

Intel® C++ Compiler Professional Edition 11.1 for Mac OS* X Installation Guide and Release Notes

Document number: 321413-002US

9 February 2010

Table of Contents

1	Introduction	3
1.1	Change History	3
1.2	Product Contents	4
1.3	System Requirements.....	4
1.4	Documentation.....	4
1.5	Technical Support.....	4
2	Installation.....	5
2.1	Installation Folders.....	5
2.2	Installing Intel® Integrated Performance Primitives Cryptography Libraries	6
2.3	Relocating Product After Install.....	6
2.4	Removal/Uninstall	6
3	Intel® C++ Compiler	7
3.1	New and Changed Features	7
3.2	New and Changed Compiler Options.....	7
3.2.1	–O0 no longer implies –mp	7
3.3	Other Changes	7
3.3.1	Optimization Reports Disabled by Default.....	7
3.3.2	Environment Setup Script Changed.....	8
3.3.3	OpenMP* Libraries Default to “compat”	8
3.3.4	OpenMP Header File Changed.....	8
4	Intel® Debugger (IDB)	9
4.1	Known Problems.....	9
4.1.1	Dwarf vs. Stabs Debug Formats	9
4.1.2	Compilation Requirements.....	9

4.1.3	Non-local Binary and Source File Access	9
4.1.4	Debugging applications that fork	10
4.1.5	Debugging applications that exec	10
4.1.6	Snapshots.....	10
4.1.7	Debugging optimized code.....	10
4.1.8	Watchpoints	10
4.1.9	Graphical User Interface (GUI)	10
4.1.10	MPP Debugging Restrictions	10
4.1.11	Function Breakpoints	10
4.1.12	Core File Debugging.....	11
4.1.13	Universal Binary Support	11
4.1.14	Debugger variable \$threadlevel	11
4.1.15	Open File Descriptors Limitation	11
4.1.16	\$cdir, \$cwd Directories	11
4.1.17	info stack Usage.....	11
4.1.18	\$stepg0 Default Value Changed.....	12
5	Intel® Integrated Performance Primitives	12
5.1	New and Changed Features	12
5.1.1	Intel® Integrated Performance Primitives 6.1 Update 4.....	12
5.1.2	Intel® Integrated Performance Primitives 6.1 Update 3.....	12
5.1.3	Intel® Integrated Performance Primitives 6.1 Update 2.....	12
5.2	Known Limitations.....	13
5.3	Intel® IPP Cryptography Libraries are Available as a Separate Download.....	13
5.4	Intel® IPP Code Samples	13
6	Intel® Math Kernel Library	13
6.1	Changes in This Version	13
6.1.1	Intel® Math Kernel Library 10.2 Update 4	13
6.1.2	Intel® Math Kernel Library 10.2 Update 3	14
6.1.3	Intel® Math Kernel Library 10.2 Update 2	15
6.2	Known Issues	17
6.3	Notices.....	17
6.4	Attributions.....	17
7	Intel® Threading Building Blocks	18

7.1	Changes in This Version	18
7.1.1	Intel® Threading Building Blocks 2.2 Update 2	18
7.1.2	Intel® Threading Building Blocks 2.2 Update 1	18
7.2	Known Issues	19
8	Disclaimer and Legal Information	20

1 Introduction

1.1 Change History

This section highlights important changes in product updates. For a list of corrections to reported problems, please read [Intel® Professional Edition Compilers 11.1 Fixes List](#), [Intel® IPP Library 6.1 Fixes List](#), [Intel® Math Kernel Library 10.2 Fixes List](#) and [Intel® Threading Building Blocks 2.2 Changes](#).

Update 5

- [Intel® Integrated Performance Primitives updated](#) to 6.1 Update 4
- [Intel® Math Kernel Library updated](#) to 10.2 Update 4
- [Intel® Threading Building Blocks Updated](#) to 2.2 Update 2
- Corrections to reported problems

Update 4 (11.1.080)

- [OpenMP header file changed](#) to improve error detection
- [Intel® Integrated Performance Primitives updated](#) to 6.1 Update 3
- [Intel® Math Kernel Library updated](#) to 10.2 Update 3
- Corrections to reported problems

Update 3 (11.1.076)

- [Intel® Threading Building Blocks updated](#) to 2.2 Update 1
- Corrections to reported problems

Update 2 (11.1.067)

- Added mention of [new compiler option `-mkl`](#)
- Mac OS* X 10.6.1 “Snow Leopard” is now supported
- Corrections to reported problems

Update 1 (11.1.058)

- Note added about [change in behavior of `-O0`](#)
- Corrections to reported problems

This document describes how to install the product, provides a summary of new and changed product features and includes notes about features and problems not described in the product documentation.

