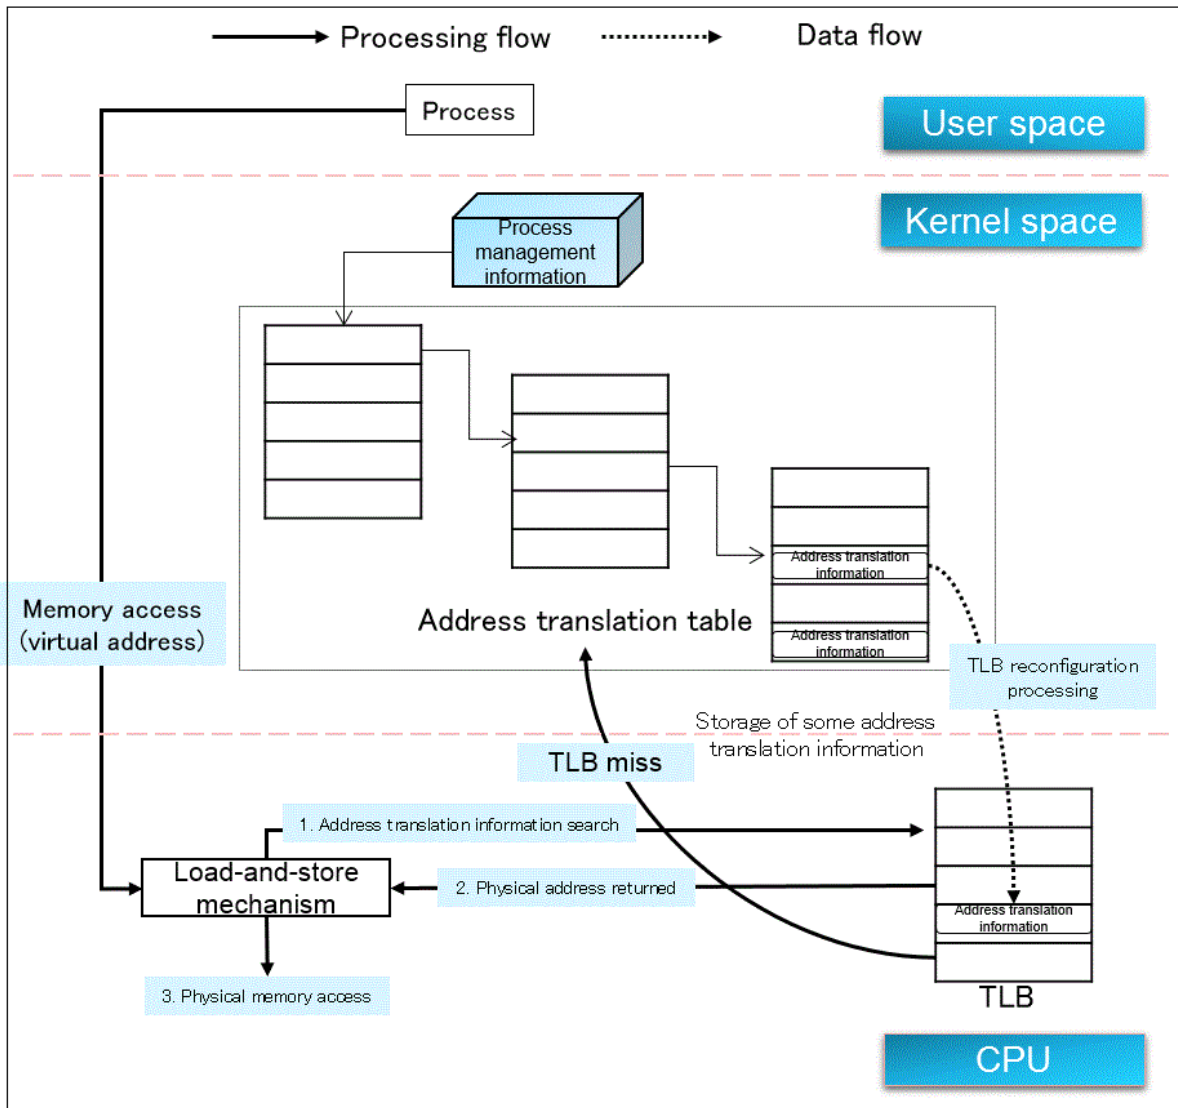
Application programs use virtual memory addresses to access memory. Therefore, virtual memory addresses must be translated into physical memory addresses. An address translation table in main memory is used to translate the addresses. At this time, to further increase the access speed, the load-and-store mechanism in the CPU uses an address translation buffer called the TLB (Translation Look-aside Buffer) to search for a physical address.

The load-and-store mechanism in the CPU receives load-and-store requests for memory from running application programs. "[Figure K.1 Translation From a Virtual Memory Address to a Physical Memory Address](#)" shows an outline of memory address translation using the load-and-store mechanism.

1. The mechanism searches for the physical memory address corresponding to the virtual memory address specified as the load-and-store destination from the TLB.  
If the TLB does not have corresponding address translation information, a TLB miss occurs and the TLB searches for and reads in the address translation information from the address translation table.
2. The mechanism obtains the physical memory address from the TLB using the corresponding address translation information.
3. Physical memory access to the load-and-store destination begins.

A TLB miss occurs when the physical memory address is not found in the TLB. If a TLB miss occurs, TLB reconfiguration processing (reading the address translation table again) is necessary. Generally, the cost (time) required for this processing is large.

Figure K.1 Translation From a Virtual Memory Address to a Physical Memory Address



### Normal Page and Large Page


The OS of the FX system can use page sizes of 64 KiB for a normal page and 2 MiB for a large page. The large page function has large pages enabled by default. You can select to enable or disable large pages by using an environment variable (XOS\_MMM\_L\_HPAGE\_TYPE) described in the section titled "K.3 Environment Variables for Large Page Library Settings".

