



Intel® Integrated Performance Primitives Developer Guide for Intel® oneAPI Base Toolkit

Developer Guide

Notices and Disclaimers

Contents

Notices and Disclaimers	4
What's New	5
Getting Help and Support	6
Notational Conventions	7
Chapter 1: Getting Started with Intel® Integrated Performance Primitives	
Finding Intel® IPP on Your System	8
Setting Environment Variables	10
Compiler Integration	10
Building Intel® IPP Applications	11
Using Intel® IPP Examples	13
Intel® IPP Examples Directory Structure	13
Building Intel® IPP Examples.....	14
Finding the Intel® IPP Documentation	14
Chapter 2: Intel® Integrated Performance Primitives Theory of Operation	
Dispatching	15
Function Naming Conventions	16
Data-domain	16
Primitive vs. Variant Name	16
Data Types.....	16
Descriptor	17
Parameters.....	18
Intel® Integrated Performance Primitives Domain Details	18
Library Dependencies by Domain.....	18
Chapter 3: Linking Your Application with Intel® Integrated Performance Primitives	
Linking Options	20
Automatically Linking Your Microsoft* Visual Studio* Project with Intel IPP	21
Chapter 4: Using Intel® Integrated Performance Primitives Platform-Aware Functions	
Chapter 5: Using Intel® Integrated Performance Primitives Threading Layer (TL) Functions	
Finding Intel® IPP TL Source Code Files	24
Building Intel® IPP TL Libraries from Source Code	25
Chapter 6: Using Custom Library Tool for Intel® Integrated Performance Primitives	
System Requirements for Custom Library Tool	27
Operation Modes	27
Building a Custom DLL with Custom Library Tool	27
Using Console Version of Custom Library Tool	28

Chapter 7: Using Integration Wrappers for Intel® Integrated Performance Primitives	
Chapter 8: Programming Considerations	
Core and Support Functions	32
Channel and Planar Image Data Layouts.....	33
Regions of Interest	34
Managing Memory Allocations	35
Cache Optimizations	36
Chapter 9: Programming with Intel® Integrated Performance Primitives in the Microsoft* Visual Studio* IDE	
Configuring the Microsoft* Visual Studio* IDE to Link with Intel® IPP	37
Using the IntelliSense* Features.....	37
Appendix A: Appendix: Intel(R) IPP Threading and OpenMP* Support	
Using Shared L2 Cache	40
Avoiding Nested Parallelization	40

Notices and Disclaimers

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described 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.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

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

Java is a registered trademark of Oracle and/or its affiliates.

What's New

This Developer Guide documents the Intel® Integrated Performance Primitives (Intel® IPP) for Intel® oneAPI Base Toolkit.

Documentation for older versions of Intel® Integrated Performance Primitives is available for download only. For a list of available documentation downloads by product version, see these pages:

- [Download Documentation for Intel Parallel Studio XE](#)
- [Download Documentation for Intel System Studio](#)

Intel® IPP 2021 document updates

Minor updates have been made to fix inaccuracies in the document.

Getting Help and 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, FAQs, tips and tricks and other support questions, please visit <http://www.intel.com/software/products/support/> and the Intel IPP forum <http://software.intel.com/en-us/forums/intel-integrated-performance-primitives/>.

NOTE

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

Notational Conventions

The following font and symbols conventions are used in this document:

<i>Italic</i>	<i>Italic</i> is used for emphasis and also indicates document names in body text, for example: see <i>Intel IPP Developer Reference</i> .
Monospace lowercase	Indicates filenames, directory names, and pathnames.
Monospace lowercase mixed with UPPERCASE	Indicates commands and command-line options, for example: <code>vars.bat ia32</code>
UPPERCASE MONOSPACE	Indicates system variables, for example: <code>\$PATH</code> .
monospace italic	Indicates a parameter in discussions, such as routine parameters, for example: <code>pSrc</code> ; makefile parameters, for example: <code>function_list</code> . When enclosed in angle brackets, indicates a placeholder for an identifier, an expression, a string, a symbol, or a value, for example: <code><ipp directory></code> .
[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.

Getting Started with Intel® Integrated Performance Primitives

1

This chapter helps you start using Intel® Integrated Performance Primitives (Intel® IPP) by giving a quick overview of some fundamental concepts and showing how to build an Intel® IPP program.

