



Intel® Integrated Performance Primitives for Windows* OS

User's Guide

IPP 7.0 Gold

318254-010US

Legal Information

Contents

Legal Information	5
What's New	7
Getting Help and Support	9
Introducing the Intel® Integrated Performance Primitives	11
Notational Conventions	13
Chapter 1: Intel(R) Integrated Performance Primitives Basics	
Cross-Architecture Alignment.....	15
Types of Input Data.....	16
Domains.....	17
Function Naming and Parameters.....	18
Chapter 2: Getting Started with Intel(R) IPP	
Checking Your Installation.....	23
Obtaining Version Information.....	23
Building Your Application.....	24
Setting Environment Variables.....	24
Including Header Files.....	25
Calling Intel IPP Functions.....	25
Before You Begin Using Intel IPP.....	25
Chapter 3: Intel(R) IPP Structure	
Supplied Libraries.....	28
Using Intel IPP Dynamic Link Libraries (DLLs).....	28
Using Intel IPP Static Libraries.....	29
Contents of the Documentation Directory.....	30
Chapter 4: Configuring Your Development Environment	
Configuring Your Microsoft* Development Environment to Link with Intel IPP.....	31
Creating Visual C++ 2005 Project Files for the Intel IPP Samples.....	32
Using the IntelliSense* Features.....	32
Chapter 5: Linking Your Application with Intel(R) IPP	
Dispatching.....	35

Detecting Processor Features and Type.....	36
Selecting Between Linking Methods.....	39
Dynamic Linking.....	39
Static Linking with Dispatching.....	40
Static Linking without Dispatching.....	42
Building a Custom DLL.....	43
Comparison of Intel IPP Linking Methods.....	44
Selecting the Intel(R) IPP Libraries Needed by Your Application.....	44
Libraries for Dynamic Linking.....	46
Libraries for Static Linking.....	47
Library Dependencies by Domain.....	48
Linking Examples.....	48

Chapter 6: Supporting Multithreaded Applications

Intel IPP Threading and OpenMP* Support.....	49
Setting Number of Threads.....	50
Using a Shared L2 Cache.....	50
Avoiding Nested Parallelization.....	51

Chapter 7: Managing Performance and Memory

Memory Alignment.....	53
Thresholding Data.....	54
Reusing Buffers.....	55
Using Fast Fourier Transform.....	56
Using the Intel IPP Performance Test Tool.....	57
Examples of Performance Test Tool Command Lines.....	58

Appendix A: Performance Test Tool Command Line Options

Appendix B: Intel(R) IPP Samples

Types of Intel IPP Sample Code.....	65
Source Code Samples.....	65
Using Intel IPP Samples.....	66
System Requirements.....	67
Building Source Code.....	68
Running the Software.....	69
Known Limitations.....	69

Appendix C: Language Support

Using Intel IPP in Java* Applications.....	72
--------------------------------------------	----

Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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 go to

<http://www.intel.com/design/literature.htm>

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

MPEG-1, MPEG-2, MPEG-4, H.261, H.263, H.264, MP3, DV, VC-1, MJPEG, AC3, AAC, G.711, G.722, G.722.1, G.722.2, AMRWB, Extended AMRWB (AMRWB+), G.167, G.168, G.169, G.723.1, G.726, G.728, G.729, G.729.1, GSM AMR, GSM FR are international standards promoted by ISO, IEC, ITU, ETSI, 3GPP and other organizations. Implementations of these standards, or the standard enabled platforms may require licenses from various entities, including Intel Corporation.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Inside, Cilk, Core Inside, i960, Intel, the Intel logo, Intel AppUp, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel Sponsors of Tomorrow., the Intel Sponsors of Tomorrow. logo, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, InTru, the InTru logo, InTru soundmark, Itanium, Itanium Inside, MCS, MMX, Moblin, Pentium, Pentium Inside, skool, the skool logo, Sound Mark, The Journey Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

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

Microsoft, Windows, Visual Studio, Visual C++, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Microsoft product screen shot(s) reprinted with permission from Microsoft Corporation.

What's New

The User's Guide provides quick guidance of using the Intel(R) Integrated Performance Primitives (Intel(R) IPP) for all supported architectures on the Windows* OS.

The User's Guide documents updates of Intel IPP for 7.0 release, including new library structure, directory paths and names, and new names of the library files.

Getting Help and Support

Getting Help



The online version of the Intel(R) Integrated Performance Primitives (Intel(R) IPP) documentation integrates into the Microsoft Visual Studio* IDE help system.

Browsing Help in the Microsoft Document Explorer

You can browse and search for topics in different ways:

- Use the **Contents** tab to browse through topics by functional domains.
- Use the **Index** tab to access an index to all topics. Either type the word you are looking for or scroll through the list.
- Click the **Search** button on the toolbar to perform a full text search.

To navigate between topics, use the navigation buttons available in the Document Explorer toolbar or from the topic pop-up menu:

- Click the **Back**  button to return to the previously viewed topic.
- Click the **Forward**  button to go to the following topic.

Using Context-sensitive Help

Context-sensitive help enables easy access to the description of a function whose name is typed in the Code Editor. You can use *F1 Help* and/or *Dynamic Help*.

To enable Dynamic Help, go to **Help > Dynamic Help**. The **Dynamic Help** window opens that displays links relevant to your current selection.


To access the function description,

1. Select the function name in the Code Editor
2. Do one of the following:
 - Click **F1**
 - Click the link to the description in the Dynamic Help window.

The function description opens in Document Explorer.

Locating Intel Topics

To display Intel documentation in Visual Studio 2005 and 2008, select **Contents** from the Visual Studio Help menu. Then, select **Intel** in the choices under **Filtered by**. In navigating the Intel documentation, you may want to know where the topic you are reading is located in the table of contents (TOC). Click the **Sync with**

Table of Contents  button on the Visual Studio toolbar to find where the topic is located in the TOC. You can also use **Sync with Table of Contents** with the context-sensitive help topics to locate them in the Intel documentation TOC.

Getting Technical Support

If you did not register your Intel software product during installation, please do so now 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 general information about Intel technical support, product updates, user forums, FAQs, tips and tricks and other support questions, please visit <http://www.intel.com/software/products/support/>.



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

For technical information about the Intel IPP library, including FAQ's, tips and tricks, and other support information, please visit the Intel IPP forum: <http://software.intel.com/en-us/forums/intel-integrated-performance-primitives/> and browse the Intel IPP knowledge base: <http://software.intel.com/en-us/articles/intel-ipp-kb/all/>.

Introducing the Intel® Integrated Performance Primitives

Intel® Integrated Performance Primitives (Intel® IPP) is a software library that provides a broad range of functionality. This functionality includes general signal and image processing, computer vision, data compression, cryptography, string manipulation, audio processing, video coding, realistic rendering and 3D data processing. It also includes more sophisticated primitives for construction of audio, video and speech codecs such as MP3 (MPEG-1 Audio, Layer 3), MPEG-4, H.264, H.263, JPEG, JPEG2000, JPEG XR, GSM-AMR, G.723.

By supporting a variety of data types and layouts for each function, the Intel IPP library delivers a rich set of options for developers to choose from when designing and optimizing an application.

To speed up performance, Intel IPP functions are optimized to use all benefits of Intel® architecture processors. In addition most of the Intel IPP functions do not use complicated data structures, which helps reduce overall execution overhead.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Notational Conventions

The document uses the following font conventions and symbols:

Notational conventions

<i>Italic</i>	<i>Italic</i> is used for emphasis and also indicates document names in body text, for example: see <i>Intel IPP Reference Manual</i>
Monospace lowercase	Indicates filenames, directory names, and pathnames, for example: <code>\tools\ia32\perfsys\ps_ippi.exe</code>
Monospace lowercase mixed with UPPERCASE	Indicates commands and command-line options, for example: <code>ps_ipp.exe -f FIRLMS_32f -r firlms.csv</code>
UPPERCASE MONOSPACE	Indicates system variables, for example, <code>\$PATH</code>
monospace italic	Indicates a parameter in discussions, such as routine parameters, for example, <i>pSrc</i> ; makefile parameters, for example, <i>function_list</i> , and so forth. When enclosed in angle brackets, indicates a placeholder for an identifier, an expression, a string, a symbol, or a value, for example, <i><ipp directory></i> .
[items]	Square brackets indicate that the items enclosed in brackets are optional.
{ item item }	Braces indicate that only one of the items listed between braces can be selected. A vertical bar () separates the items.

Intel(R) Integrated Performance Primitives Basics

1

Intel(R) Integrated Performance Primitives (Intel(R) IPP) has the following features:

- Intel IPP provides basic low-level functions for creating applications in several different domains, such as signal processing, audio coding, speech coding, image processing, video coding, operations on small matrices, and realistic rendering functionality and 3D data processing.
- The Intel IPP functions follow the same interface conventions including uniform naming rules and similar composition of prototypes for primitives that refer to different application domains.
- The Intel IPP functions use abstraction level which enables you to improve the performance of your application.

Intel IPP is well-suited for cross-platform applications. For example, the functions developed for IA-32 architecture-based platforms can be readily ported to Intel(R) 64 architecture-based platforms.

Intel IPP for Windows* OS has the following versions:

- Intel IPP for the Windows OS on IA-32 architecture
- Intel IPP for the Windows OS on Intel(R) 64 architecture

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Cross-Architecture Alignment

Intel IPP is designed to support application development on various Intel(R) architectures. This means that the API definition is common for all processors, while the underlying function implementation takes into account the variations in processor architectures.

By providing a single cross-architecture API, Intel IPP allows software application repurposing and enables you to port features across Intel(R) processor-based desktop, server, and mobile platforms. You can use your code developed for one processor architecture for many processor generations.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

See also [Dispatching](#).

Types of Input Data

Intel IPP operations are divided into groups according to the input data on which the operation is performed. Each group has its own prefix in the function name (see [Function Naming - Data-Domain](#)). The input data types are:

One-Dimensional Arrays and Signals

This group includes most functions operating on one-dimensional arrays of data. Generally these arrays are signals, and many of the operations are signal-processing operations. Examples of one-dimensional array operations include:

- vectorized scalar arithmetic, logical, statistical operations
- digital signal processing
- data compression
- audio processing and audio coding
- speech coding
- cryptography and data integrity
- string operations

Images

An image is a two-dimensional array of pixels. Their specific features distinguish them from general two-dimensional array. Examples of image operations include:

- arithmetic, logical, statistical operations
- color conversion
- image filtering
- image linear and geometric transformations
- morphological operations
- computer vision
- image compression
- video coding

Matrices

This group includes functions operating on matrices and vectors that are one- and two-dimensional arrays, and on arrays of matrices and vectors. These arrays are treated as linear equations or data vectors and subjected to linear algebra operations. Examples of matrix operations include:

- vector and matrix algebra
- solving systems of linear equations
- solving least squares problem
- computing eigenvalue problem

3D objects

This group includes functions operating with 3D objects. In this case input data depends on the used techniques. Examples of 3D operations include:

- realistic rendering
- resizing and affine transforming

Core Functions

There are several core functions that do not perform operations on one of these input data types. Examples of such operations include getting the type of CPU, aligning pointers to the specified number of bytes, controlling the dispatcher of the merged static libraries. These functions have their own header file, static libraries and DLLs.