If an application program handles a large amount of memory with normal pages, it accesses many pages, and TLB reconfiguration processing tends to occur frequently. If the application program uses large pages, it accesses fewer pages as compared to normal pages, and TLB reconfiguration processing occurs less frequently.

This means you can reduce TLB misses by allocating large pages to application programs that use a large amount of memory. In turn, this reduces the cost of TLB reconfiguration processing by the OS and improves memory access performance.



Figure K.2 Conceptual Illustration of Data Access on Normal and Large Pages



A TLB miss occurs at the first access to a memory address by an application program, because the TLB does not have translation information for the address. Looking at "Figure K.2 Conceptual Illustration of Data Access on Normal and Large Pages" as an example, the data access to four normal pages in a certain memory space caused four TLB misses, whereas the same data access to a large page in the same memory space caused only one TLB miss.

### Benefits and Drawbacks by Differences in Page Size

The following table shows the respective benefits and drawbacks of using page sizes of 64 KiB and 2 MiB.

Table K.1 Benefits and Drawbacks According to Page Size

Evaluation Item	64 KiB	2 MiB
TLB miss rate	High	Low
Memory initialization cost	Low	High
Memory use efficiency	High	Low

1. Normal pages have a page size of 64 KiB. Although their TLB miss rate is high, the memory initialization cost is low and memory use efficiency is high, so they may be effective when memory usage by application programs is small.
2. Large pages have a page size of 2 MiB. Their memory initialization cost is higher and memory use efficiency is lower as compared to 64-KiB pages. However, they have a lower TLB miss rate, so the large page function of the FX system uses large pages (2 MiB) by default.

Application programs use memory of various sizes ranging from small-scale memory for, as an example, a communication buffer used in typical communications to large-scale memory for computing.

Memory use efficiency may be low for large pages used when memory usage by application programs is small. For example, if large pages are used even when the memory usage of the heap area is only 1 MiB, 2 MiB of memory is reserved, resulting in an unused memory area of 1 MiB.

On the other hand, the TLB miss rate is high for normal pages used when memory usage by application programs is significant. Furthermore, there will be a high possibility that execution performance does not improve.

## K.1.2 Paging Methods

The two paging methods are demand paging and prepaging. For the large page function of the FX system, you can set a paging method for each of the following memory areas: static data area (.bss area), stack/thread stack areas, and reserved dynamic memory areas. This uses an environment variable (XOS\_MMM\_L\_PAGING\_POLICY) described in the section titled "K.3 Environment Variables for Large Page Library Settings".

1. The demand paging method is a method of allocating a page to main memory, as needed, if the necessary page is not in main memory during execution of an application program. The physical page is allocated at the timing of initial access to the memory area.
2. The prepaging method is a method of allocating a page to main memory beforehand. The physical page is allocated at the timing of memory area allocation.

Table K.2 Memory Access Costs According to Paging Method Differences

Paging Method	Memory Access Inside/Outside NUMA	Initial Memory Access
Demand paging method	Uses internal NUMA memory as much as possible (Cost: Low)	Loads page to physical memory during memory access (Cost: High)
Prepaging method	Uses memory regardless of whether it is inside or outside NUMA (Cost: High)	Loads page to physical memory beforehand (Cost: Low)

As shown in "Table K.2 Memory Access Costs According to Paging Method Differences", two factors affect memory access performance. The first factor is whether the physical memory accessed from a certain compute core is in the same NUMA. Access to physical memory not belonging to the same NUMA as the compute core requires a greater processing cost than access to physical memory in the same NUMA. The second factor is the cost when memory is accessed for the first time. If the demand paging method is selected, physical memory is allocated at the memory access time, so a page load processing cost is incurred when the memory is accessed for the first time.

For example, suppose that a four-core thread-parallel application program dynamically acquires memory areas for individual threads. In this case, switching from the prepaging method to the demand paging method increases the frequency of access to physical memory in the same NUMA as the compute core. You can expect an improvement in memory access performance as a result.

## K.2 Large Pages of the FX System

The large page library (libmpg.so) is provided as an extension function of the FX system. This section describes the page sizes of individual areas and the location of application program data when using the large page library of the FX system.

### K.2.1 FX System Memory Areas and Large Page Applicability

From the areas listed below, the large page library of the FX system puts the following areas on large pages (2-MiB pages): static data areas (.bss and .data areas), reserved dynamic memory areas (mmap areas), thread heap areas, stack areas, and thread stack areas. Text areas, reserved dynamic memory areas (heap areas), and shared memory areas are not intended to be put on large pages.

Table K.3 FX System Memory Areas and Large Page Applicability

Area	Large Page Applicability	Usage of Area
Text (.text) area	No (64-KiB page)	Memory area for placing a series of instructions from an application program (a.out)
Static data (.data) area	Yes (2-MiB page)	Memory area for storing the static data of an application program (a.out) (Initialized)
Static data (.bss) area	Yes (2-MiB page)	Memory area for storing the static data of an application program (a.out) (Not initialized)
Reserved dynamic memory area (heap area)	No (64-KiB page)	Process heap area/Main thread heap area This memory area is not put on large pages in the large page library. Memory areas are allocated when: Acquired with a brk(2)/sbrk(2) system call, or allocated when a dynamic memory reservation request (malloc(3)) specifying a smaller size than the size specified in the environment variable MALLOC_MMAP_THRESHOLD_ is issued without linking the large page library.