Finding Intel® IPP on Your System

Intel® Integrated Performance Primitives (Intel® IPP) installs in the subdirectory referred to as *<ipp directory>* inside *<install_dir>*. By default, the *<install_dir>* is:

- On Windows* OS: C:\Program files (x86)\Intel\oneapi (on certain systems, instead of Program Files (x86), the directory name is Program Files)
- On Linux* OS:
 - admin: /opt/intel/oneapi
 - user: ~/intel/oneapi
- On macOS*: /opt/intel/oneapi

The tables below describe the structure of the high-level directories on:

- [Windows* OS](#)
- [Linux* OS](#)
- [macOS*](#)

Windows* OS:

Directory	Contents
env	Batch files to set environmental variables in the user shell
include	Header files for the library functions
lib/ia32	Single-threaded static libraries for the IA-32 architecture
lib/intel64	Single-threaded static libraries for the Intel® 64 architecture
lib/<arch>/tl/<threading_type>, where <arch> is one of {ia32, intel64}, and <threading_type> is one of {tbb, openmp}	Threading Layer static and dynamic libraries
redist/ia32	Single-threaded DLLs for applications running on processors with the IA-32 architecture
redist/intel64	Single-threaded DLLs for applications running on processors with the Intel® 64 architecture
components	Intel IPP interfaces and example files
tools/custom_library_tool_python	Command-line and GUI tool for building custom dynamic libraries

Directory	Contents
tools/<arch>/staticlib	Header files for disabling the Intel IPP static dispatcher (linking with one CPU-optimized variant)

Linux* OS:

Directory	Contents
env	Batch files to set environmental variables in the user shell
include	Header files for the library functions
lib/ia32	Single-threaded static libraries for the IA-32 architecture
lib/intel64	Single-threaded static libraries for the Intel® 64 architecture
lib/<arch>/tl/<threading_type>, where <arch> is one of {ia32, intel64}, and <threading_type> is one of {tbb, openmp}	Threading Layer static and dynamic libraries
lib/<arch>/nonpic	Non-PIC single-threaded static libraries
lib/<arch>_and	Intel IPP libraries for Android*
components	Intel IPP interfaces and example files
tools/custom_library_tool_python	Command-line and GUI tool for building custom dynamic libraries
tools/<arch>/staticlib	Header files for disabling the Intel IPP static dispatcher (linking with one CPU-optimized variant)

macOS*:

Directory	Contents
env	Batch files to set environmental variables in the user shell
include	Header files for the library functions
lib/intel64	Single-threaded static libraries for the Intel® 64 architecture
lib/<arch>_and	Intel IPP libraries for Android*
components	Intel IPP interfaces and example files
tools/custom_library_tool_python	Header files for disabling the Intel IPP static dispatcher (linking with one CPU-optimized variant)

See Also

[Notational Conventions](#)

Setting Environment Variables

When the installation of Intel IPP is complete, set the environment variables in the command shell using one of the script files in the `env` subdirectory of the Intel IPP installation directory:

On Windows* OS:

`vars.bat`

for the IA-32 and Intel® 64 architectures.

On Linux* OS and macOS*:

`vars.sh`

Linux* OS: for the IA-32 and Intel® 64 architectures.

macOS*: for the Intel® 64 architectures.

When using the `vars` script, you need to specify the architecture as a parameter. For example:

- `vars.bat ia32`
sets the environment for Intel IPP to use the IA-32 architecture on Windows* OS.
- `. vars.sh intel64`
sets the environment for Intel IPP to use the Intel® 64 architecture on Linux* OS.

The scripts set the following environment variables:

Windows* OS	Linux* OS	Purpose
IPPROOT	IPPROOT	Point to the Intel IPP installation directory
LIB	n/a	Add the search path for the Intel IPP single-threaded libraries
PATH	LD_LIBRARY_PATH	Add the search path for the Intel IPP single-threaded DLLs
INCLUDE	n/a	Add the search path for the Intel IPP header files

Compiler Integration

Intel® C++ Compiler and Microsoft Visual Studio* compilers simplify developing with Intel® IPP.

On Windows* OS, a default installation of Intel® IPP installs integration plug-ins. These enable the option to configure your Microsoft Visual Studio* project for automatic linking with Intel IPP.

Intel® C++ Compiler also provides command-line parameters to set the link/include directories:

- On Windows* OS:
`/Qipp-link:{dynamic|static}` and `/Qipp`
- On Linux* OS:
`-ipp-link={dynamic|static}`

See Also

[Automatically Linking Your Microsoft* Visual Studio* Project with Intel IPP](#)
[Linking Your Application with Intel\(R\) IPP](#)