Code	Header file	Static Libraries	DLL	Prefix in Function Name
ippCore	ippcore.h	ippcore_1.lib ippcore_t.lib	ippcore-x.x.dll	ipp

Here *x.x* refers to the product version number, for example 7.0.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Domains

Intel IPP is divided into subdivisions of related functions. Each subdivision is called *domain*, and has its own header file, static libraries, DLLs, and tests. These domains map to the [types of input data](#) and the corresponding [prefixes](#). The Intel IPP Manual indicates in which header file each function can be found. The table below lists each domain's code, header and library names, and functional area.

Code of Domain	Header file	Static Libraries	DLL	Prefix	Description
ippAC	ippac.h	ippac[_*].lib	ippac[**]-x.x.dll	ipps	audio coding
ippCC	ippcc.h	ippcc[_*].lib	ippcc[**]-x.x.dll	ippi	color conversion
ippCH	ippch.h	ippch[_*].lib	ippch[**]-x.x.dll	ipps	string operations
ippCP	ippcp.h	ippcp[_*].lib	ippcp[**]-x.x.dll	ipps	cryptography
ippCV	ippcv.h	ippcv[_*].lib	ippcv[**]-x.x.dll	ippi	computer vision
ippDC	ippdc.h	ippdc[_*].lib	ippdc[**]-x.x.dll	ipps	data compression
ippDI	ippdi.h	ippdi[_*].lib	ippdi[**]-x.x.dll	ipps	data integrity
ippIP	ippi.h	ippi[_*].lib	ippi[**]-x.x.dll	ippi	image processing
ippJP	ippj.h	ippj[_*].lib	ippj[**]-x.x.dll	ippi	image compression
ippMX	ippm.h	ippm[_*].lib	ippm[**]-x.x.dll	ippm	small matrix operations
ippRR	ippr.h	ippr[_*].lib	ippr[**]-x.x.dll	ippr	realistic rendering and 3D data processing
ippSP	ipps.h	ipps[_*].lib	ipps[**]-x.x.dll	ipps	signal processing
ippSC	ippsc.h	ippsc[_*].lib	ippsc[**]-x.x.dll	ipps	speech coding
ippVC	ippvc.h	ippvc[_*].lib	ippvc[**]-x.x.dll	ippi	video coding
ippVM	ippvm.h	ippvm[_*].lib	ippvm[**]-x.x.dll	ipps	vector math

* refers to one of the following: l, t.

** refers to [processor-specific code](#), for example, s8.

x.x refers to the product version number, for example, 7.0

Function Naming and Parameters

Function names in Intel IPP use naming conventions to help you identify the different functions.

Intel IPP function names include a number of fields that indicate the data domain, operation, data type, and execution mode. Each field has a fixed number of pre-defined values.

Function names use the following format:

```
ipp<data-domain><name>[_<datatype>][_<descriptor>](<parameters>);
```

The elements of this format are explained in the sections that follow:

[Data-Domain](#)

[Name](#)

[Data Types](#)

[Descriptor](#)

[Parameters](#)

Data-Domain

The *data-domain* is a single character indicating [type of the input data](#). The current version of Intel IPP supports the following data-domains:

s	for signals (expected data type is a 1D array)
g	for signals of the fixed length (expected data type is a 1D array)
i	for images and video (expected data type is a 2D array of pixels)
m	for vectors and matrices (expected data type is a matrix or vector)
r	for realistic rendering functionality and 3D data processing (expected data type depends on supported rendering techniques)

The *core functions* in Intel IPP do not operate on one of these types of the input data (see [Core Functions](#)). These functions have `ipp` as a prefix without the data-domain field, for example, `ippGetString`.

Name

The *name* identifies the algorithm or operation that the function does. It has the following format:

```
<name> = <operation>[_<modifier>]
```

The *operation* field is one or more words, acronyms, and abbreviations that identify the base operation, for example `Set`, `DCTFwd`. If the operation consists of several parts, each part starts with an uppercase character without underscore, for example, `HilbertInitAlloc`.

The *modifier*, if present, denotes a slight modification or variation of the given function. For example, the modifier `CToC` in the function `ippsFFTInv_CToC_32fc` signifies that the inverse fast Fourier transform operates on complex data, performing complex-to-complex (CToC) transform. Functions for matrix operation have an object type description as a modifier, for example, `ippmMul_mv` - multiplication of a matrix by a vector.

Data Types

The *datatype* field indicates data types used by the function in the following format:

```
<datatype> = <bit_depth><bit_interpretation>
```

where

`bit_depth = <1|8|16|32|64>`

and

`bit_interpretation = < u|s|f>[c]`

Here `u` indicates "unsigned integer", `s` indicates "signed integer", `f` indicates "floating point", and `c` indicates "complex".

For functions that operate on a single data type, the `datatype` field contains only one value.

If a function operates on source and destination objects that have different data types, the respective data type identifiers are listed in the function name in order of source and destination as follows:

`<datatype> = <src1Datatype>[src2Datatype][dstDatatype]`

For example, the function `ippsDotProd_16s16sc` computes the dot product of 16-bit short and 16-bit complex short source vectors and stores the result in a 16-bit complex short destination vector. The `dstDatatype` modifier is not present in the name because the second operand and the result are of the same type.

Descriptor

The optional `descriptor` field describes the data associated with the operation. It can contain implied parameters and/or indicate additional required parameters.

To minimize the number of code branches in the function and thus reduce potentially unnecessary execution overhead, most of the general functions are split into separate primitive functions, with some of their parameters entering the primitive function name as descriptors.

However, general functions with large number of permutations may still have parameters that determine internal operation (for example, `ippiThreshold`).

The following descriptors are used in Intel IPP:

A	Image data contains an alpha channel as the last channel, requires C4 descriptor, the alpha channel is not processed.
A0	Image data contains an alpha channel as the first channel, requires C4 descriptor, the alpha channel is not processed.
Axx	Advanced arithmetic operations with <code>xx</code> bits of accuracy.
C	The function operates on a specified channel of interest (COI) for each source image.
Cn	Image data is made of <code>n</code> discrete interleaved channels (<code>n = 1, 2, 3, 4</code>)
Dx	Signal is <code>x</code> -dimensional (default is D1)
I	Operation is performed in-place (default is not-in-place).
L	For signal processing indicates that one pointer is used for each row in two-dimensional array. For matrix operation: function with layout description of the objects.
M	Operation uses a mask to determine the pixels to be processed.

P	F or signal processing: specified number of vectors to be processed. For matrix operation: function with pointer description of the objects.
P _n	Image data is made of <i>n</i> discrete planar (non-interleaved) channels (<i>n</i> = 1, 2, 3, 4) with a separate pointer to each plane.
R	Function operates on a defined region of interest (ROI) for each source image.
S	Function with standard description of the objects (for matrix operation)
s	Saturation and no scaling mode (default mode)
Sfs	Saturation and fixed scaling mode.

The descriptors in function names are always presented in alphabetical order.

Some data descriptors are implied as "default" for certain operations. Such descriptors are not added to the function names. For example, the image processing functions always operate on a two-dimensional image and saturate the results without scaling them. In these cases, the implied descriptors D2 (two-dimensional signal) and s (saturation and no scaling) are not included in the function name.

Parameters

The `parameters` field specifies the function parameters (arguments).

The order of parameters is as follows:

1. All source operands. Constants follow arrays
2. All destination operands. Constants follow arrays
3. Other, operation-specific parameters

The parameter's name has the following conventions:

- All arguments defined as pointers start with *p*, for example, *pPhase*, *pSrc*, *pSeed*; arguments defined as double pointers start with *pp*, for example, *ppState*; and arguments defined as values start with a lowercase letter, for example, *val*, *src*, *srcLen*.
- Each new part of an argument name starts with an uppercase character, without underscore, for example, *pSrc*, *lenSrc*, *pDlyLine*.
- Each argument name specifies its functionality. Source arguments are named *pSrc* or *src*, sometimes followed by numbers or names, for example, *pSrc2*, *srcLen*. Output arguments are named *pDst* or *dst*, sometimes followed by numbers or names, for example, *pDst1*, *dstLen*. For in-place operations, the input/output argument contains the name *pSrcDst*.

Examples of function syntax:

- `ippsIIR_32f_I(Ipp32f* pSrcDst, int len, IppsIIRState_32f* pState);`
- `ippiConvert_8ulu_C1R(const Ipp8u* pSrc, int srcStep, Ipp8u* pDst, int dstStep, int dstBitOffset, IppiSize roiSize, Ipp8u threshold);`

- `ippmSub_vac_32f(const Ipp32f* pSrc, int srcStride0, int srcStride2, Ipp32f val, Ipp32f* pDst, int dstStride0, int dstStride2, int len, int count);`

Getting Started with Intel(R) IPP

This section helps you start using Intel(R) Integrated Performance Primitives (Intel(R) IPP) by giving basic information you need to know and describing the necessary steps you need to follow after installation of the product.

Checking Your Installation

After completing the installation of Intel IPP, confirm proper installation and configuration of the library with the following steps:

1. Check the installation directory. The default installation directory is `C:\Program Files\Intel\composerxe-2011`.
2. Check that file `ippvars.bat` is placed in the `bin` directory. You can use this file to set the environment variables `PATH`, `LIB`, and `INCLUDE` in the user shell.
3. Check that the dispatching and processor-specific libraries are on the path.
4. Run `ippiDemo.exe` (or `ippsDemo.exe`) from the `<ipp directory>\demo\ia32`, or run `ippiDemo_em64t.exe` (or `ippsDemo_em64t.exe`) from the `<ipp directory>\demo\intel64`.

If you receive error messages "This application has failed to start because `ippcore.dll` was not found" or "No DLL were found in the waterfall procedure", this means that the operating system is unable to determine the location of the Intel IPP dynamic libraries. To solve this issue:

- Ensure that the Intel IPP directory is in the path. Before using the Intel IPP dynamic libraries, add `<ipp directory>\redist\<arch>\ipp` to the `PATH` environment variable (see [Setting Environment Variables](#));
- Copy the contents of `<ipp directory>\redist\<arch>\ipp` to the `\system32` directory;
- Copy the contents of `<ipp directory>\redist\<arch>\ipp` to the application directory.

Note that you need to delete all Intel IPP DLLs from previous releases from the `C:\winnt\system` and `C:\winnt\system32` directories. Verify that paths to older library versions are not listed in the `PATH` environment variable.

Obtaining Version Information

To obtain information about the active library version including the version number, package ID, and the licensing information, call the `ippGetLibVersion` function. See the "Support Functions" chapter in the "Intel(R) IPP Reference Manual" (vol.1) for the function description and calling syntax.

You may also get the version information by running the `ippversion.h` file located in the `\include` directory.