Area	Large Page Applicability	Usage of Area
Thread heap area	Yes (2-MiB page)	Subthread heap area
Stack area	Yes (2-MiB page)	Process stack area/Main thread stack area (Use of large pages must be explicitly specified in the environment variable XOS_MMM_L_LPG_MODE.)
Thread stack area	Yes (2-MiB page)	Subthread stack area (Use of large pages must be explicitly specified in the environment variable XOS_MMM_L_LPG_MODE.)
Reserved dynamic memory area (mmap area)	Yes (2-MiB page)	Memory area allocated when mmap(2) is issued Memory areas are allocated when: Allocated when the large page library is linked and a dynamic memory reservation request (malloc(3)) specifying a larger size than the size specified in the environment variable MALLOC_MMAP_THRESHOLD_ is issued without linking the large page library, or alternatively allocated as a thread heap area/ thread stack area too.
Shared memory	No (64-KiB page)	Used for memory sharing between processes

The large page library, libmpg.so, is a large page function optimized for the FX system. By using the overcommit function of HugeTLBs, it improves the usability of the user I/F without reserving a memory pool beforehand. By specifying the `-Klargepage` option at compilation, the application program will use this library. To use the standard large page function in Linux without using libmpg.so, specify the `-Knolargepage` option at compilation. To use normal pages, explicitly specify the environment variable `XOS_MMM_L_HPAGE_TYPE=none`.

For details on the environment variables described in this section, see the section titled "[K.3 Environment Variables for Large Page Library Settings](#)".

### Note

#### - Link order of the large page library

The large page library specification (`-lmpg` option specified by the compiler) must always precede the library explicitly specified by the end user (e.g., `-lc`, `-lpthread`, etc.). Otherwise, large pages may not be allocated successfully.

#### - Memory area acquired by `brk(2)`/`sbrk(2)`

Due to kernel specifications, memory acquired by `brk(2)`/`sbrk(2)` is not on large page. So linking the large page library libmpg.so together with `brk(2)`/`sbrk(2)` and `malloc(3)`(Size request greater than environment variable `MALLOC_MMAP_THRESHOLD_`) would result in a mixture of normal and large pages, and depending on the amount and timing of memory acquisition, `brk(2)`/`sbrk(2)` could end with `ENOMEM` without getting any memory.

It is recommended that you use `malloc(3)` to link the large page library libmpg.so.

If you want to use `brk(2)`/`sbrk(2)` in application programs that link the large page library libmpg.so, you must explicitly specify the environment variable `XOS_MMM_L_HPAGE_TYPE=none`.

#### - Applications that signal frequently

If you link a large page library(libmpg) with a prepaging method and use the `timer_create(2)` system call to generate a `SIGALRM`/`SIGVTALRM` or any other signal frequently, the application-issued `fork(2)`(`clone(2)`) receives the signal and the system call continues to return with `ERESTARTNOINTR`. As a result, frequent retries of `fork(2)`(`clone(2)`) may occur and it does not terminate.

In such cases, setting the paging method to demand paging avoids the problem.

Run the application with the environment variable `XOS_MMM_L_PAGING_POLICY=demand:demand:demand` explicitly.

## Information

You can create large-page-enabled application programs, even with a general compiler (e.g., gcc), by linking this large page library, libmpg.so. If you explicitly link to other libraries (example below, "-lc -lpthread"), specify that the large page libraries always precede those libraries.

To enable large pages, the large page library also provides a large page linker script in the path shown below. This script is intended to put static data (.data and .bss) on large pages. Specify the linker script appropriately for the compiler too.

```
/opt/FJSVxos/mmm/util/bss-2mb.lds
```

The following example compiles an application program on gcc.

```
Example)
gcc -Wl,-T/opt/FJSVxos/mmm/util/bss-2mb.lds -L/opt/FJSVxos/mmm/lib64 -lmpg -lc -lpthread
test_program.c
```

If an application program is compiled in PIE (Position Independent Executable) format, the .data/.bss areas are not on large pages. In this case, the large page library only outputs a warning log, normal pages are used for the .data/.bss areas, and the application program continues running. As an example, to compile an application program on gcc while explicitly avoiding the PIE format, use the -no-pie option.

## K.2.2 Memory Areas for Placing Application Program Data

Variables in application programs are placed in different memory areas according to how they are defined. Therefore, when specifying the page size of each memory area, you need to be aware of which memory areas that the variables are placed in.

This section describes which memory areas have variables placed in them along with how they are declared in examples of Fortran, C, and C++ source code. The page size for each allocated area is as shown in "Table K.3 FX System Memory Areas and Large Page Applicability" in "K.2.1 FX System Memory Areas and Large Page Applicability".

- Fortran

```
program main
integer*8, parameter::N=(1024_8)
real*8 a(N)                ! a: Local array without initial values
real*8 :: b(N)=1.0         ! b: Array with initial values
real*8,allocatable::c(:)   ! c: Allocatable array
allocate(c(N))
...
end
```

Local array a is placed in a static data area (.bss). However, for an application program compiled with the -Kauto or -Kthreadsafe enabling option, it is placed in the process stack area. Array b is placed in a static data area (.data). Array c is placed in a reserved dynamic memory area (heap or mmap area).

- C

```
#include <stdlib.h>
#include <pthread.h>
#define N 1024
double a[N];                // a: Global variable without initial value
double b[N]={1.0};         // b: Global variable with initial value
double *c;
double *d;                  // c, d: Pointer variable
void *func_thread(void *args) {
    double f[N];           // f: Local variable
    d=(double *)malloc(sizeof(double)*N);
    ...
}
```

```

int main(void){
    double e[N]; // e: Local variable
    c=(double *)malloc(sizeof(double)*N);
    ...

    pthread_t pthread;
    pthread_create( &pthread, NULL, &func_thread, (void *)NULL);
    ...
}

```

Global variable a is placed in a static data area (.bss). Global variable b is placed in a static data area (.data). Pointer variables c and d are placed in a reserved dynamic memory area (heap or mmap area). Local variable e is placed in the process stack area. Local variable f is placed in the thread stack area.

### Note

When the thread-parallel application program in each programming language is executed, a dedicated thread stack area is prepared as the stack area for each thread.