Building Intel® IPP Applications

The code example below represents a short application to help you get started with Intel® IPP:

```
#include "ipp.h"
#include <stdio.h>
int main(int argc, char* argv[])
{
    const IppLibraryVersion *lib;
    IppStatus status;
    Ipp64u mask, emask;

    /* Init IPP library */
    ippInit();
    /* Get IPP library version info */
    lib = ippGetLibVersion();
    printf("%s %s\n", lib->Name, lib->Version);

    /* Get CPU features and features enabled with selected library level */
    status = ippGetCpuFeatures( &mask, 0 );
    if( ippStsNoErr == status ) {
        emask = ippGetEnabledCpuFeatures();
        printf("Features supported by CPU\tyby IPP\n");
        printf("-----\n");
        printf(" ippCPUID_MMX      = ");
        printf("%c\t%c\t", ( mask & ippCPUID_MMX ) ? 'Y':'N', ( emask & ippCPUID_MMX ) ? 'Y':'N');
        printf("Intel(R) Architecture MMX technology supported\n");
        printf(" ippCPUID_SSE      = ");
        printf("%c\t%c\t", ( mask & ippCPUID_SSE ) ? 'Y':'N', ( emask & ippCPUID_SSE ) ? 'Y':'N');
        printf("Intel(R) Streaming SIMD Extensions\n");
        printf(" ippCPUID_SSE2     = ");
        printf("%c\t%c\t", ( mask & ippCPUID_SSE2 ) ? 'Y':'N', ( emask & ippCPUID_SSE2 ) ? 'Y':'N');
        printf("Intel(R) Streaming SIMD Extensions 2\n");
        printf(" ippCPUID_SSE3     = ");
        printf("%c\t%c\t", ( mask & ippCPUID_SSE3 ) ? 'Y':'N', ( emask & ippCPUID_SSE3 ) ? 'Y':'N');
        printf("Intel(R) Streaming SIMD Extensions 3\n");
        printf(" ippCPUID_SSSE3    = ");
        printf("%c\t%c\t", ( mask & ippCPUID_SSSE3 ) ? 'Y':'N', ( emask & ippCPUID_SSSE3 ) ?
'Y':'N');
        printf("Intel(R) Supplemental Streaming SIMD Extensions 3\n");
        printf(" ippCPUID_MOVBE    = ");
        printf("%c\t%c\t", ( mask & ippCPUID_MOVBE ) ? 'Y':'N', ( emask & ippCPUID_MOVBE ) ?
'Y':'N');
        printf("The processor supports MOVBE instruction\n");
        printf(" ippCPUID_SSE41    = ");
        printf("%c\t%c\t", ( mask & ippCPUID_SSE41 ) ? 'Y':'N', ( emask & ippCPUID_SSE41 ) ?
'Y':'N');
        printf("Intel(R) Streaming SIMD Extensions 4.1\n");
        printf(" ippCPUID_SSE42    = ");
        printf("%c\t%c\t", ( mask & ippCPUID_SSE42 ) ? 'Y':'N', ( emask & ippCPUID_SSE42 ) ?
'Y':'N');
        printf("Intel(R) Streaming SIMD Extensions 4.2\n");
        printf(" ippCPUID_AVX      = ");
        printf("%c\t%c\t", ( mask & ippCPUID_AVX ) ? 'Y':'N', ( emask & ippCPUID_AVX ) ? 'Y':'N');
        printf("Intel(R) Advanced Vector Extensions instruction set\n");
        printf(" ippAVX_ENABLEDBYOS = ");
        printf("%c\t%c\t", ( mask & ippAVX_ENABLEDBYOS ) ? 'Y':'N', ( emask & ippAVX_ENABLEDBYOS ) ?
'Y':'N');
```

```

printf("The operating system supports Intel(R) AVX\n");
printf(" ippCPUID_AES          = ");
printf("%c\t%c\t", ( mask & ippCPUID_AES ) ? 'Y':'N', ( emask & ippCPUID_AES ) ? 'Y':'N');
printf("Intel(R) AES instruction\n");
printf(" ippCPUID_SHA          = ");
printf("%c\t%c\t", ( mask & ippCPUID_SHA ) ? 'Y':'N', ( emask & ippCPUID_SHA ) ? 'Y':'N');
printf("Intel(R) SHA new instructions\n");
printf(" ippCPUID_CLMUL         = ");
printf("%c\t%c\t", ( mask & ippCPUID_CLMUL ) ? 'Y':'N', ( emask & ippCPUID_CLMUL ) ?
'Y':'N');
printf("PCLMULQDQ instruction\n");
printf(" ippCPUID_RDRAND        = ");
printf("%c\t%c\t", ( mask & ippCPUID_RDRAND ) ? 'Y':'N', ( emask & ippCPUID_RDRAND ) ?
'Y':'N');
printf("Read Random Number instructions\n");
printf(" ippCPUID_F16C          = ");
printf("%c\t%c\t", ( mask & ippCPUID_F16C ) ? 'Y':'N', ( emask & ippCPUID_F16C ) ? 'Y':'N');
printf("Float16 instructions\n");
printf(" ippCPUID_AVX2          = ");
printf("%c\t%c\t", ( mask & ippCPUID_AVX2 ) ? 'Y':'N', ( emask & ippCPUID_AVX2 ) ? 'Y':'N');
printf("Intel(R) Advanced Vector Extensions 2 instruction set\n");
printf(" ippCPUID_AVX512F       = ");
printf("%c\t%c\t", ( mask & ippCPUID_AVX512F ) ? 'Y':'N', ( emask & ippCPUID_AVX512F ) ?
'Y':'N');
printf("Intel(R) Advanced Vector Extensions 3.1 instruction set\n");
printf(" ippCPUID_AVX512CD      = ");
printf("%c\t%c\t", ( mask & ippCPUID_AVX512CD ) ? 'Y':'N', ( emask & ippCPUID_AVX512CD ) ?
'Y':'N');
printf("Intel(R) Advanced Vector Extensions CD (Conflict Detection) instruction set\n");
printf(" ippCPUID_AVX512ER      = ");
printf("%c\t%c\t", ( mask & ippCPUID_AVX512ER ) ? 'Y':'N', ( emask & ippCPUID_AVX512ER ) ?
'Y':'N');
printf("Intel(R) Advanced Vector Extensions ER instruction set\n");
printf(" ippCPUID_ADCOX         = ");
printf("%c\t%c\t", ( mask & ippCPUID_ADCOX ) ? 'Y':'N', ( emask & ippCPUID_ADCOX ) ?
'Y':'N');
printf("ADCX and ADOX instructions\n");
printf(" ippCPUID_RDSEED        = ");
printf("%c\t%c\t", ( mask & ippCPUID_RDSEED ) ? 'Y':'N', ( emask & ippCPUID_RDSEED ) ?
'Y':'N');
printf("The RDSEED instruction\n");
printf(" ippCPUID_PREFETCHW     = ");
printf("%c\t%c\t", ( mask & ippCPUID_PREFETCHW ) ? 'Y':'N', ( emask & ippCPUID_PREFETCHW ) ?
'Y':'N');
printf("The PREFETCHW instruction\n");
printf(" ippCPUID_KNC           = ");
printf("%c\t%c\t", ( mask & ippCPUID_KNC ) ? 'Y':'N', ( emask & ippCPUID_KNC ) ? 'Y':'N');
printf("Intel(R) Xeon Phi(TM) Coprocessor instruction set\n");
}
return 0;
}