Building Your Application

Follow the procedure described below to build the Windows OS application:

- set environment variables;
- include header files;
- call Intel IPP functions

Setting Environment Variables

The batch file `ippvars.bat` in the `\ipp\bin` directory sets the `LIB`, `INCLUDE`, and `PATH` environment variables for Intel IPP on the specified architecture. To do this run this file with the one of the following options in the command line:

```
ippvars.bat ia32|intel64.
```

Alternatively you can run the architecture specific batch file `ippvars_ia32.bat` in the `\ipp\bin\ia32` directory, or `ippvars_intel64.bat` in the `\ipp\bin\intel64` to set the Intel IPP `LIB`, `INCLUDE`, and `PATH` environment variables for the corresponding architecture.

To set the environment variables using the graphic user interface, see the operating system online help, or complete the following steps in the Windows XP* OS:

1. Right-click the **My Computer** icon on your desktop or from the Windows Explorer* and select **Properties**(or open **Control Pane** and select **System**),
2. Select the **Advanced** tab,
3. Select the **Environment Variables** button,
4. Use the interface to set the environment variables for only the current user (top dialog box) or for anyone who uses the system (bottom dialog box),
5. Select a variable and click the **Edit** button,
6. Add the path to the related Intel IPP files to the existing list. For example:
 - Select `LIB` and type in the directory for the Intel IPP libraries (default is: `<ipp_directory>\ipp\lib\ia32` and `<ipp_directory>\ipp\lib\intel64`),
 - Select `INCLUDE` and type in the directory for the Intel IPP header files (default is: `<ipp_directory>\ipp\include`),
 - Select `PATH` and type in the directory for the Intel IPP binaries (default is: `<ipp_directory>\redist\ia32\ipp` and `<ipp_directory>\redist\intel64\ipp`).
7. Click **OK** in the **Edit User Variable** dialog box,
8. Click **OK** in the **Environment Variables** dialog box,
9. Click **OK** in the **Systems Properties** dialog box.
10. Click **OK** in the **System Properties** dialog box.

For information on how to set up environment variables for threading, refer to [Supporting Multithreaded Applications](#).

Including Header Files

Intel IPP functions and types are defined in several header files that are organized by the function [domains and groups](#). They are located in the `\include` directory. For example, the `ippac.h` file contains declarations for all primitives for audio processing and audio coding.

The file `ipp.h` includes Intel IPP header files with the exception of cryptography and generated functions.

If you do not use cryptography and generated functions, use only `ipp.h` in your application for forward compatibility.

If you want to use cryptography and generated functions, you must directly include their header files `ippcp.h` and `ippgen.h` in your application.

Calling Intel IPP Functions

Due to the DLL dispatcher and merged static library mechanisms described in [Linking Your Application with Intel\(R\) IPP](#), calling Intel IPP functions is as simple as calling any other C function.

To call an Intel IPP function, do the following:

1. Include the `ipp.h` header file
2. Set up the function parameters
3. Call the function

The multiple versions of optimized code for each function are concealed under a single entry point. Refer to the "*Intel(R) IPP Reference Manual*" for detailed function descriptions, lists of required parameters, return values and so on.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Before You Begin Using Intel IPP

Before you start using Intel IPP, it is helpful to understand some basic concepts summarized in the table below:

What you need to know before you get started

Function Domain	Identify the Intel IPP functional domain that questions to answer belong to. Reason: If you know the functional domain you intend to use, it will narrow the search in the Reference Manuals for specific routines you need.
-----------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Besides, you may easily find a sample codes you would like to run from <http://www.intel.com/software/products/ipp/samples.htm>.

Refer to the sections "[Domains](#)" and "[Selecting Intel IPP Libraries Needed for your Application](#)" to understand what function domains are and what libraries are needed, and to the table "[Library Dependencies by Domain](#)" to understand what kind of cross-domain dependency is introduced.

Linking model

Decide what linking method is appropriate for linking.

Reason: If you choose a linking method that suits you, you will get the best linking results. For information on the benefits of each linking method, linking command syntax and examples, see [Linking Your Application with Intel\(R\) IPP](#)

Threading model

Select among the following options to determine how you are going to thread your application:

- Your application is already threaded.
- You may want to use the Intel(R) threading capability, that is, Compatibility OpenMP* run-time library (`libiomp`), or a threading capability provided by a third-party compiler.
- You do not want to thread your application.

Reason: By default, Intel IPP uses the OpenMP* software to set the number of threads that will be used. If you need a different number, you have to set it yourself using one of the available mechanisms. For more information, see [Supporting Multithreaded Applications](#).

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Intel(R) IPP Structure

This section discusses the structure of Intel(R) IPP after installation as well as the library types supplied.

Intel IPP Directory Structure

The table below shows the directory structure of Intel IPP for the given architecture after installation.

Directory	File types
<code><ipp directory></code>	Main directory (by default: C:\Program Files\Intel\composerxe-2011)
<code><ipp directory>\redist\<arch>\ipp</code>	Intel IPP dynamic link libraries (DLLs)
<code><ipp directory>\Documentation\en_US\ipp</code>	Intel IPP documentation files
<code><ipp directory>\Samples\en_US\IPP</code>	Intel IPP application-level samples (see Using Intel IPP Samples)
<code><ipp directory>\ipp\demo\<arch></code>	Executable programs that demonstrate various image and signal processing functionalities
<code><ipp directory>\ipp\include</code>	Intel IPP header files
<code><ipp directory>\ipp\lib\<arch></code>	Intel IPP static libraries, and static import libraries (<i>stub</i> static libraries)
<code><ipp directory>\ipp\interfaces</code>	High-level interfaces for data-compression
<code><ipp directory>\ipp\bin\<arch></code>	batch files to set environment variables
<code><ipp directory>\ipp\tools\<arch>\perfsys</code>	Performance Test tool
<code><ipp directory>\ipp\tools\<arch>\staticlib</code>	tools to create single-processor executable

Here `<arch>` refer to the architecture specific folder.

Supplied Libraries

The following table lists the types of libraries in Intel IPP and shows examples of the library files supplied.

Types of Libraries of Intel IPP

Library types	Description	Folder location	Example
Dynamic	Dynamic link libraries (DLL) include both processor dispatchers and function implementations	\redist\ <arch>\ipp</arch>	ipp _s -7.0.dll ipp _s p8-7.0.dll
Static (import)	"Stub" static library files. They load the required DLLs and link to the correct entry points	\ipp\lib\ <arch>< td=""> <td>ipp_s.lib</td> </arch><>	ipp _s .lib
Static merged	Contain function implementations for all supported processor types	\ipp\lib\ <arch>< td=""> <td>ipp_s_l.lib</td> </arch><>	ipp _s _l.lib
Threaded static merged	Contain threaded function implementations	\ipp\lib\ <arch>< td=""> <td>ipp_s_t.lib</td> </arch><>	ipp _s _t.lib
Static emerged	Contain dispatchers for the merged libraries	\ipp\lib\ <arch>< td=""> <td>ipp_s_l.lib</td> </arch><>	ipp _s _l.lib

Here <arch> refer to the architecture specific folder.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Using Intel IPP Dynamic Link Libraries (DLLs)

Intel IPP includes the dynamic link libraries (DLLs) in the \redist\\ipp directory.

To load the Intel IPP DLLs and link to the correct entry points, use static import library ("stub" library) files in the \ipp\lib\Intel IPP directory structure).

To dynamically link to the DLL for the appropriate processor, link to the ipp*.lib (* denotes the appropriate function domain) files. You must set your LIB environment variable using the required batch file (see [Setting Environment Variables](#)), or refer to these files using their full path.

The DLLs `ipp*-x.x.dll` (* denotes the appropriate function domain, `x.x` - denotes Intel IPP version number) are "dispatcher" dynamic libraries. At run time, they detect the processor and load the correct processor-specific DLLs. This enables you to write the code to call the Intel IPP functions without worrying about which processor the code will execute on - the appropriate processor-specific library is automatically used. These processor-specific libraries contain [processor specific codes](#) in their names. For example, `ippiv8-7.0.dll` in the `\redist\ia32\ipp` directory reflects the imaging processing libraries optimized for the Intel® Core™ 2 Duo processors.

See also [Selecting Intel® IPP Libraries Needed by Your Application](#).



NOTE. You must include the appropriate `libiomp5md.dll` in your `PATH` environment variable. You can find this file in the `\compiler\lib\<arch>` directory.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Using Intel IPP Static Libraries

The Intel IPP includes the *merged* static library files that contain optimized code for different processor types of each function. These files reside in the `\ipp\lib\<arch>` directory (see [Intel IPP directory structure](#)).

Just as with the dynamic dispatcher, the appropriate version of a function is executed when the function is called. This mechanism is not as convenient as the dynamic mechanism, but it can result in a smaller total code size in spite of the big size of the static libraries.

To use these static libraries, link to the appropriate files `ipp*_l.lib` or `ipp*_t.lib` in the `\ipp\lib\<arch>` directory (* refers to the functional [domain](#)). You need to set your `LIB` environment variable (see [Setting Environment Variables](#)), or refer to these files using their full path.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

See also [Selecting Intel® IPP Libraries Needed by Your Application](#).

Contents of the Documentation Directory

The <ipp_directory>\Documentation\en_US\ipp directory includes all the documentation related to Intel(R) IPP. See the `ipp_documentation.htm` file for a listing of all the available document with links or pointers to their location.

Configuring Your Development Environment

4

This section explains how to configure your development environment for use with Intel(R) IPP.

Configuring Your Microsoft* Development Environment to Link with Intel IPP

Follow these steps to configure your Microsoft* development environment, that were tested on Microsoft Visual C++* 2005, Microsoft Visual C++* 2008, or Microsoft Visual C++* 2010 . If you have another version of Microsoft Visual Studio*, the options may differ slightly.

1. Select **View > Solution Explorer** (make sure this window is active).
2. For Microsoft Visual C++* 2005 or Microsoft Visual C++* 2008:
Select **Projects > Intel C++ Composer 2011 > Select Build Components** and check **Use IPP** checkbox (Intel Integrated Performance Primitives section).
3. For Microsoft Visual C++* 2010:
Select **Projects > Properties > Configuration Properties > Intel Performance Libraries** and set **Yes** to **Use IPP**.
4. Type of the Intel IPP libraries to link to depends on the **Runtime Library** option in C/C++ configurations. Setting multi-threaded DLL (/MD) or multi-threaded debug DLL (/MDd) options you add `_IPP_PARALLEL_DYNAMIC` pre-processor definition and dynamic libraries. Setting multi-threaded (/MT) and multi-threaded debug (/MTd) options you add `_IPP_SEQUENTIAL_STATIC` pre-processor definition and static libraries. For more information on choosing the best linkage model for your Intel IPP application, please refer to [Linking Your Application with Intel® IPP](#).

On processors with the Intel® 64 architecture, currently many of the 64-bit operating systems like Microsoft Windows XP* Professional x64 Edition and Microsoft Windows Server* 2003 family x64 editions can support running 32-bit or 64-bit applications. If you want to build an application for the Intel® 64 architecture-based machine in the 32-bit mode, you must install Intel IPP for IA-32 architecture and link to 32-bit libraries.

If you use the Intel IPP for Intel 64 architecture, make sure the active solution platform is "x64" in **Menu > Build > Configuration Manager**.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this

Optimization Notice

product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Creating Visual C++ 2005 Project Files for the Intel IPP Samples

To create Microsoft* Visual* C++ 2005 project files for the Intel IPP samples, follow the steps below, which use jpegview sample code as an example.