1.2 Product Contents

Intel® C++ Compiler Professional Edition 11.1 for Mac OS X* includes the following components:

- Intel® C++ Compilers for building applications that run on Intel-based Mac systems running the Mac OS* X operating system
- Intel® Debugger
- Intel® Integrated Performance Primitives 6.1 Update 4
- Intel® Math Kernel Library 10.2 Update 4
- Intel® Threading Building Blocks 2.2 Update 2
- Integration into the Xcode* development environment
- On-disk documentation

1.3 System Requirements

- An Intel®-based Apple* Mac* system
- 1GB RAM minimum, 2GB RAM recommended
- 3GB free disk space
- One of the following combinations of Mac OS* X, Xcode* and the Xcode SDK:
 - OS X 10.6.2 with Xcode 3.2.1 and SDK 10.6 or 10.5
 - OS X 10.6 with Xcode 3.2 and SDK 10.6 or 10.5
 - OS X 10.5.8 with Xcode 3.1.4 and SDK 10.5
 - OS X 10.5.6 with Xcode 3.1.2 and SDK 10.5
- gcc* 4

Note: Advanced optimization options or very large programs may require additional resources such as memory or disk space.

1.4 Documentation

Product documentation can be found in the `Documentation` folder as shown under [Installation Folders](#).

1.5 Technical Support

If you did not register your compiler during installation, please do so at the [Intel® Software Development Products Registration Center](#). Registration entitles you to free technical support, product updates and upgrades for the duration of the support term.

For information about how to find Technical Support, Product Updates, User Forums, FAQs, tips and tricks, and other support information, please visit:

<http://www.intel.com/software/products/support/>

Note: If your distributor provides technical support for this product, please contact them for support rather than Intel.

2 Installation

If you are installing the product for the first time, please be sure to have the product serial number available as you will be asked for it during installation. A valid license is required for installation and use.

If you will be using Xcode*, please make sure that a supported version of Xcode is installed. If you install a new version of Xcode in the future, you must reinstall the Intel C++ Compiler afterwards.

You will need to have administrative or “sudo” privileges to install, change or uninstall the product.

If you received the compiler product on DVD, insert the DVD. Locate the disk image file (`m_cproc_p_11.1.xxx.dmg`) on the DVD and double-click. If you received the compiler product as a download, double-click the downloaded file, which will have a name of the form `m_cproc_p_11.1.xxx.dmg`.

Follow the prompts to complete installation.

2.1 Installation Folders

The 11.1 product installs into a different arrangement of folders than in previous versions. The new arrangement is shown in the diagram below. Not all folders will be present in a given installation.

- `<root>/intel/Compiler/11.1/xxx/`
 - `bin`
 - `ia32`
 - `intel64`
 - `include`
 - `ia32`
 - `intel64`
 - `lib`
 - `perf_headers`
 - `Frameworks`
 - `ipp`
 - `mkl`
 - `tbb`
 - `Documentation`
 - `man`
 - `Samples`

Where `<root>` is `/opt` by default, `xxx` is the three-digit build number and the folders under `bin`, `include` and `lib` are used as follows:

- `ia32`: Compilers that build applications that run on a 32-bit Intel-based Mac OS* X system
- `intel64`: Compilers that build applications that run on a 64-bit Intel-based Mac OS* X system (also referred to as Intel® 64 architecture)

If you have both the Intel C++ and Intel Fortran compilers installed, they will share folders for a given version.

2.2 Installing Intel® Integrated Performance Primitives Cryptography Libraries

The Intel® Integrated Performance Primitives product provides an optional component containing libraries of cryptography routines. Installation and use of the cryptography libraries requires a separate license that is available at no charge from Intel once your license for Intel Integrated Performance Primitives has been registered. Export restrictions apply. For details, please see <http://software.intel.com/en-us/articles/intel-integrated-performance-primitives-cryptography-library/>

2.3 Relocating Product After Install

If you wish to move the installed product's command line interface to a different location on disk, you can do so using a supplied script.

1. Open a terminal window
2. Change directory (`cd`) to the compiler install folder (for example, `/opt/intel/Compiler/11.1/xxx`)
3. Type the command:
`./move_cproc.sh <new-install-location>`
where `<new-install-location>` is the new directory path

This script will move all the files and update symlinks, environment variables and startup scripts as needed. If you have both Intel C++ and Intel Fortran installed in the old path, both products will be moved to the new location.