- C++

```

#include <vector>
const int N = 1024;
struct Klass {
    double k;
    Klass() : k (0.0) {}
    Klass(double K) : k (K) {}
};
std::vector<Klass> a(N); // Secure area by using vector class
int main(){
    Klass* b = new Klass[N]; // Secure area by using new operator
    std::vector<double> c(N); // Secure area by using vector class
    return 0;
}

```

Variables a, b, and c are all placed in a reserved dynamic memory area (heap or mmap area).

### Note

The memory areas allocated for the variables in C++ are the same as those in C, but the memory areas dynamically acquired by vector class, etc. are allocated as reserved dynamic memory areas (heap area or mmap area).

## K.3 Environment Variables for Large Page Library Settings

You can use the environment variables shown in this section to adjust the behavior of the large page library.

### K.3.1 Basic Settings of the Large Page Library

This section shows the environment variables for the basic settings of the large page library.

Table K.4 Environment Variables for the Basic Settings of the Large Page Library

Variable Name	Setting Value	Default Value	Details
XOS_MMM_L_HPAGE_TYPE	hugetlbfs none	hugetlbfs	The setting is the selection to enable/disable operations with large page allocation by the large page library.

Variable Name	Setting Value	Default Value	Details
			<p>"hugetlbfs" enables large pages with HugeTLBfs.</p> <p>"none" disables large pages with the large page library. If it is specified, all specified environment variables beginning with "XOS_MMM_L_" for the large page library are invalid.</p> <p>If any other value is specified, "hugetlbfs" is assumed specified.</p> <p>There is a precaution on enabling large pages. See the note below this table ("Stack area display of /proc/pid/maps when large pages are enabled").</p>
XOS_MMM_L_LPG_MODE	base+stack base	base+stack	<p>The setting is the selection to enable/disable operation with large page allocation for the stack area and thread stack area.</p> <p>"base+stack" enables large pages for not only the static data and reserved dynamic memory areas but also the stack area and thread stack area.</p> <p>"base" enables large pages for only the static data and reserved dynamic memory areas. Large pages are not enabled for the stack area and thread stack area.</p> <p>If any other value is specified, "base+stack" is assumed specified.</p> <p>There is a precaution on enabling large pages for the stack area. See the note below this table ("Alignment with the stack area on large pages").</p>
XOS_MMM_L_PRINT_ENV	on off 1 0	0	<p>The setting is used for debugging application programs.</p> <p>If "1" or "on" is specified in this environment variable, a list of the performance tuning environment variables provided by the large page library is output to the standard error output. The list is output only once, at the application program start time (before the main function is executed). If "0" or "off" is specified, the list of performance tuning environment variables is not output to the standard error output.</p> <p>If any other value is specified, "0" is assumed specified.</p> <p>The output after the execution of the main function does not reflect the settings changed by mallopt(3).</p>

## Note

- Stack area display of `/proc/pid/maps` when large pages are enabled (*pid*: Process ID)

The displayed contents of `/proc/pid/maps` are partly rewritten when large pages are enabled. This is shown in the following examples. (The same applies to `/proc/pid/numa_maps`.)

- Static data (.bss/.data area) or reserved dynamic memory area (heap/mmap area)

(Before enabling large pages)

```
00430000-00890000 rw-p 00000000 00:00 0 [heap]
```

(After enabling large pages)

```
aaaae2400000-aaaae2600000 rw-p 00000000 00:0e 452989 /anon_hugepage (deleted)
```

- Stack area (process stack area/main thread stack area)

(Before enabling large pages)

```
fffffffd0000-1000000000000000 rw-p 00000000 00:00 0 [stack]
```

(After enabling large pages)

```
fffaeb800000-1000000000000000 rw-p 00000000 00:0e 274404 /memfd: [stack] by libmpg (deleted)
```

- Alignment with the stack area on large pages

Basically, the size of a stack area on a large page is aligned to the size of an allocated HugeTLBfs page, based on the soft limit value that is set in `ulimit -s`.

The start and end addresses of an allocated HugeTLBfs page are also similarly aligned up or down.

As a result of aligning mapped addresses and sizes, the stack area may overlap an adjacent area. If there is an overlap, large pages are disabled. Then, a warning message is output, and application programs continue as is with normal pages.

## K.3.2 Paging Method Setting

This section shows the environment variable for a paging method setting of the large page library.

Table K.5 Environment Variable for a Paging Method Setting

Variable Name	Setting Value	Default Value	Details
XOS_MMM_L_PAGING_POLICY	[demand prepage]: [demand prepage]: [demand prepage]	prepage:demand:prepage	<p>The setting is the selection of a paging method (page allocation time) for each memory area.</p> <p>"demand" means the demand paging method, and "prepage" means the prepaging method.</p> <p>This variable specifies paging methods for three memory areas by delimiting the methods with a colon (:).</p> <p>The first specified method is for the .bss area for static data. ("prepage" is always set for the .data area for static data, which is not subject to specification of a paging method.)</p> <p>The second specified method is for the stack area and thread stack area.</p> <p>The third specified method is for the reserved dynamic memory areas.</p>

Variable Name	Setting Value	Default Value	Details
			<p>If any other value is specified, "prepage:demand:prepage" is assumed specified.</p> <p>The prepagging method prevents subsequent occurrences of page faults in an application program and tends to improve performance and also reduce performance variations. However, in some cases, it is better not to select the prepagging method for any application program that implements prepagging with its own mechanism (for example, zero-clearing or the like) or accesses only a partial memory area.</p>

### K.3.3 Settings for Tuning

This section shows the environment variables for tuning related to large page allocation.

Of these environment variables, those with names beginning with "XOS\_MMM\_L\_" have uniquely been added to the large page library. The other environment variables are glibc variables.