1. Download the media codec sample jpegview from <http://www.intel.com/software/products/ipp/samples.htm>.
2. Select **File > New > Project from Existing Code:**.
3. Select **Visual C++** as the type of project from the drop-down menu in the pop-up window. Fill in the project file location and project name. Check the **Add files to the project from these folders** check box. Click **Add**, select the jpegview folder, and click **OK**.
4. Select **View > Solution Explorer** (make sure this window is active).
5. Select **Projects > Properties**.
6. Under **Configuration Properties > C/C++** in the drop down menu titled **Show directories for:**, select **Include Files** and type in the directory for the Intel IPP include files (for example, the default is: `<ipp_directory>\ipp\include`).
7. In the **Show directories for:** drop-down menu, select **Library Files** and type in the directory for the Intel IPP library files. By default it is: `<ipp_directory>\ipp\lib\<arch>`.
8. In the **Show directories for:** drop-down menu, select **Executable Files** and type in the directory for the Intel IPP executable files. By default it is: `<ipp_directory> \redist\<ipp_directory>\ipp`.

Using the IntelliSense* Features

Intel IPP supports two Microsoft Visual Studio intellisense features that support language references: [Complete Word](#) and [Parameter Info](#).



NOTE. Both features require header files. Therefore, to benefit from IntelliSense, specify the path to the include files in the Visual Studio or solution settings. See above sections on how to do this.

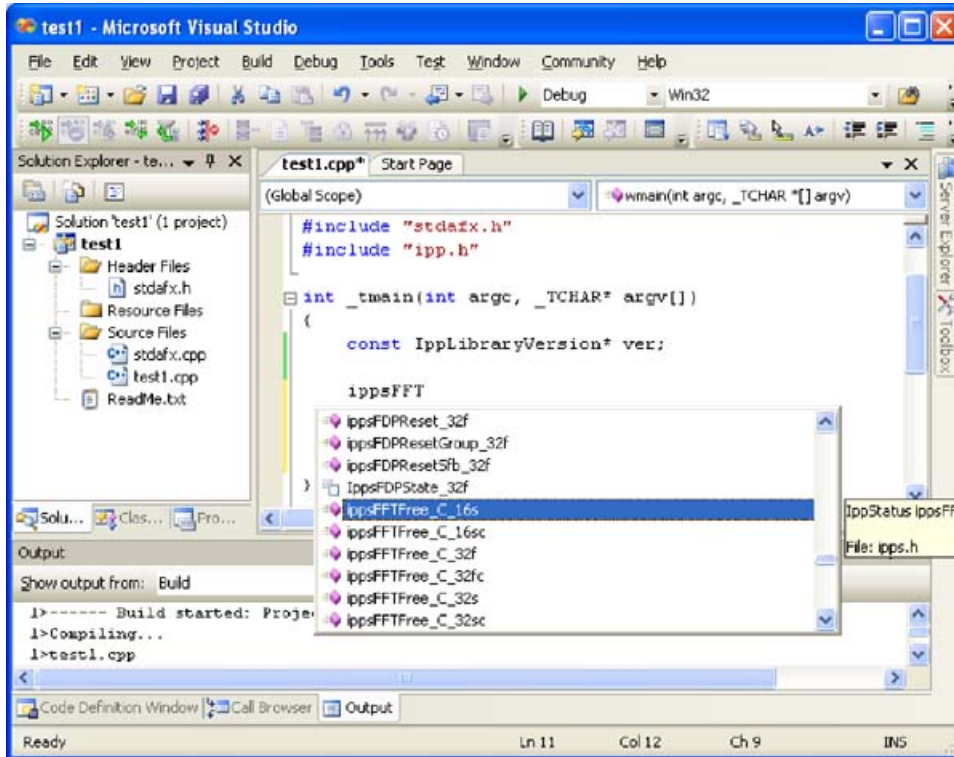
Complete Word

For a software library, the *Complete Word* feature types or prompts for the rest of the name defined in the header file once you type the first few characters of the name in your code.

Provided your C/C++ code contains the include statement with the appropriate Intel IPP header file, to complete the name of the function or named constant specified in the header file, follow these steps:

1. Type the first few characters of the name (for example, `ipp_sFFT`).

2. Press **Alt + RIGHT ARROW** or **Ctrl + SPACEBAR** If you have typed enough characters to eliminate ambiguity in the name, the rest of the name is typed automatically. Otherwise, the pop-up list of the names specified in the header file opens - see the figure below.



3. Select the name from the list, if needed.

Parameter Info

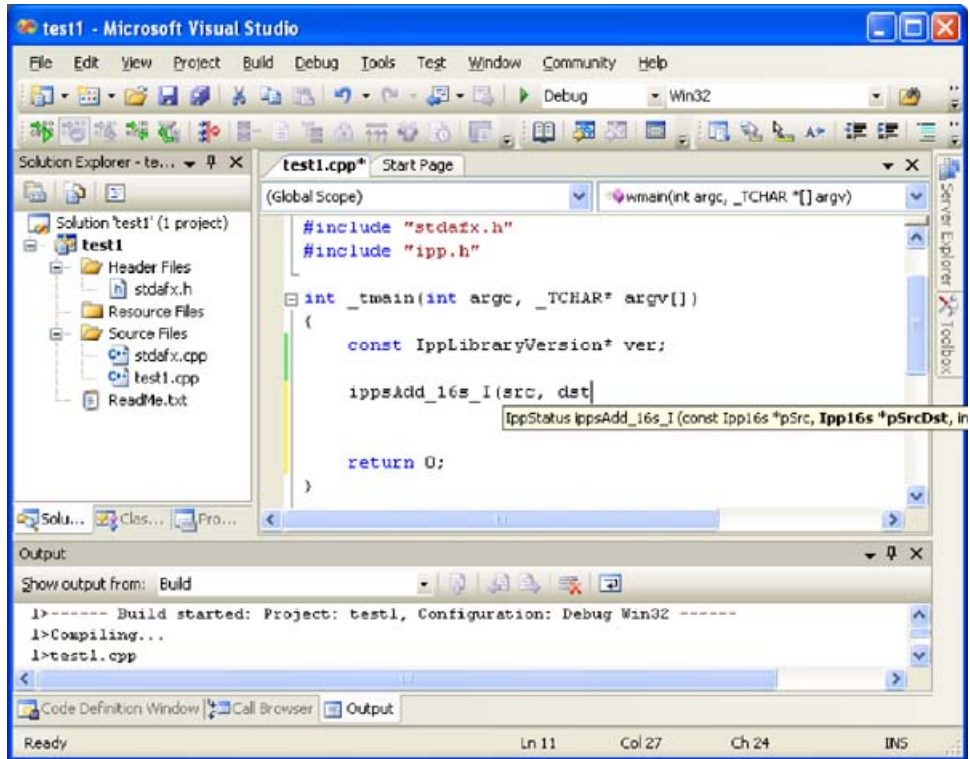
The *Parameter Info* feature displays the parameter list for a function to give information on the number and types of parameters.

To get the list of parameters of a function specified in the header file, follow these steps:

1. Type the function name
2. Type the opening parenthesis

4 *Intel® Integrated Performance Primitives for Windows* OS User's Guide*

A tooltip appears with the function API prototype, and the current parameter in the API prototype is highlighted



- see the figure below.

Linking Your Application with Intel(R) IPP

5

This section discusses dispatching of the Intel(R) IPP libraries to specific processors using various models for linking Intel(R) IPP to an application, considers differences between the linking methods regarding development and target environments, installation specifications, run-time conditions, and other application requirements to help you select the best linking method for your application, shows linking procedure for each linking method, and provides examples.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Dispatching

Intel IPP uses codes optimized for various central processing units (CPUs). Dispatching refers to detection of your CPU and selecting the corresponding Intel IPP binary. For example, the `ippiv8-7.0.dll` file in the `\redist\ia32\ipp` directory, reflects the imaging processing libraries optimized for the Intel(R) Core™ 2 Duo processors.

A single Intel IPP function, for example `ippsCopy_8u()`, may have many versions, each one optimized to run on a specific Intel(R) processor with specific architecture, for example, the version of this function optimized for the Intel(R) Core™ 2 Duo processor is `v8_ippsCopy_8u()`, and version optimized for 64-bit applications on processors with Intel(R) SSE4.1 is `u8_ippsCopy_8u()`.

The following table Table 5-1 shows processor-specific codes used in Intel IPP:

Table 5-1 Identification of Codes Associated with Processor-Specific Libraries

Abbreviation	Meaning
	IA-32 Intel® architecture
px	C-optimized for all IA-32 processors
v8	Optimized for processors with Intel® Supplemental Streaming SIMD Extensions 3

Abbreviation	Meaning
p8	Optimized for processors with Intel® Streaming SIMD Extensions 4.1
g9	Optimized for processors that support Intel® Advanced Vector Extensions instruction set
s8	Optimized for the Intel® Atom™ processor Intel® 64 architecture
mx	C-optimized for processors with Intel® 64 instructions set architecture
u8	Optimized for 64-bit applications on processors with Intel® Supplemental Streaming SIMD Extensions 3
y8	Optimized for 64-bit applications on processors with Intel® Streaming SIMD Extensions 4.1
n8	Optimized for the Intel® Atom™ processor
e9	Optimized for processors that support Intel® Advanced Vector Extensions instruction set

Supporting Intel(R) Xeon(R) Processors with the Intel(R) 64 Architecture

Intel IPP supports Intel(R) Xeon(R) processors with the Intel(R) 64 architecture running 32-bit and 64-bit modes. For the 32-bit mode, download and install Intel IPP for the Pentium(R) processor-based systems. When running the 64-bit mode, install Intel IPP for the Intel(R) 64 architecture.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Detecting Processor Features and Type

To obtain information about the features of the processor used in your computer system, use the function `ippGetCpuFeatures`. To obtain information about the features of the processor that are enabled for the loaded libraries, use the function `ippGetEnabledCpuFeatures`. Both functions are declared in the `ippcore.h` file. These function retrieve main processor features returned by the function `CPUID.1` and store them consecutively in the mask that is returned by the function. The table below lists all CPU features that can be retrieved.

Table 5-3 Processor Features Mask

Mask Value	Bit Name	Supported Feature
1	<code>ippCPUID_MMX</code>	MMX™ technology
2	<code>ippCPUID_SSE</code>	Intel® Streaming SIMD Extensions
4	<code>ippCPUID_SSE2</code>	Intel® Streaming SIMD Extensions 2
8	<code>ippCPUID_SSE3</code>	Intel® Streaming SIMD Extensions 3
16	<code>ippCPUID_SSSE3</code>	Intel® Supplemental Streaming SIMD Extensions 3
32	<code>ippCPUID_MOVBE</code>	MOVBE instruction
64	<code>ippCPUID_SSE41</code>	Intel® Streaming SIMD Extensions 4.1
128	<code>ippCPUID_SSE42</code>	Intel® Streaming SIMD Extensions 4.2
256	<code>ippCPUID_AVX</code>	Intel® Advanced Vector Extensions (Intel® AVX) instruction set
512	<code>ippAVX_ENABLEDBYOS</code>	The operating system supports Intel® AVX
1024	<code>ippCPUID_AES</code>	Intel® Advanced Encryption Standard (AES) instructions set
2048	<code>ippCPUID_CLMUL</code>	PCLMULQDQ instruction

To detect the processor type used in your computer system, use the function `ippGetCpuType`, which is declared in the `ippcore.h` file. It returns an appropriate `IppCpuType` variable value. All of the enumerated values are given in the `ippdefs.h` header file. For example, the return value `ippCpuPII` means that your system uses Intel(R) Core™ 2 Duo processor.