The Xcode integration is relocatable simply by dragging and dropping the Xcode directory tree to another location. If you wish to use `idb` from a command prompt using a relocated Xcode directory tree, please see <http://software.intel.com/en-us/articles/running-idb-from-command-line-after-relocating-xcode-environment/> for additional steps that are required. Note that `idb` is not available from within the Xcode IDE.

2.4 Removal/Uninstall

It is not possible to remove the compiler while leaving any of the performance library components installed.

1. Open Terminal and set default (`cd`) to any folder outside `<install-dir>`

2. Type the command: `<install-dir>/uninstall_cproc.sh`
3. Follow the prompts

If you are not currently logged in as `root` you will be asked for the `root` password. If you also have the same-numbered version of Intel® Fortran Compiler installed, it may also be removed.

3 Intel® C++ Compiler

This section summarizes changes, new features and late-breaking news about the Intel C++ Compiler.

3.1 New and Changed Features

Please refer to the compiler documentation for details

- Additional features from C++ 0x
- C++ lambda functions
- Decimal floating point
- `valarray` implementation using IPP option
- `#pragma vector_nontemporal`
- `#pragma unroll_and_jam`
- Support for OpenMP* 3.0
- The default mode of the C++ compiler now more closely matches the default mode of gcc. Some C99 features, such as mixed declarations and code, may no longer be turned on by default, but can be enabled using `-std=c99`

3.2 New and Changed Compiler Options

- `-mkl[=lib]`

For a list of deprecated compiler options, see the Compiler Options section of the documentation.

3.2.1 `-O0` no longer implies `-mp`

In version 11.1, the `-O0` option for disabling optimizations no longer implies `-mp` for maximizing floating-point precision. The `-mp` switch is deprecated, so we recommend using an explicit `-fp-model` option for applications that are sensitive to floating-point precision changes.

3.3 Other Changes

3.3.1 Optimization Reports Disabled by Default

As of version 11.1, the compiler no longer issues, by default, optimization report messages regarding vectorization, automatic parallelization and OpenMP threaded loops. If you wish to see these messages you must request them by specifying `-diag-enable vec`, `-diag-enable par` and/or `-diag-enable openmp`, or by using `-vec-report`, `-par-report` and/or `-openmp-report`.

Also, as of version 11.1, optimization report messages are sent to `stderr` and not `stdout`.

3.3.2 Environment Setup Script Changed

The `icc.sh` (`icc.csh`) script, used to set up the command-line build environment, changed in version 11.0. In previous versions, you chose the target platform by selecting either the `cc` or `cce` directory root. In version 11.1, there is one version of these scripts and they now take an argument to select the target platform.

The command takes the form:

```
source /opt/intel/Compiler/11.1/xxx/bin/iccvars.sh argument
```

Where `xxx` is the package identifier and `argument` is either `ia32` or `intel64` as described above under [Installation Folders](#). If you have installed the compiler into a different path, make the appropriate adjustments in the command. Establishing the compiler environment also establishes the Intel® Debugger (`idb`) environment.

3.3.3 OpenMP* Libraries Default to “compat”

In version 10.1, a new set of OpenMP* libraries was added that allowed applications to use OpenMP code from both Intel and `gcc*` compilers. These “compatibility” libraries can provide higher performance than the older “legacy” libraries. In version 11.x, the compatibility libraries are used by default for OpenMP applications, equivalent to `-openmp-lib compat`. If you wish to use the older libraries, specify `-openmp-lib legacy`

The “legacy” libraries will be removed in a future release of the Intel compilers.

3.3.4 OpenMP Header File Changed

The OpenMP header file `omp.h` has been improved with additional error checking in version 11.1 Update 4.

The definitions of `omp_lock_t` and `omp_nest_lock_t` types have changed. With this release, the compiler distinguishes these types at compile time. This change will not affect OpenMP programs written in conformance with the OpenMP specification. However, a non-conforming OpenMP application may generate a compiler warning. For example:

```
$ cat sample.c
#include <omp.h>
int main() {
    omp_lock_t lk;
    omp_init_nest_lock( & lk );
    return 0;
} // main

$ icc -openmp sample.c

sample.c(4): warning #167: argument of type "omp_lock_t *" is
incompatible with parameter of type "omp_nest_lock_t *"
```



```
omp_init_nest_lock( & lk );  
                    ^
```

4 Intel® Debugger (IDB)

4.1 Known Problems

4.1.1 Dwarf vs. Stabs Debug Formats

The debugger only supports debugging of executables whose debug information is in Dwarf format, and does not support the Stabs debug format. Use the `-gdwarf-2` flag on the compile command to have gcc and g++ generate Dwarf output. The Intel compilers (icc and ifort) produce Dwarf debug format with the `-g` flag.