Table K.6 Environment Variables for Tuning

Variable Name	Setting Value	Default Value	Details
XOS_MMM_L_ARENA_FREE	1 2	1	<p>The setting relates to handling of the heap area released by free(3).</p> <p>If "1" is specified, the memory that can be released is immediately released. If "2" is specified, no memory is released, and all the memory is reused in pools. If a value other than "1" or "2" is specified, "1" is assumed specified.</p> <p>If "2" is specified, memory is allocated from the heap area even for a memory request for a size equal to or greater than the value specified in MALLOC_MMAP_THRESHOLD_. Even after being released, the memory is retained as a free memory area. Release processing of the heap area is not executed. Since memory of all sizes must be allocated in a single heap area to prevent decreases in memory use efficiency, the thread heap area is not generated. In other words, "2" is equivalent to the combination of the following settings:</p> <p>XOS_MMM_L_ARENA_LOCK_TYPE=1  XOS_MMM_L_MAX_ARENA_NUM=1  MALLOC_MMAP_THRESHOLD_=ULONG_MAX  MALLOC_TRIM_THRESHOLD_=ULONG_MAX</p>
XOS_MMM_L_ARENA_LOCK_TYPE	0 1	1	<p>The setting relates to the memory allocation policy. "0" means that memory acquisition performance takes priority, whereas "1" means memory use efficiency takes priority. If a value other than "0" or "1" is specified, "1" is assumed specified.</p>



Variable Name	Setting Value	Default Value	Details
			<p>If "0", a new thread heap area is generated when the main arena contention occurs. If "1", when there is a main arena contention, it will generate a thread heap area for each thread if it is within the maximum arena count (XOS_MMM_L_MAX_ARENA_NUM), otherwise it will wait for a lock to be acquired.</p> <p>Suppose that, in a multithread application program, threads simultaneously call a memory acquisition request. In this case, "0" will enable parallel processing but decrease memory use efficiency. "1" will increase memory use efficiency even though memory acquisition will be processed sequentially.</p> <p>In the case of frequent contention among memory acquisition requests from threads in a multithread application program, the "0" setting may improve performance.</p>
XOS_MMM_L_MAX_ARENA_NUM	Integer between 1 and INT_MAX [Decimal number]	1	<p>You can set the number of arenas that can be generated (total number of process heap and thread heap areas) with this variable, which is valid only when XOS_MMM_L_ARENA_LOCK_TYPE=1. Use it when you want to limit the number of thread heap areas generated.</p> <p>With the default setting ("1"), only process heap areas are used, and no thread heap areas are generated. This case is equivalent to XOS_MMM_L_ARENA_LOCK_TYPE=1. If the setting is n ( 2), in addition to process heap areas, (n-1) thread heap areas may be generated.</p>
XOS_MMM_L_HEAP_SIZE_MB	Integer between MALLOC_MMAP_THRESHOLD_ x 2 and ULONG_MAX <in MiB> [Decimal number]	MALLOC_MMAP_THRESHOLD_ x 2	<p>The variable, for using a thread heap area, sets the size of memory acquired when generating or extending the thread heap area.</p> <p>The default value is MALLOC_MMAP_THRESHOLD_ multiplied by 2. A thread heap area with the default value is generated when memory for the thread is acquired for the first time. Furthermore, a thread heap area with the specified value is generated (expanded) when the total amount of acquired memory per thread exceeds the default value.</p> <p>If the expected amount of memory acquired for each thread is small, setting a small value for this environment variable may increase memory use efficiency.</p>
XOS_MMM_L_COLORING	0 1	1	<p>The setting enables/disables cache coloring.</p> <p>Cache coloring reduces conflicts in the L1 cache of the processor.</p> <p>If "0" is specified, cache coloring is not used.</p>

Variable Name	Setting Value	Default Value	Details
			<p>If "1" is specified, cache coloring is used when memory is acquired by <code>mmap(2)</code> using the <code>MALLOC_MMAP_THRESHOLD_</code> size (default value: 128 MiB) or greater. Cache coloring is not used when memory is acquired from the heap area using a smaller size than <code>MALLOC_MMAP_THRESHOLD_</code>.</p> <p>The conditions under which cache coloring must be enabled are as follows:</p> <ol style="list-style-type: none"> <li>1) <code>MALLOC_MMAP_THRESHOLD_</code> or larger size of <code>malloc(3)</code> request, and</li> <li>2) <code>XOS_MMM_L_FORCE_MMAP_THRESHOLD=1</code> (This forces <code>mmap(2)</code> to allocate memory)</li> </ol> <p>If a value other than "0" or "1" is specified, "1" is assumed specified.</p> <p>Generally, cache coloring tends to improve performance. However, disabling cache coloring with this library can be said to be better for any application program that implements cache coloring with its own mechanism.</p>
<code>XOS_MMM_L_FORCE_MMAP_THRESHOLD</code>	0 1	0	<p>The setting specifies whether <code>mmap(2)</code> takes priority when acquiring memory of a size equal to or greater than <code>MALLOC_MMAP_THRESHOLD_</code> (default value: 128 MiB).</p> <p>If "0" is specified, <code>mmap(2)</code> does not take priority. First, there is a search for free space in the heap area. Free memory in the heap area is returned when free space is found. <code>mmap(2)</code> is used to acquire memory only when no free space is found in the heap area.</p> <p>If "1" is specified, <code>mmap(2)</code> takes priority. <code>mmap(2)</code> is used to acquire memory (even if there is free space) without searching for free space in the heap area.</p> <p>If a value other than "0" or "1" is specified, "0" is assumed specified.</p>
<code>MALLOC_CHECK_</code>	0 1 2 3 5 7 [Decimal number]	3	<p>The setting relates to detection of programming errors (memory destruction, double free, etc.). According to the setting, the following actions are taken when errors are detected.</p> <p>0: Ignore the error, and continue processing.</p> <p>1: Output a detailed error message, and continue processing.</p> <p>2: Abort the application program.</p> <p>3: Output a detailed error message, stack trace data, and memory mapping data, and abort the application program.</p> <p>5: Output a simple error message, and continue processing.</p>