The following table Table 5-4 shows possible return values of `ippGetCpuType` and their meaning.

Table 5-4 Detecting Processor Type. Return Values and Their Meaning

Return Value	Processor Type
<code>ippCpuPP</code>	Intel® Pentium® processor
<code>ippCpuPMX</code>	Pentium® processor with MMX™ technology
<code>ippCpuPPR</code>	Pentium® Pro processor
<code>ippCpuPII</code>	Pentium® II processor
<code>ippCpuPIII</code>	Pentium® III processor and Pentium® III Xeon® processor
<code>ippCpuP4</code>	Pentium® 4 processor and Intel® Xeon® processor
<code>ippCpuP4HT</code>	Pentium® 4 processor with Hyper-Threading Technology
<code>ippCpuP4HT2</code>	Pentium® Processor with Intel® Streaming SIMD Extensions 3

Return Value	Processor Type
ippCpuCentrino	Intel® Centrino® processor technology
ippCpuCoreSolo	Intel® Core™ Solo processor
ippCpuCoreDuo	Intel® Core™ Duo processor
ippCpuITP	Intel® Itanium® processor
ippCpuITP2	Intel® Itanium® 2 processor
ippCpuEM64T	Intel® 64 Instruction Set Architecture (ISA)
ippCpuC2D	Intel® Core™ 2 Duo processor
ippCpuC2Q	Intel® Core™ 2 Quad processor
ippCpuPenryn	Intel® Core™ 2 processor with Intel® Streaming SIMD Extensions 4.1 instruction set
ippCpuBonnell	Intel® Atom™ processor
ippCpuNehalem	Intel® Core™ i7 processor
ippCpuSSE	Processor with Intel® Streaming SIMD Extensions instruction set
ippCpuSSE2	Processor with Intel® Streaming SIMD Extensions 2 instruction set
ippCpuSSE3	Processor with Intel® Streaming SIMD Extensions 3 instruction set
ippCpuSSSE3	Processor with Intel® Supplemental Streaming SIMD Extensions 3 instruction set
ippCpuSSE41	Processor with Intel® Streaming SIMD Extensions 4.1 instruction set
ippCpuSSE42	Processor with Intel® Streaming SIMD Extensions 4.2 instruction set
ippCpuAVX	Processor supports Intel® Advanced Vector Extensions instruction set
ippCpuAES	Processor supports Intel® Advanced Encryption Standard (AES) instructions set
ippCpuX8664	Processor supports 64 bit extension
ippCpuUnknown	Unknown Processor

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this

Optimization Notice

product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Selecting Between Linking Methods

You can use different linking methods for Intel IPP:

- Dynamic linking using the run-time dynamic link libraries (DLLs)
- Static linking with dispatching by using emerged and merged static libraries
- Static linking without automatic dispatching using merged static libraries
- Dynamically building your own - custom - DLL.

Answering the following questions helps you select the linking method which best suites you:

- Are there limitations on the size of the application executable? Are there limitations on the size of the application installation package?
- Is the Intel IPP-based application a device driver or similar "ring 0" software that executes in the kernel mode at least some of the time?
- Will the application be installed on a range of processor types, or is the application explicitly supported only on a single type of processor? Is the application part of an embedded computer with only one type of processor?
- What resources are available for maintaining and updating customized Intel IPP components? What level of effort is acceptable for incorporating new processor optimizations into the application?
- How often will the application be updated? Will application components be distributed independently or will they always be packaged together?

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Dynamic Linking

Dynamic linking is the simplest method and the most commonly used. It takes full advantage of the dynamic dispatching mechanism in the dynamic link libraries (DLLs) (see also [Intel® IPP Structure](#)). The following table summarizes the features of dynamic linking to help you understand the benefits and drawbacks of this linking method.

Table 5-5 Summary of Dynamic Linking Features

Benefits	Drawbacks
<ul style="list-style-type: none"> Automatic run-time dispatch of processor-specific optimizations Enables updates with new processor optimizations without recompile/relink Reduces disk space requirements for applications with multiple Intel IPP-based executables Enables more efficient shared use of memory at run-time for multiple Intel IPP-based applications 	<ul style="list-style-type: none"> The application executable requires access to Intel IPP run-time dynamic link libraries (DLLs) Not appropriate for kernel-mode/device-driver/ring-0 code Not appropriate for web applets/plugin-ins that require very small download A one-time performance penalty when the Intel IPP DLLs are first loaded

To dynamically link with Intel IPP, follow these steps:

1. Include `ipp.h` in your application. This header includes the header files for all Intel IPP functional domains.
2. Call IPP functions using normal, undecorated function names.
3. Link corresponding domain import libraries. For example, if you use the function `ippsCopy_8u`, link to `ipps.lib`.
4. Make sure that the run-time libraries, for example `ipps.dll`, are on the executable search path at run time. Run the `ippvars.bat` from directory `\ipp\bin` to ensure that the application loads the appropriate processor-specific library.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Static Linking with Dispatching

Some applications use only a few Intel IPP functions and require a small memory footprint. Using the static link libraries offers both the benefits of a small footprint and optimization on multiple processors. The static libraries (such as `ipps_l.lib`) provide an entry point for the non-decorated (with normal names) Intel IPP functions, and the jump table to each processor-specific implementation. When linked with your application, the function calls corresponding functions in accordance with the CPU setting detected by functions in `ippcore_l.lib`.

If you want to use the threaded functions of Intel IPP, you need to link to the threaded versions of the static libraries (such as `ipps_t.lib`), and threaded version `ippcore_t.lib`.

You can choose one of the following functions to initialize the libraries:

- `ippStaticInit()` to use the best available optimization, or
- `ippStaticInitCpu()` that lets you specify directly the CPU.

One of these functions must be called before any other IPP functions. Otherwise, a C-optimized version of the IPP functions is called. This can decrease the performance of your application. The following example illustrates the performance difference. This example appears in the `t2.cpp` file.

Example: 5-1 Performance difference with and without calling `StaticInit`

```
#include <stdio.h>
#include <ipp.h>

int main() {
    const int N = 20000, loops = 100;
    Ipp32f src[N], dst[N];
    unsigned int seed = 12345678, i;
    Ipp64s t1,t2;
    // no StaticInit call, means PX code, not optimized
    ippsRandUniform_Direct_32f(src,N,0.0,1.0,&seed);
    t1=ippGetCpuClocks();
    for(i=0; i<loops; i++)
        ippsSqrt_32f(src,dst,N);
    t2=ippGetCpuClocks();
    printf("without StaticInit: %.1f clocks/element\n",
        (float)(t2-t1)/loops/N);
    ippStaticInit();
    t1=ippGetCpuClocks();
    for(i=0; i<loops; i++)
        ippsSqrt_32f(src,dst,N);
    t2=ippGetCpuClocks();
    printf("with StaticInit: %.1f clocks/element\n",
        (float)(t2-t1)/loops/N);
    return 0;
}
```

```
t2.cpp
cmdlinetest>t2
without StaticInit: 61.3 clocks/element
with StaticInit: 4.5 clocks/element
```

The following table summarizes the features of static linking to help you understand the benefits and drawbacks of this linking method.

Table 5-6 Summary of features of the static linking (with dispatching)

Benefits	Drawbacks
<ul style="list-style-type: none"> Dispatches processor-specific optimizations during run-time Creates a self-contained application executable Generates a smaller footprint than the full set of Intel IPP DLLs 	<ul style="list-style-type: none"> Intel IPP code is duplicated for multiple Intel IPP-based applications because of static linking An additional function call for dispatcher initialization is needed (once) during program initialization

Follow these steps to use static linking with dispatching:

1. Include `ipp.h` in your application. This header includes the header files of all IPP domains.
2. Initialize the static dispatcher using either function `ippStaticInit()` or `ippInitCPU()`, which are declared in the header file `ippcore.h`.

3. Call IPP functions using normal, undecorated function names.
4. Link corresponding static libraries, and then `ippcore_l.lib` or `ippcore_t.lib`. For example, if you use the function `ippsCopy_8u()`, the linked libraries are `ipps.lib` and `ippcore_l.lib`.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Static Linking without Dispatching

Static linking links directly with the merged static libraries. Use this method to link a self contained application that is supported on a specific processor type. Static linking is useful for embedded applications that are bundled with one type of processor.

The following table summarizes the features of static linking to help you understand the benefits and drawbacks of this linking method.

Table 5-7 Summary of Features of the Static Linking (without dispatching)

Benefits	Drawbacks
<ul style="list-style-type: none"> • Small executable size with support for only one processor type • Suitable for kernel-mode/device-driver/ring-0 use *) • Suitable for a Web applet or a plug-in requiring very small file download and support for only one processor type • Self-contained application executable that does not require the Intel IPP run-time DLLs • Smallest footprint for application package • Smallest installation package 	<ul style="list-style-type: none"> • The executable is optimized for only one processor type • Updates to processor-specific optimizations require rebuild and/or relink

*) for unthreaded libraries only

You may want to use your own static dispatcher instead of the provided emerged dispatcher. The IPP sample `mergelib` demonstrates how to do this.

Refer to the latest updated sample from the Intel IPP samples directory:
`\ipp-samples\advanced-usage\linkage\mergelib` at
<http://www.intel.com/software/products/ipp/samples.htm>.

The Intel IPP package includes a set of processor-specific header files (such as `ipp_v8.h`). You can use these header files instead of the `IPPCALL` macro. Refer to *Static linking to Intel(R) IPP Functions for One Processor* in `\ipp\tools<arch>\static.lib\readme.htm`.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Building a Custom DLL

A custom dynamic link library (DLL) is useful for an application that has few internal modules, and only these modules share Intel IPP code. In this case, you can use dynamic linking with a customized DLL containing only those Intel IPP functions that the application uses.

The following table summarizes features of the custom DLL.

Custom DLL Features

Benefits	Drawbacks
<ul style="list-style-type: none"> Run-time dispatching of processor-specific optimizations Reduced hard-drive footprint compared with a full set of Intel IPP DLLs Smallest installation package to accommodate use of some of the same Intel IPP functions by multiple applications 	<ul style="list-style-type: none"> Application executable requires access to the Intel compiler specific run-time libraries that are delivered with Intel IPP Developer resources are needed to create and maintain the custom DLLs Integration of new processor-specific optimizations requires rebuilding the custom DLLs Not appropriate for kernel-mode/device-driver/ring-0 code

To create a custom DLL, you need to create a separate build step or project that generates the DLL and stubs. The specially developed sample demonstrates how to do it. Refer to the latest updated `custom_dll` sample from the Intel IPP samples directory: `\ipp-samples\advanced-usage\linkage\customdll` at <http://www.intel.com/software/products/ipp/samples.htm>.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this

Optimization Notice

product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Comparison of Intel IPP Linking Methods

The following table gives a quick comparison of the Intel IPP linking methods.