4.1.2 Compilation Requirements

Starting with Xcode 2.3, the Dwarf debugging information is stored in the object (.o) files. These object files are accessed by the debugger to obtain information related to the application being debugged and thus must be available for symbolic debugging.

In cases where a program is compiled and linked in one command, such as:

```
icc -g -o hello.exe hello.c
```

the object files are generated by the compiler but deleted before the command completes. The binary file produced by this command will have no debugging information. To make such an application debuggable users have two options.

Users may build the application in two steps, explicitly producing a .o file:

```
icc -c -g -o hello.o hello.c
```

```
icc -g -o hello.exe hello.o
```

Alternatively, users may use the compiler switch `-save-temps` to prevent the compiler from deleting the .o files it generates:

```
icc -g -save-temps -o hello.exe hello.c
```

The debugger does not use the output of the “dsymutil” utility.

4.1.3 Non-local Binary and Source File Access

The debugger cannot access binary files from a network-mounted file system (such as NFS). The error message will look like this:

```
Internal error: cannot create absolute path for: /home/me/hello
```

```
You cannot debug "/home/me/hello" because its type is "unknown".
```

The debugger cannot access source files from a network-mounted file system (such as NFS). The error message will look like this:

```
Source file not found or not readable, tried...  
  
./hello.c  
  
/auto/mount/site/foo/usr1/user_me/c_code/hello.c  
  
(Cannot find source file hello.c)
```

The file-path specified will be correct.

The workaround in both cases is to copy the files to a local file system (i.e., one which is not mounted over the network).

4.1.4 Debugging applications that fork

Debugging the child process of an application that calls fork is not yet supported.

4.1.5 Debugging applications that exec

The \$catchexecs control variable is not supported.

4.1.6 Snapshots

Snapshots are not yet supported as described in the manual.

4.1.7 Debugging optimized code

Debugging optimized code is not yet fully supported. The debugger may not be able to see some function names, parameters, variables, or the contents of the parameters and variables when code is compiled with optimizations turned on.

4.1.8 Watchpoints

Watchpoints that are created to detect write access don't trigger when a value identical to the original has been written. These restrictions are due to a limitation in the Mac OS* X operating system.

Because the SIGBUS signal rather than the SIGSEGV signal is used by the debugger to implement watchpoints, you cannot create a signal detector which will catch a SIGBUS signal.

4.1.9 Graphical User Interface (GUI)

This version of the debugger does not support the GUI

4.1.10 MPP Debugging Restrictions

MPP debugging is not supported as described in the manual.

4.1.11 Function Breakpoints

Debugger breakpoints set in functions (using the "stop in" command) may not halt user program execution at the first statement. This is due to insufficient information regarding the function prologue in the generated Dwarf debug information. As a work-around, use the "stop at" command to set a breakpoint on the desired statement.

The compiler generates a call to "`__dyld_func_lookup`" as part of the prologue for some functions. If you set a breakpoint on this function the debugger will stop there, but local variable values may not be valid. The work-around is to set a breakpoint on the first statement inside the function.

4.1.12 Core File Debugging

Debugging core files is not yet supported.

4.1.13 Universal Binary Support

Debugging of universal binaries is supported. The debugger supports debugging the IA-32 Dwarf sections of binaries on IA-32 and either the IA-32 or the Intel® 64 sections on Intel® 64.

4.1.14 Debugger variable `$threadlevel`

The manual's discussion of the debugger variable "`$threadlevel`" says "On Mac OS* X, the debugger supports POSIX threads, also known as pthreads." This sentence might be read as implying that other kinds of threads might be supported. This is not true; only POSIX threads are supported on Mac OS* X.

4.1.15 Open File Descriptors Limitation

Because the debugger opens the .o files of a debuggee to read debug information, you should raise the open file limit.

Mac OS* limits the number of open file descriptors to 256. You can increase this limit as follows:

```
ulimit -n 2000
```

Please use this command to increase the number of open descriptors before starting the debugger.

This is a workaround until the debugger can better share a limited number of open file descriptors over many files.

4.1.16 `$cdir`, `$cwd` Directories

`$cdir` is the compilation directory (if recorded). This is supported in that the directory is set; but `$cdir` is not itself supported as a symbol.

`$cwd` is the current working directory. Neither the semantics nor the symbol are supported.

The difference between `$cwd` and `'.'` is that `$cwd` tracks the current working directory as it changes during a debug session. `'.'` is immediately expanded to the current directory at the time an entry to the source path is added.

4.1.17 `info stack` Usage

The debugger command "info stack" does not currently support negative frame counts in the optional syntax below:

```
info stack [num]
```

A positive frame count num will print the innermost num frames. A negative or zero count will print no frames rather than the outermost num frames.