Variable Name	Setting Value	Default Value	Details
			<p>7: Output a simple error message, stack trace data, and memory mapping data, and abort the application program.</p> <p>If any other value is specified, action is taken according to the lower 3 bits of the value as a binary number.</p> <p>Not all errors can be detected. Errors such as a memory leak cannot be detected.</p>
MALLOC_TOP_PAD_	<p>Integer between 0 and ULONG_MAX &lt;in bytes&gt;</p> <p>[Decimal number]</p>	<p>131072 (= 128 KiB)</p>	<p>The setting specifies the extended size for each time that the heap area is extended. A value rounded up to the nearest page size is used.</p> <p>If this setting value exceeds the amount of memory acquired or released at one time by an application program, the number of issued system calls decreases, and memory acquisition performance may improve. (However, memory use efficiency may decrease.)</p>
MALLOC_PERTURB_	<p>Integer between INT_MIN and INT_MAX</p> <p>[Decimal number]</p>	0	<p>The setting is used for debugging application programs. Based on the value specified in this environment variable, a memory area is filled when memory is acquired (except by calloc(3)) or released. When memory is acquired, the complement of the lowest byte of the specified value is written. When memory is released, the lowest byte of the specified value is written.</p> <p>Use this variable to detect problems like the use of an uninitialized memory area or a reference to a released memory area.</p>
MALLOC_MMAP_MAX_	<p>Integer between INT_MIN and INT_MAX</p> <p>[Decimal number]</p>	<p>2097152 (= 2*1024*1024)</p>	<p>The setting is the value of the upper limit on the number of times that memory is acquired by mmap(2) during acquisition of memory of a size equal to or greater than MALLOC_MMAP_THRESHOLD_ (default size: 128 MiB).</p> <p>The current count value is incremented by 1 when mmap(2) acquires memory. The current count value is decremented by 1 when free(3) releases memory.</p> <p>An upper limit value is set for each process.</p> <p>Set "0" to not acquire memory with mmap(2) but instead acquire memory from the heap area regardless of the memory acquisition size and set value of MALLOC_MMAP_THRESHOLD_.</p>
MALLOC_MMAP_THRESHOLD_	<p>Integer between 0 and ULONG_MAX &lt;in bytes&gt;</p> <p>[Decimal or hexadecimal number]</p>	<p>134217728 (= 128 MiB)</p>	<p>For a memory acquisition request with a memory size equal to or greater than the size specified in this environment variable, mmap(2) is used to acquire memory. For a memory acquisition request with a memory size smaller than the size specified in the environment variable, memory is acquired from the heap area.</p>

Variable Name	Setting Value	Default Value	Details
			<p>When executed for memory acquired with <code>mmap(2)</code>, <code>free(3)</code> immediately releases the memory. On the other hand, when executed for memory acquired from the heap area, <code>free(3)</code> does not immediately release the memory but instead pools it, unless it is in a contiguous free area of a size equal to or greater than <code>MALLOC_TRIM_THRESHOLD_</code> at the top of the heap area. The pooled memory is reusable.</p> <p>If this setting value is large, the maximum memory size managed in the heap area increases, reuse of memory areas may be promoted, and memory acquisition performance may improve. (However, memory use efficiency may decrease.)</p>
<code>MALLOC_TRIM_THRESHOLD_</code>	Integer between 0 and <code>ULONG_MAX</code> <in bytes>  [Decimal or hexadecimal number]	134217728 (= 128 MiB)	<p>The setting is the threshold for immediately releasing memory when <code>free(3)</code> is executed for memory acquired from the heap area. If it is contiguous free memory of a size equal to or greater than the specified threshold and positioned at the top of the heap area, the memory is immediately released. If it is contiguous free memory smaller than the threshold, the memory is not released but pooled instead.</p> <p>If this setting value is large, memory areas acquired at one time are released less frequently, reuse of memory areas may be promoted, and memory acquisition performance may improve. (However, memory use efficiency may decrease.)</p>



## Example

### Example of checking the effect of tuning with the environment variable `XOS_MMM_L_ARENA_FREE`

#### - Application program description

1. Acquire (`malloc(3)`) and release (`free(3)`) memory 1,024 times from a heap area with a size of 8 MiB.
2. Execute the processing of 1 in two loops.
3. Measure the time taken by `malloc(3)/free(3)` each time.

#### - Application program

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <errno.h>
#include <time.h>

#define N 1024
#define MALLOC_CNT 1024
#define DATA_CNT (1024*1024*1024)

clock_t time_start;
clock_t time_end;
```

```

double *c[MALLOC_CNT]; //heap memory
double a[DATA_CNT]; //data memory

int main(int argc, char *argv[]){
    int i;
    long sec;
    long nsec;
    int loop=0;
    struct timespec time1 ={0,0};
    struct timespec time2 ={0,0};

    while(loop <2){
        printf("malloc start.\n");
        clock_gettime(CLOCK_REALTIME, &time1);
        for(i=0;i<MALLOC_CNT;i++){
            c[i]=(double *)malloc(sizeof(double)*N*N);
            if (c[i] == NULL) {
                fprintf(stderr, "malloc error: cnt=%d, errno=%d\n", i, errno);
                exit(1);
            }
        }
        clock_gettime(CLOCK_REALTIME, &time2);
        printf("malloc end.\n");
        sec = (time2.tv_sec - time1.tv_sec);
        nsec= (time2.tv_nsec-time1.tv_nsec);
        if(nsec<0){
            sec--;
            nsec += 1000000000L;
        }
        printf("MALLOC TIME:%d:%010d\n", sec, nsec);
        sleep(10);

        printf("free start.\n");
        clock_gettime(CLOCK_REALTIME, &time1);
        for(i=0;i<MALLOC_CNT;i++){
            free(c[i]);
        }
        clock_gettime(CLOCK_REALTIME, &time2);
        printf("free end.\n");
        sec = (time2.tv_sec - time1.tv_sec);
        nsec= (time2.tv_nsec-time1.tv_nsec);
        if(nsec<0){
            sec--;
            nsec += 1000000000L;
        }
        printf("FREE TIME:%d:%010d\n", sec, nsec);
        loop++;
    }
    return EXIT_SUCCESS;
}

```

## - Tuning method

Using the following two patterns, set the environment variable XOS\_MMM\_L\_ARENA\_FREE, and compare performance:

1. Setting for immediately releasing the memory pages that can be released at the free(3) time