Table 5-9 Intel IPP Linking Method Summary Comparison

Feature	Dynamic Linking	Static Linking with Dispatching	Static Linking without Dispatching	Using Custom DLL
Processor Updates	Automatic	Recompile & redistribute	Release new processor-specific application	Recompile & redistribute
Optimization	All processors	All processors	One processor	All processors
Build	Link to stub static libraries	Link to static libraries and static dispatchers	Link to merged libraries	Build separate DLL
Calling	Regular names	Regular names	Processor-specific names	Regular names
Total Binary Size	Large	Small	Smallest	Small
Executable Size	Smallest	Small	Small	Smallest
Kernel Mode	No	Yes	Yes	No

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Selecting the Intel(R) IPP Libraries Needed by Your Application

The following table shows functional domains and functions groups, and the relevant header files and libraries used for each linking method.

Table 5-10 Libraries Used for Each Linking Method

Description	Header Files	Dynamic Linking	Static Linking (with/without Dispatching) and Custom Dynamic Linking
Audio Coding	ippac.h	ippac.lib	ippac_l.lib ippac_t.lib
Color Conversion	ippcc.h	ippcc.lib	ippcc_l.lib ippcc_t.lib
String Operation	ippch.h	ippch.lib	ippch_l.lib ippch_t.lib
Cryptography	ippcp.h	ippcp.lib	ippcp_l.lib ippcp_t.lib
Computer Vision	ippcv.h	ippcv.lib	ippcv_l.lib ippcv_t.lib
Data Compression	ippdc.h	ippdc.lib	ippdc_l.lib ippdc_t.lib
Data Integrity	ippdi.h	ippdi.lib	ippdi_l.lib ippdi_t.lib
Image Processing	ippi.h	ippi.lib	ippi_l.lib ippi_t.lib
Image Compression	ippj.h	ippj.lib	ippj_l.lib ippj_t.lib
Realistic Rendering and 3D Data Processing	ippr.h	ippr.lib	ippr_l.lib ippr_t.lib
Small Matrix Operations	ippm.h	ippm.lib	ippm_l.lib ippm_t.lib
Signal Processing	ipps.h	ipps.lib	ipps_l.lib ipps_t.lib
Speech Coding	ippsc.h	ippsc.lib	ippsc_l.lib ippsc_t.lib
Video Coding	ippvc.h	ippvc.lib	ippvc_l.lib

Description	Header Files	Dynamic Linking	Static Linking (with/without Dispatching) and Custom Dynamic Linking
			ippvc_t.lib
Vector Math	ippvm.h	ippvm.lib	ippvm_l.lib ippvm_t.lib
Generated Functions	ippgen.h	ippgen.lib	ippgen_l.lib ippgen_t.lib
Core Functions	ippcore.h	ippcore.lib	ippcore_l.lib ippcore_t.lib

Libraries for Dynamic Linking

To use the dynamic linking libraries, you must link to the domain libraries `ipp*.lib` used in your application (here * denotes the appropriate function domain). You can find them in the `\ipp\lib\<arch>` directory. Additionally you must link to the libraries `ipps.lib`, `ippcore.lib`, and `libiomp5md.lib`.

For example, your application uses three Intel IPP functions `ippiCopy_8u_C1R`, `ippiCanny_16s8u_C1R`, and `ippmMul_mc_32f`. These three functions belong to the image processing, computer vision, and small matrix operations domains, respectively. To include these functions into your application, you must link to the following libraries:

```

ippi.lib
ippcv.lib
ippm.lib
ipps.lib
ippcore.lib
libiomp5md.lib

```

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Libraries for Static Linking

To use the static linking libraries, you need to link to all required domain libraries `ipp*_l.lib`, core library `ippcore_l.lib`, and libraries on which domain libraries depend (see next section). The * denotes the appropriate function domain.

If you want to use the Intel IPP functions threaded with the OpenMP*, you need to link to the threaded versions of the libraries `ipp*_t.lib`, `ippcore_t.lib`, and `libiomp5md.lib`.

All domain-specific and core libraries are located in the `\ipp\lib<arch>` directory.

For example, your application uses three Intel IPP functions `ippiCopy_8u_C1R`, `ippiCanny_16s8u_C1R`, and `ippmMul_mc_32f`. These three functions belong to the image processing, computer vision, and small matrix operations domains respectively. All these domain libraries depend on signal processing library.



NOTE. The order in which libraries are linked must correspond to the library dependencies by domain (see the next section).

To include these functions into your application, link to the following libraries:

```
ippcv_l.lib
ippm_l.lib
ippi_l.lib
ipps_l.lib
ippcore_l.lib
```

or if you want to use the threaded functions:

```
ippcv_t.lib
ippm_t.lib
ippi_t.lib
ipps_t.lib
ippcore_t.lib
libiomp5md.lib
```

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Library Dependencies by Domain

The following table lists library dependencies by domain for static linking.

Table 5-11 Library Dependencies by Domain

Domain	Library	Dependent on
Audio Coding	ippac	ippdc, ipp, ippcore
Color Conversion	ippcc	ippi, ipp, ippcore
Cryptography	ippcp	ippcore
Computer Vision	ippcv	ippi, ipp, ippcore
Data Compression	ippdc	ipp, ippcore
Data Integrity	ippdi	ippcore
Image Processing	ippi	ipp, ippcore
Image Compression	ippj	ippi, ipp, ippcore
Small Matrix Operations	ippm	ippi, ipp, ippcore
Realistic Rendering and 3D Data Processing	ipprr	ippi, ipp, ippcore
Signal Processing	ipp	ippcore
Speech Coding	ippsc	ipp, ippcore
String Operations	ippch	ipp, ippcore
Video Coding	ippvc	ippi, ipp, ippcore
Vector Math	ippvm	ippcore

Refer to the *Intel IPP Reference Manual* to find which domain your function belongs to.

Generated functions (ippgen library) depends on ipp and ippcore libraries.

When you use static linking to a certain library (for example, data compression domain ippdc), you must link to the libraries on which it depends (in our example, the signal processing ipp and core functions ippcore).

Linking Examples

For more linking examples, please go to <http://www.intel.com/software/products/ipp/samples.htm>

For information on using sample code, see [Intel\(R\) IPP Samples](#).

Supporting Multithreaded Applications

6

This section discusses the use of Intel(R) IPP in multithreading applications and helps you:

[set number of threads in multithreaded applications](#)

[get information on number of threads](#)

[use shared L2 cache](#)

[avoid nested parallelization](#)

[disable multithreading.](#)

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Intel IPP Threading and OpenMP* Support

All Intel IPP functions are thread-safe. They supports multithreading in both dynamic and static libraries and can be used in multithreaded applications.

Some Intel IPP functions contain OpenMP* code, which increase performance on multi-processor and multi-core systems. These functions include color conversion, filtering, convolution, cryptography, cross correlation, matrix computation, square distance, bit reduction, and so on.

Refer to the *ThreadedFunctionsList.txt* document to see the list of all threaded APIs in the documentation directory of the Intel IPP installation.

See also <http://www.intel.com/software/products/support/ipp> for more topics related to Intel IPP threading and OpenMP* support, including older Intel IPP versions of threaded API.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this

Optimization Notice

product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Setting Number of Threads

By default, the number of threads for Intel IPP threaded libraries equals the number of processors in the system. If the value of the `OMP_NUM_THREADS` environment variable is less than the number of processors in the system, then the number of threads for Intel IPP threaded libraries equals the value of the `OMP_NUM_THREADS` environment variable.

To configure the number of threads used by Intel IPP internally, call the function `ippSetNumThreads(n)` at the very beginning of an application. `n` is the desired number of threads (1, ...). To disable internal parallelization, call `ippSetNumThreads(1)`.

Getting Information on Number of Threads

To find the number of threads created by the Intel IPP, call function `ippGetNumThreads`.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Using a Shared L2 Cache

Some functions in the signal processing domain are threaded on two threads intended for the Intel(R) Core™ 2 processor family, and make use of the merged L2 cache. These functions (single and double precision `FFT`, `Div`, `Sqrt` and so on) achieve the maximum performance if both two threads are executed on the same die. In this case, these threads work on the same shared L2 cache. For processors with two cores on the die, this condition is satisfied automatically. For processors with more than two cores, set the following OpenMP environmental variable to avoid performance degradation:

```
KMP_AFFINITY=compact
```

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this

Optimization Notice

product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Avoiding Nested Parallelization

Nested parallelization may occur if you use a threaded Intel IPP function on a multithreaded application. Nested parallelization may cause performance degradation.

For applications that use OpenMP threading, nested threading is disabled by default, so this is not an issue.

However, if your application uses threading created by a tool other than OpenMP, you must disable multithreading in the threaded Intel IPP function to avoid this issue.

Disabling Multithreading

To disable multithreading, call function `ippSetNumThreads` with parameter 1, or link your application with IPP unthreaded static libraries.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Managing Performance and Memory

7

This section describes several methods for managing performance and memory to optimize your Intel(R) IPP application: aligning memory, thresholding denormal data, reusing buffers, using FFT for algorithmic optimization. Finally, it explains how to test the performance of the Intel IPP functions, using the Intel(R) IPP Performance Test Tool and it gives some examples of using the Performance Tool Command Lines.

Memory Alignment

Access to memory is faster if pointers to the data are aligned, and Intel IPP functions perform better if they process data with aligned pointers.

The following Intel IPP functions can be used for pointer alignment, memory allocation and deallocation:

```
void* ippAlignPtr(void* ptr, int alignBytes)
```

Aligns a pointer, can align to the powers of 2, that is 2, 4, 8, 16 and so on.

```
void* ippMalloc(int length)
```

32-byte aligned memory allocation. Memory can be freed only with the function `ippFree`.

```
void ippFree(void* ptr)
```

frees memory allocated by the function `ippMalloc`.

```
Ipp<datatype>* ippMalloc_<datatype>(int len)
```

32-byte aligned memory allocation for signal elements of different data types. Memory can be freed only with the function `ippFree`.

```
void ippFree(void* ptr)
```

Frees memory allocated by the function `ippMalloc`.

```
Ipp<datatype>* ippiMalloc_<mod>(int widthPixels, int heightPixels, int* pStepBytes) -
```

32-byte aligned memory allocation for images where every line of the image is padded with zeros. Memory can be freed only with the function `ippiFree`.

```
void ippiFree(void* ptr) -
```

Frees memory allocated by the function `ippiMalloc`.

The amount of memory that can be allocated is determined by the operating system and system hardware, but it cannot exceed 2GB.



- Intel IPP functions `ippFree`, `ippFree`, and `ippiFree` can only be used to free memory allocated by the functions `ippMalloc`, `ippMalloc`, and `ippiMalloc` respectively.