4.1.18 `$stepg0` Default Value Changed

The debugger has changed the default value of the debugger variable `$stepg0` from 1 to 0. With the value "0" the debugger will step over code without debug information if you do a "step" command. Set the debugger variable to 1 to have compatibility with previous debugger versions as follows:

```
(idb) set $stepg0 = 1
```

5 Intel® Integrated Performance Primitives

This section summarizes changes, new features and late-breaking news about the Intel® Integrated Performance Primitives (Intel® IPP) as part of Intel C++ Compiler Professional Edition. For detailed information about IPP see the following links:

- **New features:** see the information below and visit the main Intel IPP product page on the Intel web site at: <http://software.intel.com/en-us/intel-ipp>.
- **Documentation, help, and samples:** see the documentation links on the IPP product page at: <http://software.intel.com/en-us/intel-ipp>.

5.1 New and Changed Features

5.1.1 Intel® Integrated Performance Primitives 6.1 Update 4

- New string processing code examples in the IPP signal processing reference manual.
- Optimizations for RSA-1024 based decryption added to the library.
- OpenSSL performance improvements and support for version 0.9.8j of OpenSSL.

5.1.2 Intel® Integrated Performance Primitives 6.1 Update 3

- New code examples in chapter 11 of the IPP signal processing reference manual.
- UMC documentation now includes motion estimation and mode decision components.
- Approximate 5% performance improvement to the BZIP2 decoder.

5.1.3 Intel® Integrated Performance Primitives 6.1 Update 2

- Support Intel® Advanced Vector Extensions (Intel® AVX)
- Support Intel® Core™ i7 processor with new optimization and threading control/optimization
- 3D Image Processing: 3D Geometric Transforms, 3D Filters
- New Data Compression Functions APIs
- New Intel IPP Crypto support to RSA_SSA1.5 and RSA_PKCSv1.5
- Unified Image Classes (UIC) to add PNG format support and new features to support DXT1, DXT3, DXT5 texture compression
- Advanced lighting functions including Spherical Harmonic and Perlin Noise generation Functions

- Windows Media* Photo Support (HD Photo): IPP PCT Functions
- New video coding areas improvement including Scene Analyzer in MPEG-2, Intensity Compensation & Overlap Smoothing in VC1
- Samples for Signal Processing, Image Processing, String Processing and C++, C# language support have been added in /Samples folder. Others can be downloadable by clicking “Free Code Samples” at <http://software.intel.com/en-us/intel-ipp>
- The deprecated APIs have been added more reference information in reference manual and header files.

5.2 Known Limitations

- For a list of bug fixes, known issues, and limitations please see the following knowledge base article: <http://software.intel.com/en-us/articles/intel-ipp-library-61-fixes-list/>.

5.3 Intel® IPP Cryptography Libraries are Available as a Separate Download

The Intel® IPP cryptography libraries are available as a separate download. For download and installation instructions, please read

<http://software.intel.com/en-us/articles/download-ipp-cryptography-libraries/>

5.4 Intel® IPP Code Samples

The Intel® IPP code samples are organized into downloadable packages for Windows*, Linux* and Mac OS* at

<http://software.intel.com/en-us/articles/intel-integrated-performance-primitives-code-samples/>

The samples include source code for audio/video codecs, image processing and media player applications, and for calling functions from C++, C# and Java*. Instructions on how to build the sample are described in a readme file that comes with the installation package for each sample.

6 Intel® Math Kernel Library

This section summarizes changes, new features and late-breaking news about the Intel® Math Kernel Library (Intel® MKL) as part of Intel C++ Compiler Professional Edition.

6.1 Changes in This Version

For further information on improvements in this and previous releases, see

<http://software.intel.com/en-us/articles/new-in-intel-mkl-10-2/>

6.1.1 Intel® Math Kernel Library 10.2 Update 4

New Features

- Introduced the single precision complex absolute value function SCABS1
- Introduced the solver ?DTSVB for diagonally dominant tri-diagonal systems which is up to 2x faster than the general solver with partial pivoting (?GTSV)
- Added routines for factorization (?DTTRFB) and the forward/backward substitution (?DTTRSB) of the diagonally dominant tri-diagonal systems

Performance improvements

- FFTs
 - Enhanced performance for transforms which are a multiple of 8 or 13
 - Optimized 1D complex cluster FFTs for non-power-of-2 vector lengths
- VSL
 - Convolution and Correlation computations that require decimation show significant improvements (re-link required, see [Known Issues](#))