```
export XOS_MMM_L_ARENA_FREE=1
```

2. Setting for not immediately releasing the memory pages that can be released at the free(3) time

```
export XOS_MMM_L_ARENA_FREE=2
```

### - Explanation and performance prediction

If `XOS_MMM_L_ARENA_FREE=1` is set, the reserved memory area in the mmapped chunk is released immediately after `free(3)` is executed for the first time. Later, when `malloc(3)` is executed for the second time, its processing is considered to take nearly the same length of time as the processing of the first `malloc(3)` execution since a memory area has to be secured in the mmapped chunk again.

If `XOS_MMM_L_ARENA_FREE=2` is set, the reserved memory area in the heap area is not released after `free(3)` is executed for the first time. The processing of the second `malloc(3)` execution is considered to take a shorter length of time since the reserved memory area in the heap area is reused, reducing the memory allocation cost.

## Example

### Example of checking the effect of tuning with the environment variable `XOS_MMM_L_ARENA_LOCK_TYPE`

#### - Application program description

1. Create multiple threads, and allocate each thread to a different CPU.
2. Issue `malloc(3)` individually for each thread, and secure the heap area.
3. Execute the processing of 1 and 2 in 10 loops, and measure the time taken by `malloc(3)/free(3)` each time.

#### - Application program

```
#define _GNU_SOURCE
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <pthread.h>
#include <sched.h>
#include <errno.h>
#include <sys/time.h>

#define KBYTES (1024)
#define CPU_NUM (16)
#define MAX_MALLOC_CNT (5*64*KBYTES)

pthread_t thread[CPU_NUM];
pthread_barrier_t barrier;
int malloc_size = 64*KBYTES;
int malloc_cnt = MAX_MALLOC_CNT/CPU_NUM;
int loop_cnt = 10;
char *strp[MAX_MALLOC_CNT];
int cpuid[CPU_NUM]={};
int cpunum[CPU_NUM]={};

void* thread_main(void *arg) {
    int cpuid = *(int*)(arg);
    int i;
    pthread_barrier_wait(&barrier);
    for (i=0;i<malloc_cnt;i++) {
        strp[i+cpuid*malloc_cnt] = (char*)malloc(malloc_size);
    }
    return NULL;
}

int main(int argc, char *argv[]) {
    int i, ret, loop;
    cpu_set_t cpu;
    struct timeval st, et, rt;

    if ( 2 == argc ) {
```

```

    malloc_size = atoi(argv[1]);
} else if ( 3 == argc ) {
    malloc_size = atoi(argv[1]);
    malloc_cnt = atoi(argv[2]);
} else if ( 4 == argc ) {
    malloc_size = atoi(argv[1]);
    malloc_cnt = atoi(argv[2]);
    loop_cnt = atoi(argv[3]);
}

loop = 0;
while ( loop < loop_cnt ) {
    ret = pthread_barrier_init(&barrier, NULL, CPU_NUM+1);
    for (i=0;i<CPU_NUM;i++) {
        cpuid[i] = i;
        cpunum[i] = i;
        ret = pthread_create(&thread[i], NULL, thread_main, (void*)&cpunum[i]);
        if ( 0 != ret ) {
            fprintf(stderr, "pthread_create: errno=%d\n", ret);
            exit(1);
        }

        CPU_ZERO(&cpu);
        CPU_SET(cpuid[i], &cpu);
        ret = pthread_setaffinity_np(thread[i], sizeof(cpu_set_t), &cpu);
        if ( 0 != ret ) {
            fprintf(stderr, "pthread_setaffinity_np: errno=%d\n", ret);
            exit(1);
        }
    }

    pthread_barrier_wait(&barrier);
    gettimeofday(&st, NULL);
    for (i=0;i<CPU_NUM;i++) {
        ret = pthread_join(thread[i], NULL);
        if ( 0 != ret ) {
            fprintf(stderr, "pthread_join: errno=%d\n", ret);
            exit(1);
        }
    }

    gettimeofday(&et, NULL);
    timersub(&et, &st, &rt);
    printf("%d %d %ld.%ld\n", malloc_size, malloc_cnt, rt.tv_sec, rt.tv_usec);
    fflush(NULL);
    pthread_barrier_destroy(&barrier);
    for (i=0;i<malloc_cnt*CPU_NUM;i++) {
        free(strp[i]);
    }
    loop++;
}
return 0;
}

```

#### - Tuning method

Using the following two patterns, set the environment variable XOS\_MMM\_L\_ARENA\_LOCK\_TYPE, and compare performance:

1. Prioritizing memory acquisition performance (processing dynamic memory reservation requests in parallel)

```
export XOS_MMM_L_ARENA_LOCK_TYPE=0
```

## 2. Prioritizing memory use efficiency (processing dynamic memory reservation requests sequentially)