- Intel IPP functions `ippFree`, `ippsFree`, and `ippiFree` cannot be used to free memory allocated by standard functions like `malloc` or `calloc`. The memory allocated by the Intel IPP functions `ippMalloc`, `ippsMalloc`, and `ippiMalloc` cannot be freed by the standard function `free`.

The following code example shows how to call the function `ippiMalloc`.

Example 7-1. Calling the Function `ippiMalloc`

```
#include "stdafx.h"
#include "ipp.h"
#include "tools.h"

int main(int argc, char *argv[])
{
    IppiSize size = {320, 240};
    int stride;
    Ipp8u* pSrc = ippiMalloc_8u_C3(size.width, size.height, &stride);
    ippiImageJaehne_8u_C3R(pSrc, stride, size);

    int dstStride;
    Ipp8u* pDst = ippiMalloc_8u_C3(size.width, size.height, &dstStride);
    ippiCopy_8u_C3R(pSrc, stride, pDst, dstStride, size);

    IppiSize ROIsize = { size.width/2, size.height/2 };
    ippiCopy_8u_C3R(pSrc, stride, pDst, dstStride, ROIsize);

    IppiPoint srcOffset = { size.width/4, size.height/4 };
    ippiCopy_8u_C3R(pSrc + srcOffset.x*3 + srcOffset.y*stride, stride,
    pDst, dstStride, ROIsize);

    return 0;
}
```

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Thresholding Data

Denormal numbers are the border values in the floating-point format and special case values for the processor. Denormal data occurs, for example, in filtering by Infinite Impulse Response (IIR) and Finite Impulse Response (FIR) filters of the signal captured in fixed-point format and converted to the floating-point format. Operations

on denormal data make processing slow, even if corresponding interrupts are disabled. To avoid the slowdown effect in denormal data processing, the Intel IPP threshold functions can be applied to the input signal before filtering. For example:

```
if (denormal_data)
    ippThreshold_LT_32f_I( src, len, 1e-6f );
    ippSFIR_32f( src, dst, len, st );
```

The `1e-6f` value is the threshold level; the input data below that level are set to zero. Because the Intel IPP threshold function is very fast, the execution of these two functions is faster than execution of filter if denormal numbers meet in the source data. Of course, if the denormal data occurs while using the filtering procedure, the threshold functions do not help.

For Intel(R) Pentium(R) 4 processor and later processors, you can set special computation modes - *flush-to-zero* (FTZ) and the *denormals-are-zero* (DAZ). Use the functions `ippSetFlushToZero` and `ippSetDenormAreZeros` to enable these modes. Note that this setting takes effect only when computing is done with the Intel(R) Streaming SIMD Extensions (Intel(R) SSE) and Intel(R) Streaming SIMD Extensions 2 (Intel(R) SSE2) instructions. (See more details

<http://software.intel.com/en-us/articles/x87-and-sse-floating-point-assists-in-ia-32-flush-to-zero-ftz-and-denormals-are-zero-daz/>).

The following table illustrates how denormal data may affect performance and shows the effect of thresholding denormal data. As you can see, thresholding takes only three clocks more. On the other hand, denormal data can cause the application performance to drop x250.

Table 7-1 Performance Resulting from Thresholding Denormal Data

Data/Method	Normal	Denormal	Denormal + Threshold
CPU cycles per element	46	11467	49

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Reusing Buffers

Some Intel IPP functions require internal memory for various optimization strategies. However, memory allocation inside of the function may have a negative impact on performance in some situations, for example, cache misses. To avoid or minimize memory allocation and keep your data in a warm cache, some functions, for example, Fourier transform functions (FFT), can use or reuse memory given as a parameter to the function.

If you call such a function, many times, you can reuse of an external buffer and get better performance.

The following example reuses memory buffers to compute FFT as two FFTs in two separate threads:

```

ippsFFTInitAlloc_C_32fc(&ctxN2, order-1, IPP_FFT_DIV_INV_BY_N, ippAlgHintAccurate);
ippsFFTGetBufSize_C_32fc( ctxN2, &sz );

buffer = sz > 0 ? ippsMalloc_8u( sz ) : 0;

// prepare source data for two FFTs
int phase = 0;
ippsSampleDown_32fc( x, fftlen, xleft, &fftlen2, 2, &phase );
phase = 1;
ippsSampleDown_32fc( x, fftlen, xrght, &fftlen2, 2, &phase );

ippsFFTFwd_CToC_32fc( xleft, Xleft, ctxN2, buffer );
ippsFFTFwd_CToC_32fc( xrght, Xrght, ctxN2, buffer );

```

The external buffer is not necessary. If the pointer to the buffer is 0, the function allocates memory inside.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Using Fast Fourier Transform

Fast Fourier Transform (FFT) is a universal method to increase performance of data processing, especially in the field of digital signal processing where filtering is essential.

The convolution theorem states that filtering of two signals in the spatial domain can be computed as point-wise multiplication in the frequency domain. The data transformation to and from the frequency domain is usually performed using the Fourier transform. You can apply the Finite Impulse Response (FIR) filter to the input signal by using Intel IPP FFT functions, which are optimized for Intel® processors. You can also increase the data array length to the next greater power of two by padding the array with zeroes and then applying the forward FFT function to the input signal and the FIR filter coefficients. Fourier coefficients obtained in this way are multiplied point-wise and the result can easily be transformed back to the spatial domain. The performance gain achieved by using FFT may be very significant.

If the applied filter is the same for several processing iterations, then the once calculated filter coefficients can be reused in each iteration. The twiddle tables and the bit reverse tables are created in the initialization function for the forward and inverse transforms at the same time. The following example presents the main operations of this kind of filtering:

```

ippsFFTInitAlloc_R_32f( &pFFTSpec, fftord, IPP_FFT_DIV_INV_BY_N, ippAlgHintNone );

// perform forward FFT to put source data xx to frequency domain
ippsFFTFwd_RToPack_32f( xx, XX, pFFTSpec, 0 );

// perform forward FFT to put filter coefficients hh to frequency domain
ippsFFTFwd_RToPack_32f( hh, HH, pFFTSpec, 0 );

```

```
/// point-wise multiplication in freq-domain is convolution
ippsMulPack_32f_I( HH, XX, fftlen );

/// perform inverse FFT to get result in time-domain
ippsFFTInv_PackToR_32f( XX, yy, pFFTSpec, 0 );

/// free FFT tables
ippsFFTFree_R_32f( pFFTSpec );
```

The zeros in the example above could be pointers to the external memory, which is another way to increase performance. Another way to significantly improve performance is using FFT and multiplication for processing large size data.

The signal processing FIR filter in Intel IPP is implemented using FFT, and you do not need to create a special implementation of the FIR functions.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Using the Intel IPP Performance Test Tool

Intel IPP installation includes a command-line tool for performance testing - The Intel IPP Performance Test Tool (PTT). It does performance testing for each Intel IPP functions on the same hardware platforms that are valid for the related Intel IPP libraries.

The performance test executable files `ps_ipp*.exe` files (* - functional domain suffix) are placed in the `\ipp\tools\<arch>\perfsys` directory. For example, `ps_ipps.exe` is a tool to measure performance of all Intel IPP signal processing functions.

With the command-line options you can create a list of functions to test and set required parameters with which the function is called during the performance test. You can define the functions and parameters in the `.ini` file, or enter them directly from the console.

The results are saved in a `.csv` file. The course of test is displayed on the console, you can be save it in the `.txt` file.

In the enumeration mode, the Intel IPP PTT creates only a list of the tested functions on the console and can store it in the `.txt` or `.csv` files.

The command-line format is:

```
<ps_FileName>.exe [option_1] [option_2]... [option_n]
```

A short reference for the command-line options can be displayed on the console. To invoke the reference, just enter `-?` or `-h` in the command-line:

```
ps_ipps .exe -h
```

The command-line options are divided into several groups by functionality. You can enter options in arbitrary order with at least one space between each option name. Some options (like `-r`, `-R`, `-o`, `-O`) may be entered several times with different file names, and option `-f` may be entered several times with different function patterns. For detailed descriptions of the performance test tool command-line options, see [Performance Test Tool Command-Line Options](#).

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Examples of Performance Test Tool Command Lines

The following are examples of Intel IPP performance test tool command lines.

Example 1. Running in the standard mode:

```
ps_ippch.exe -B -r
```

This command tests all Intel IPP string functions by the default timing method on standard data (`-B` option) and stores results in file `ps_ippch.csv` (`-r` option).

Example 2. Testing selected functions:

```
ps_ipps.exe -f FIRLMS_32f -r firllms.csv
```

This command tests the signal processing function `FIRLMS_32f` (`-f` option), and stores the results in `firllms.csv` (`-r` option).

Example 3. Retrieving function lists:

```
ps_ippvc.exe -e -o vc_list.txt
```

This command lists all video coding functions (`-e` option) in the output file `vc_list.txt` (`-o` option).

```
ps_ippvc.exe -e -r H264.csv -f H264
```

This command displays the list of functions with names containing `H264` (`-f` option) that can be tested (`-e` option) on the console and stores the list in `H264.csv` (`-r` option).

Example 4. Launching performance test tool with the .ini file:

```
ps_ipps.exe -B -I
```

This command creates `ps_ippi.ini` file after the first run (`-I` option) to test all signal processing functions using the default timing method on standard data (`-B` option).

```
ps_ippi.exe -i -r
```

This command performs the second run to test all functions with parameters specified in the `ps_ippi.ini` file (`-i` option), and generates the output file `ps_ippi.csv` (`-r` option).

For detailed descriptions of performance test tool command line options, see [Performance Test Tool Command Line Options](#).

Performance Test Tool Command Line Options



The following table lists the Intel(R) IPP Performance Test tool (PTT) command-line options by functional groups, and briefly describes each option.

Performance Test Tool Command Line Options

Groups	Options	Descriptions
1. Adjusting Console Input	-A	Prompt for the parameters before every test from console
	-B	Batch mode
2. Managing Output	-r[<file-name>]	Create <file-name>.csv file and write PTT results to it
	-R[<file-name>]	Add test results to the file <file-name>.csv
	-H[ONLY]	Add 'Interest' column to .csv file [and run only hot tests]
	-o[<file-name>]	Create <file-name>.txt file and write console output to it
	-O[<file-name>]	Add console output to the file <file-name>.txt
	-L <ERR/WARN/PARM/INFO/TRACE>	Set detail level of the console output
	-u[<file-name>]	Create <file-name>.csv file and write summary table ('_sum' is added to the file name)
	-U[<file-name>]	Add summary table to the file <file-name>.csv ('_sum' is added to the file name)
	-e	Enumerate tests and exit
	-g[<file-name>]	Create signal file just at the end of the whole testing
	-s[-]	Sort or don't sort functions (sort mode is default)
3. Selecting Functions for Testing	-f <or-pattern>	Run tests for functions with pattern in their names, case sensitive

Groups	Options	Descriptions
	-f-<not-pattern>	Do not test functions with <code>pattern</code> in their names, case sensitive
	-f+<and-pattern>	Run tests only for functions with <code>pattern</code> in their names, case sensitive
	-f=<eq-pattern>	Run tests of functions with name <code>eq-pattern</code> , case sensitive
	-F<func-name>	Start testing from function with the full name <code>func-name</code> , case sensitive
4. Operation with .ini Files	-i[<file-name>]	Read PTT parameters from the file <code><file-name>.ini</code>
	-I[<file-name>]	Write PTT parameters to the file <code><file-name>.ini</code> and exit
	-P	Read tested function names from the <code>.ini</code> file
5. Adjust default directories and file names for input and output	-n<title-name>	Set default title name for <code>.ini</code> file and output files
	-p<dir-name>	Set default directory for <code>.ini</code> file and input test data files
	-l<dir-name>	Set default directory for output files
6. Direct Data Input	-d<name>=<value>	Set PTT parameter value
7. Process priority	-Y<HIGH/NORMAL>	Set high or normal process priority (normal is default)
8. Setting environment	-N<num-threads>	Call <code>ippSetNumThreads(<num-threads>)</code>
9. Getting help	-h	Display short help and exit
	-hh	Display extended help and exit
	-h<option>	Display extended help for the specified option and exit

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Optimization Notice

Notice revision #20110804

Intel(R) IPP Samples

This appendix describes the types of Intel IPP sample code available for developers to learn how to use Intel IPP, gives the source code example files by categories with links to view the sample code, and explains how to build and run sample applications.