6.1.2 Intel® Math Kernel Library 10.2 Update 3

Performance Improvements

- BLAS
 - Threaded the 32-bit OS versions of the following BLAS Level 1 and 2 functions for Intel® Core™ i7 processors and Intel® Xeon® processor 5300, 5400, and 5500 series: (D,S,C,Z)COPY, (D,S,C,Z)SWAP, (D,S,C,Z)AXPY, (S,C)ROT, (S,C)DOT, CDOTC, (D,S,C,Z)GEMV, (D,S,C,Z)TRMV, (S,C)SYMV, (S,C)SYR, (S,C)SYR2
 - Improved 32-bit and 64-bit OS versions of the following BLAS level 1 functions for Intel® Xeon® processors 5300, 5400, 5500: ZAXPY, ZSCAL, ZDOT(U,C), and (D,S)ROT
 - Improved DGEMM threading efficiency for matrices with many more rows than columns for Intel® Xeon® processor 5300
- LAPACK
 - Improved scalability of the following LAPACK functions: ?POTRF, ?GEBRD, ?SYTRD, ?HETRD, and ?STEDC divide and conquer eigensolvers
- FFTs
 - Updated underlying kernels to provide widespread performance improvements in FFTs
 - Improved threading of 3D FFTs when a small number of transforms are calculated with a single function call
 - Extended threading to small size multidimensional transforms
- VML
 - Further optimization for these VML functions on Intel® Xeon® processor 5500 series: v(s,d)Asin, v(s,d)Acos, v(s,d)Ln, v(s,d)Log10, vsLog1p, v[s/d]Hypot
- VSL
 - Improved performance of viRngPoisson and viRngPoissonV random number generators

Usability and Interface Improvements

- Improved example programs for uBLAS, Java, FFTW3, LAPACK95, and BLAS95
- Some examples in the reference manual were removed where identical examples in source code form also appeared in the examples directory

- New 64-bit integer (ILP64) `fftw_mpi` interfaces for cluster FFTs

6.1.3 Intel® Math Kernel Library 10.2 Update 2

New Features

- LAPACK 3.2
 - 238 new LAPACK functions
 - Extra Precise Iterative Refinement
 - Non-Negative Diagonals from Householder QR factorization
 - High Performance QR and Householder Reflections on Low-Profile Matrices
 - New fast and accurate Jacobi SVD
 - Routines for Rectangular Full Packed format
 - Pivoted Cholesky
 - Mixed precision iterative refinement (Cholesky)
 - More robust DQDS algorithm
- Introduced implementation of the DZGEMM Extended BLAS function (as described at <http://www.netlib.org/blas/blast-forum/>). See the description of the `*gemm` family of functions in the BLAS section of the reference manual.
- PARDISO now supports real and complex, single precision data

Usability/Interface improvements

- Sparse matrix format conversion routines which convert between the following formats:
 - CSR (3-array variation) ↔ CSC (3-array variation)
 - CSR (3-array variation) ↔ diagonal format
 - CSR (3-array variation) ↔ skyline
- Fortran95 BLAS and LAPACK compiled module files (`.mod`) are now included
 - Modules are pre-built with the Intel Fortran Compiler and are located in the include directory (see Intel® MKL User's Guide for full path)
 - Source is still available for use with other compilers
 - Documentation for these interfaces can be found in the Intel® MKL User's Guide
- The FFTW3 interface is now integrated directly into the main libraries
 - Source code is still available to create wrappers for use with compilers not compatible with the default Intel® Fortran compiler convention for name decoration
 - See Appendix G of the Reference Manual for information
- `DFTI_DESCRIPTOR_HANDLE` now represents a true type name and can now be referenced as a type in user programs
- Added parameter to Jacobi matrix calculation routine in the optimization solver domain to allow access to user data (see the description of the `djacobi` function in the reference manual for more information)

- Added an interface mapping calls to single precision BLAS functions in Intel® MKL (functions with “s” or “c” initial letter) to 64-bit floating point precision functions has been added on 64-bit architectures (See “sp2dp” in the Intel® MKL User Guide for more information)
- Compatibility libraries (also known as “dummy” libraries) have been removed from this version of the library