```
export XOS_MMM_L_ARENA_LOCK_TYPE=1
```

The environment variable XOS\_MMM\_L\_ARENA\_FREE=2 is also set with each pattern. This facilitates processing cost comparison in cases of contention among dynamic memory reservation requests by multiple threads. Also, the resulting processing cost for physical page allocation is low.

The above example program gets 20480 times of memory of size 64 KiB per 1 thread. It competes in 16 threads and measures 10 times.

### - Explanation and performance prediction

The possibility of contention in memory acquisition processing with malloc(3) is high when multiple threads issue malloc(3) simultaneously.

If XOS\_MMM\_L\_ARENA\_LOCK\_TYPE=0 is set, a thread heap area is generated when malloc(3) contention occurs, thereby enabling dynamic memory reservation requests to be processed in parallel. This is considered to shorten the time taken for memory acquisition.

If XOS\_MMM\_L\_ARENA\_LOCK\_TYPE=1 is set, a thread heap area is not generated when malloc(3) contention occurs. Also, threads share the process heap area. For this reason, their processing must be done sequentially while waiting for the malloc(3) processing currently in progress to complete. This is considered to lengthen the time taken for memory acquisition.

.....

## K.4 Messages

---

This section describes the messages output by the large page library. The messages appear at the standard error output.

### K.4.1 How to Read Messages

---

The messages output by the large page library have the following format.

```
message-type component-name - message-text
```

The following table shows the components of a message from the large page library.

Table K.7 Message Components

Component	Meaning
Message type	Displays the following message type: [WARN] Warning message
Component name	Displays the xos LPG number. The number is a unique message identification number.
' '	Delimiter character
Message text	Displays the details of the event that occurred.

### K.4.2 Messages

---

#### Warning Message

**[WARN] xos LPG 2001 - Failed to allocate HugeTLBfs pages, going to try allocating normal pages.**

#### Meaning

Normal pages were acquired because acquisition of large pages (HugeTLBfs) failed.

#### Action

The possible factors in the failure to acquire large pages include insufficient free memory and fragmented memory. Check the amount of memory used by the job and the amount of free memory. As needed, ask the system administrator about the amount of memory available to a job. If this phenomenon recurs despite an appropriate amount of memory used by the application program and appropriate system settings, collect a core dump and contact a Fujitsu systems engineer (SE) or Fujitsu Support Desk.



---

**[WARN] xos LPG 2002 - Failed to map HugeTLBfs for data/bss: *a.out***  
**The *e\_type* of elf header must be ET\_EXEC when using libmpg. You can check it on your load module by *readelf -h* command.**

#### Meaning

The *.data/.bss* area cannot be on a large page (HugeTLBfs) because the binary format of the application program is a PIE (Position Independent Executable). Using a normal page for the *.data/.bss* area, execution continued.

*a.out*: Application program

#### Action

To put the *.data/.bss* area on a large page, compile the application program while preventing the binary format of the program from becoming PIE. (Using the *readelf -h* command, create the application program in such a way that *e\_type* becomes ET\_EXEC).

---

**[WARN] xos LPG 2003 - Failed to map HugeTLBfs for data/bss: Layout problem with segments *seg\_index1* and *seg\_index2*: Segments would overlap.**

#### Meaning

The *.data/.bss* area for the application program cannot be on a large page (HugeTLBfs) because the address range of the *.data/.bss* area conflicts with another segment. Using a normal page for the *.data/.bss* area, execution continued.

*seg\_index1, seg\_index2*: Segment index value

#### Action

To put the *.data/.bss* area on a large page, specify an appropriate linker script (linker option) for the application program compile time to align the virtual address of the *.data/.bss* area to the large page size.

---

**[WARN] xos LPG 2004 - Failed to map HugeTLBfs for thread stack: specified stack address in *pthread\_attr: addr=stack\_addr size=stack\_size***

#### Meaning

The thread stack cannot be on a large page (HugeTLBfs) because the virtual address of the thread stack is specified in the second argument *attr* (attribute) of *pthread\_create(3)*. Using a normal page for the thread stack, execution continued.

*stack\_addr*: Specified stack address

*stack\_size*: Specified stack size

#### Action

To put the thread stack on a large page, do not specify the virtual address of the thread stack in the second argument *attr* (attribute) of *pthread\_create(3)*.

---

**[WARN] xos LPG 2005 - Failed to map HugeTLBfs for process stack: confirmed the existence of multiple threads.**

#### Meaning

The process stack cannot be on a large page (HugeTLBfs) because multiple threads were detected during the processing to put the process stack on a large page. Using a normal page for the process stack, execution continued.

#### Action

There must be only a single thread during the processing to put the process stack on a large page. Check whether your program generates a thread. Also consider that other libraries, not just the large page library, can also generate threads.

---

**[WARN] xos LPG 2006 - Failed to map HugeTLBfs for process stack: The bottom of the process stack is invading the next vma: *bottom=stack\_bottom\_addr next-vma=next\_vma\_addr***

## Meaning

The process stack cannot be on a large page (HugeTLBfs) because the process stack conflicts with another memory area, which is located after (at a higher address than) the stack. Using a normal page for the process stack, execution continued.

*stack\_bottom\_addr*: Bottom address of the stack

*next\_vma\_addr*: Start address of the memory area after the stack

## Action

To put the process stack on a large page, the address range must be guaranteed not to overlap with another memory area. Check whether your program explicitly uses `mmap(2)` to map a memory area near the process stack. Also consider that dynamic libraries other than the large page library can also map memory areas.

---

**[WARN] xos LPG 2007 - Failed to map HugeTLBfs for process stack: The top of the process stack is invading the previous vma: top=stack\_top\_addr prev-vma=prev\_vma\_addr**

## Meaning

The process stack cannot be on a large page (HugeTLBfs) because the process stack conflicts with another memory area, which is located before (at a lower address than) the stack. Using a normal page for the process stack, execution continued.

*stack\_top\_addr*: Start address of the stack

*prev\_vma\_addr*: Bottom address of the memory area before the stack

## Action

To put the process stack on a large page, the address range must be guaranteed not to overlap with another memory area. Check whether your program explicitly uses `mmap(2)` to map a memory area near the process stack. Also consider that dynamic libraries other than the large page library can also map memory areas.