For information on configuring Microsoft* Visual* C++ project files for the Intel IPP samples, see "[Creating Visual C++ 2005 Project Files for the Intel® IPP Samples](#)".

Types of Intel IPP Sample Code

Intel IPP sample code are provided to help you to learn how to use the Intel Integrated Performance Primitives. They are intended only to demonstrate how to use the APIs and how to build applications in different development environments.

Table B-1 Types of Intel IPP Sample Code

Type	Description
Application-level samples	These samples illustrate how to build a wide variety of applications such as encoders, decoders, viewers, and players using the Intel IPP APIs (see Using Intel IPP Samples).
Source Code Samples	These platform independent examples show basic techniques for using Intel IPP functions to perform such operations as performance measurement, time-domain filtering, affine transformation, canny edge detection, and more. Each example consists of 1-3 source code files (.cpp) (see Source Code Samples).
Code Examples	These code examples (or code <i>snippets</i>) are very short programs demonstrating how to call a particular Intel IPP function. Numerous code examples are contained in the <i>Intel IPP Reference Manual</i> as part of the function descriptions.

Source Code Samples

The following table presents the files with the source code for the Intel IPP samples by categories and links to them.

Category	Summary	Description and Links
Basic Techniques	Introduction to programming with Intel IPP functions	Performance measurement: GetClocks.cpp Copying data: Copy.cpp Optimizing table-based Functions:

Category	Summary	Description and Links
		LUT.cpp
Digital Filtering	Fundamentals of signal processing	Executing the: DFT.cpp Filtering with FFT: FFTFilter.cpp Time-domain filtering: FIR.cpp
Audio Processing	Audio signal generation and manipulation	Generating DTMF tones: DTMF.cpp Using IIR to create an echo: IIR.cpp Using FIRMR to resample a signal: Resample.cpp
Image Processing	Creating and processing a whole image or part of an image	Allocating, initializing, and copying an image: Copy.cpp Rectangle of interest sample wrapper: ROI.h ROI.cpp ROITest.cpp Mask image sample wrapper: Mask.h Mask.cpp MaskTest.cpp
Image Filtering and Manipulation	General image affine transformations	Wrapper for resizing an image: Resize.h Resize.cpp ResizeTest.cpp Wrapper for rotating an image: Rotate.h Rotate.cpp RotateTest.cpp Wrapper for doing an affine transform on an image: Affine.h Affine.cpp AffineTest.cpp
Graphics and Physics	Vector and small matrix arithmetic functions	ObjectViewer application: ObjectViewerDoc.cpp ObjectViewerDoc.h ObjectViewerView.cpp ObjectViewerView.h Transforming vertices and normals: CTestView::OnMutateModel Projecting an object onto a plane: CTestView::OnProjectPlane Drawing a triangle under the cursor: CTestView::Draw Performance comparison, vector vs. scalar: perform.cpp Performance comparison, buffered vs. unbuffered: perform2.cpp
Special-Purpose Domains	Cryptography and computer vision usage	RSA key generation and encryption: rsa.cpp rsa.h rsatest.cpp bignum.h bignum.cpp Canny edge detection class: canny.cpp canny.h cannytest.cpp filter.h filter.cpp Gaussian pyramids class: pyramid.cpp pyramid.h pyramidtest.cpp

Using Intel IPP Samples

Download the Intel IPP application-level samples from <http://www.intel.com/software/products/ipp/samples.htm>. These samples are updated in each version of Intel IPP. For best results upgrade the Intel IPP samples when a new version of Intel IPP is available.

Several common samples are included with the product. You can find them in the `<ipp directory>\Samples\en_US\IPP`.

System Requirements

The specific system requirements for each sample are listed in the `readme.htm` document in the root directory of each sample. Most common requirements are listed below.

Hardware requirements:

- A system based on an Intel® Pentium® processor, Intel® Xeon® processor, or a subsequent IA-32 architecture-based processor
- A system based on processors supported the Intel® 64 architecture

Software requirements:

- Intel(R) IPP for the Windows* OS, version 7.0 or higher
- Microsoft* Windows 7, Microsoft* Windows Vista*, Microsoft Windows XP*, Microsoft Windows Server* 2008, or Microsoft Windows Server 2003 operating system
- Microsoft* DirectX* API: 9.0 SDK Update (February 2005) or SDK (December 2005)
- Intel(R) C++ Compiler for Windows OS: versions 12.0, 11.1; Intel(R) Parallel Composer
- Microsoft Visual C++* .NET 2010, Microsoft Visual C++ .NET 2008 or Microsoft Visual C++ .NET 2005 development systems
- To develop an application for a processor supporting the Intel(R) 64 architecture, the Microsoft EM64T Platform SDK is required

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Optimization Notice

Notice revision #20110804

Building Source Code

The specific building procedure is described in the `readme.htm` document in the root directory of each sample. Most common steps are described below.

Set up your build environment by creating an environment variable named `IPPROOT` that points to the root directory of your Intel IPP installation. For example, default `C:\Program Files\Intel\Compiler`.

To build all samples, change your current directory to the root media samples directory and use batch file `build*.bat[option]`.

By default, the batch file searches the last version of compiler (assuming that compiler is installed in the default directory). If you wish to use specific version of the Intel C/C++ compiler or Microsoft C/C++ .NET 2005 compiler, set an option for the batch file in accordance with the following table:

Table B-2 Options for Batch Files

Compiler	Option
Intel C++ Composer XE 2011 for Windows OS	<code>ic1120</code>
Intel C++ Compiler 11.1 for Windows OS	<code>ic1111</code>
Intel Parallel Composer	<code>ipc2011</code> <code>ipc2009</code>
Microsoft Visual C++ .NET 2010 software	<code>c110</code>
Microsoft Visual C++ .NET 2008 software	<code>c19</code>
Microsoft Visual C++ .NET 2005 software	<code>c18</code>

After the successful build, the executable file is placed in the corresponding sample directory:

`<install_dir>\ipp-samples\<sample-name>\bin\win*_compiler,`

where `compiler = c18|c19|c110|ipc2009|ipc2011|ic1111|ic1120`.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Running the Software

To run each sample application, the Intel IPP dynamic link libraries must be on the system's path. See "[Setting Environment Variables](#)" for more details.

Detailed instructions on how to run the application, the list of command line options or menu commands are given in the `readme.htm` document in the root directory of each sample.

Known Limitations

The applications created with the Intel IPP Samples are intended to demonstrate how to use the Intel IPP functions and help you to create your own software. These sample applications have some limitations. The specific limitations for each sample are described in the section "*Known Limitations*" in the `readme.htm` document in the root directory of each sample.



Language Support

The following table lists the different languages that Intel(R) IPP supports. It also provides information on the Intel IPP samples available for each language.

Language Support

Language	Environment	The Sample Description
C++	Makefile	Overloads the Intel IPP C-library functions to create classes for easy signal and image manipulation.
	Microsoft* eMbedded Visual C++* 4.0, or higher WinCE 5.0 SDK for x86	Uses the Intel IPP libraries in the development of applications for the Windows* CE environment for IA-32 architecture.
Fortran	Makefile	N/A
C#*	Microsoft .NET C#*	Uses Intel IPP functions in a C# wrapper class.
Visual Basic*	Microsoft .NET Visual Basic*	The demo-application shows how to call Intel IPP functions from a Visual Basic wrapper class.
Object Pascal	Borland Delphi*	Uses the Intel IPP image processing functions in Borland Delphi*.
Java*	Java Development Kit 1.5.0	Uses the Intel IPP image processing functions in a Java wrapper class.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

See Also

Download the samples from <http://www.intel.com/software/products/ipp/samples.htm>

Using Intel IPP in Java* Applications

You can call Intel IPP functions in your Java application by using the Java* Native Interface (JNI*). There is some overhead associated with JNI use, especially when the input data size is small. Combining several functions into one JNI call and using managed memory can improve the overall performance.

Index

A

avoiding nested parallelization 51

B

basic concepts 25
basic features 15
buffers, reusing 55
building application 24
building samples 68

C

calling functions 25
checking your installation 23
command line options 61
concepts, basic 25
configuring development environment 31
 Microsoft Visual Studio 31
core functions 16
cross-architecture alignment 15

D

data
 thresholding 54
 types 16
detecting processor features 36
detecting processor type 36
dispatching 35
documentation directory 30
domains 17

F

function
 naming 18
 parameters 18

H

hardware requirement 67
header files 25
high-level directory structure 27

I

input data
 types 16
Intel IPP basics 15

J

java applications 72

L

language support 71
libraries, types of 28
library dependencies by domain 48
linking 35, 39, 40, 42, 43, 44, 46, 47, 48
 custom dynamic 43
 dynamic 39, 46
 libraries 46
 examples 48
 methods 39
 selecting libraries 44
 static
 dependencies 48
 libraries 47
 static with dispatching 40
 static without dispatching 42
linking methods comparison 44

M

managing memory 53, 55
 alignment 53
 reusing buffers 55
managing performance 53, 54, 56
 thresholding data 54
 using FFT 56
memory alignment 53
memory, managing 53
multithreading 49
 support 49

N

naming, function 18
nested parallelization, avoiding 51

notational conventions 13
number of threads
 getting 50
 setting 50

O

OpenMP support 49

P

parameters 18
performance test tool 57, 58, 61
 command line examples 58
 command line options 61
performance, managing 53
processor-specific codes 35

R

reusing buffers 55

S

sample code 65
samples 32, 65, 68, 69
 building 68
 creating project 32
 known limitations 69

samples (*continued*)
 running 69
 source code 65
 types of 65
selecting libraries 44
selecting linking methods 39
setting environment variables 24
shared L2 cache, using 50
software requirement 67
static libraries, using 29
supplied libraries 28

T

technical support 9
threading 49
thresholding data 54
types of samples 65

U

using FFT 56
using Intel IPP with Java 72
using Intel IPP with Visual C++ .NET 31
using shared L2 cache 50

V

version information 23