Performance improvements

- Further threading in BLAS level 1 and 2 functions for Intel® 64 architecture
 - Level 1 functions (vector-vector): (C,Z,S,D)ROT, (C,Z,S,D)COPY, and (C,Z,S,D)SWAP
 - Increase in performance by up to 1.7-4.7 times over version 10.1 Update 1 on 4-core Intel® Core™ i7 processor depending on data location in cache
 - Increase in performance by up to 14-130 times over version 10.1 Update 1 on 24-core Intel® Xeon® processor 7400 series system, depending on data location in cache
 - Level 2 functions (matrix-vector): (C,Z,S,D)TRMV, (S,D)SYMV, (S,D)SYR, and (S,D)SYR2
 - Increase in performance by up to 1.9-2.9 times over version 10.1 Update 1 on 4-core Intel® Core™ i7 processor, depending on data location in cache
 - Increase in performance by up to 16-40 times over version 10.1 Update 1 on 24-core Intel® Xeon® processor 7400 series system, depending on data location in cache
- Introduced recursive algorithm in 32-bit sequential version of DSYRK for up to 20% performance improvement on Intel® Core™ i7 processors and Intel® Xeon® processors in 5300, 5400, and 7400 series.
- Improved LU factorization (DGETRF) by 25% over Intel MKL 10.1 Update 1 for large sizes on the Intel® Xeon® 7460 Processor; small sizes are also dramatically improved
- BLAS *TBMV/*TBSV functions now use level 1 BLAS functions to improve performance by up to 3% on Intel® Core™ i7 processors and up to 10% on Intel® Core™2 processor 5300 and 5400 series.
- Improved threading algorithms to increase DGEMM performance
 - up to 7% improvement on 8 threads and up to 50% on 3,5,7 threads on the Intel® Core™ i7 processor
 - up to 50% improvement on 3 threads on Intel® Xeon® processor 7400 series.
- Threaded 1D complex-to-complex FFTs for non-prime sizes
- New algorithms for 3D complex-to-complex transforms deliver better performance for small sizes (up to 64x64x64) on 1 or 2 threads
- Implemented high-level parallelization of out-of-core (OOC) PARDISO when operating on symmetric positive definite matrices.

- Reduced memory use by PARDISO for both in-core and out-of-core on all matrix types
- PARDISO OOC now uses less than half the memory previously used in Intel MKL 10.1 for real symmetric, complex Hermitian, or complex symmetric matrices
- Parallelized Reordering and Symbolic factorization stage in PARDISO/DSS
- Up to 2 times better performance (30% improvement on average) on Intel® Core® i7 and Intel® Core™2 processors for the following VML functions: $v(s,d)Round$, $v(s,d)Inv$, $v(s,d)Div$, $v(s,d)Sqrt$, $v(s,d)Exp$, $v(s,d)Ln$, $v(s,d)Atan$, $v(s,d)Atan2$
- Optimized versions of the following functions available for Intel® Advanced Vector Extensions (Intel® AVX)
 - BLAS: DGEMM
 - FFTs
 - VML: exp, log, and pow
 - See important information in the Intel® MKL User's Guide regarding the `mkl_enable_instructions()` function for access to these functions

6.2 Known Issues

A full list of the known limitations of this release can be found in the Knowledge Base for the Intel® MKL at <http://software.intel.com/en-us/articles/known-limitations-in-intel-mkl-10-2>

6.3 Notices

The following change is planned for future versions of Intel MKL. Please contact [Technical Support](#) if you have concerns:

- Content in the libraries containing `solver` in the filenames will be moved to the core library in a future version of Intel MKL. These `solver` libraries will then be removed.

6.4 Attributions

As referenced in the End User License Agreement, attribution requires, at a minimum, prominently displaying the full Intel product name (e.g. "Intel® Math Kernel Library") and providing a link/URL to the Intel® MKL homepage (www.intel.com/software/products/mkl) in both the product documentation and website.

The original versions of the BLAS from which that part of Intel® MKL was derived can be obtained from <http://www.netlib.org/blas/index.html>.

The original versions of LAPACK from which that part of Intel® MKL was derived can be obtained from <http://www.netlib.org/lapack/index.html>. The authors of LAPACK are E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. Our FORTRAN 90/95 interfaces to LAPACK are similar to those in the LAPACK95 package at <http://www.netlib.org/lapack95/index.html>. All interfaces are provided for pure procedures.

The original versions of ScaLAPACK from which that part of Intel® MKL was derived can be obtained from <http://www.netlib.org/scalapack/index.html>. The authors of ScaLAPACK are

L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley.

PARDISO in Intel® MKL is compliant with the 3.2 release of PARDISO that is freely distributed by the University of Basel. It can be obtained at <http://www.pardiso-project.org>.

Some FFT functions in this release of Intel® MKL have been generated by the SPIRAL software generation system (<http://www.spiral.net/>) under license from Carnegie Mellon University. Some FFT functions in this release of the Intel® MKL DFTI have been generated by the UHFFT software generation system under license from University of Houston. The Authors of SPIRAL are Markus Puschel, Jose Moura, Jeremy Johnson, David Padua, Manuela Veloso, Bryan Singer, Jianxin Xiong, Franz Franchetti, Aca Gacic, Yevgen Voronenko, Kang Chen, Robert W. Johnson, and Nick Rizzolo.

7 Intel® Threading Building Blocks

This section summarizes changes, new features and late-breaking news about Intel® Threading Building Blocks (Intel® TBB) as part of Intel® C++ Compiler Professional Edition.

7.1 Changes in This Version

7.1.1 Intel® Threading Building Blocks 2.2 Update 2

- `parallel_invoke` and `parallel_for_each` now take function objects by const reference, not by value.
- Improvements in exception support: new exception classes introduced, all exceptions are thrown via an out-of-line internal method.
- Improvements and fixes in the TBB allocator and malloc replacement, including robust memory identification, and more reliable dynamic function substitution on Windows*.
- Method `swap()` added to class `tbb_thread`.
- Methods `rehash()` and `bucket_count()` added to `concurrent_hash_map`.
- Other fixes and improvements in code, tests, examples, and docs.

7.1.2 Intel® Threading Building Blocks 2.2 Update 1

- Documentation was updated.
- TBB scheduler auto-initialization now covers all possible use cases.
- `concurrent_queue`: made argument types of `sizeof` used in padding consistent with those actually used.
- Memory allocator was improved: supported corner case of user's malloc calling `scalable_malloc` (non-Windows), corrected processing of memory allocation requests during tbb memory allocator startup (Linux).
- Windows malloc replacement has got better support for static objects.
- In pipeline setups that do not allow actual parallelism, execution by a single thread is guaranteed, idle spinning eliminated, and performance improved.
- RML refactoring and clean-up.

- New constructor for `concurrent_hash_map` allows reserving space for a number of items.
- Operator `delete()` added to the TBB exception classes.
- Lambda support was improved in `parallel_reduce`.
- gcc 4.3 warnings were fixed for `concurrent_queue`.
- Fixed possible initialization deadlock in modules using TBB entities during construction of global static objects.
- Copy constructor in `concurrent_hash_map` was fixed.
- Fixed a couple of rare crashes in the scheduler possible before in very specific use cases.
- Fixed a rare crash in the TBB allocator running out of memory.
- New tests were implemented, including `test_lambda.cpp` that checks support for lambda expressions.
- Fixed known exception safety issues in `concurrent_vector`.
- Better concurrency of simultaneous grow requests in `concurrent_vector`.
- TBB allocator further improves performance of large object allocation.
- Problem with source of text relocations was fixed on Linux
- Fixed bugs related to `malloc` replacement under Windows
- A few other small changes in code, tests, and documentation.

7.2 Known Issues

Please note the following with respect to this particular release of Intel Threading Building Blocks.

- Unhandled exceptions in the user code executed in the context of TBB algorithms or containers may lead to segmentation faults when Intel(R) C++ Compiler 10.x is used with `glibc` 2.3.2, 2.3.3, or 2.3.4.
- To allow more accurate results to be obtained with Intel® Thread Checker or Intel® Thread Profiler, download the latest update releases of these products before using them with Intel® Threading Building Blocks.
- If you are using Intel® Threading Building Blocks and OpenMP* constructs mixed together in rapid succession in the same program, and you are using Intel compilers for your OpenMP* code, set `KMP_BLOCKTIME` to a small value (e.g., 20 milliseconds) to improve performance. This setting can also be made within your OpenMP* code via the `kmp_set_blocktime()` library call. See the compiler OpenMP* documentation for more details on `KMP_BLOCKTIME` and `kmp_set_blocktime()`.
- In general, non-debug ("release") builds of applications or examples should link against the non-debug versions of the Intel® Threading Building Blocks libraries, and debug builds should link against the debug versions of these libraries. See the Tutorial in the product documentation sub-directory for more details on debug vs. release libraries.

8 Disclaimer and Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL(R) PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site.

MPEG-1, MPEG-2, MPEG-4, H.263, H.264, MP3, DV SD/25/50/100, VC-1, G.722.1, G.723.1A, G.726, G.728, G.729, GSM/AMR, GSM/FR, JPEG, JPEG 2000, Aurora, TwinVQ, AC3 and AAC are international standards promoted by ISO, IEC, ITU, SMPTE, ETSI and other organizations. Implementations of these standards or the standard enabled platforms may require licenses from various entities, including Intel Corporation.

Celeron, Centrino, Intel, Intel logo, Intel386, Intel486, Intel Atom, Intel Core, Itanium, MMX, Pentium, VTune, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2010 Intel Corporation. All Rights Reserved.