```

This application consists of three sections:

1. Initialize the Intel IPP library. This stage is required to take advantage of full Intel IPP optimization. The `ippInit()` function detects the processor type and sets the dispatcher to use the processor-specific code of the Intel® IPP library corresponding to the instruction set capabilities available. If your application runs without `ippInit()`, the Intel IPP library is auto-initialized with the first call of the Intel IPP function from any domain that is different from `ippCore`.

In certain debugging scenarios, it is helpful to force a specific implementation layer using `ippSetCpuFeatures()`, instead of the best as chosen by the dispatcher.

2. Get the library layer name and version. You can also get the version information using the `ippversion.h` file located in the `/include` directory.
3. Show the hardware optimizations used by the selected library layer and supported by CPU.

Building the First Example with Microsoft Visual Studio* Integration on Windows* OS

On Windows* OS, Intel IPP applications are significantly easier to build with Microsoft* Visual Studio*. To build the code example above, follow the steps:

1. Start Microsoft Visual Studio* and create an empty C++ project.
2. Add a new c file and paste the code into it.
3. Set the include directories and the linking model as described in [Automatically Linking Your Microsoft* Visual Studio* Project with Intel IPP](#).
4. Compile and run the application.

If you did not install the integration plug-in, configure your Microsoft* Visual Studio* IDE to build Intel IPP applications following the instructions provided in [Configuring the Microsoft Visual Studio* IDE to Link with Intel® IPP](#).

Building the First Example on Linux* OS

To build the code example above on Linux* OS, follow the steps:

1. Paste the code into the editor of your choice.
2. Make sure the compiler and Intel IPP variables are set in your shell. For information on how to set environment variables see [Setting Environment Variables](#).
3. Compile with the following command: `icc ipptest.cpp -o ipptest -I $IPPROOT/include -L $IPPROOT/lib/<arch> -lippi -lipps -lippcore`. For more information about which Intel IPP libraries you need to link to, see [Library Dependencies by Domain](#) and [Linking Options](#).
4. Run the application.

See Also

[Automatically Linking Your Microsoft* Visual Studio* Project with Intel IPP](#)

[Configuring the Microsoft Visual Studio* IDE to Link with Intel® IPP](#)

[Setting Environment Variables](#)

[Library Dependencies by Domain](#)

[Linking Options](#)

[Dispatching](#)

[Intel® IPP Examples Directory Structure](#)

Using Intel® IPP Examples

This section provides information on Intel IPP examples directory structure and examples build system.

Intel® IPP Examples Directory Structure

The Intel IPP package includes code examples, located in the `components_and_examples_<os>.zip` archive at the `<ipp_directory>/components/` subdirectory. The `examples_core` subdirectory inside the archive contains the following files and directories:

Directory	Contents
<code>common</code>	Common code files for all examples
<code>documentation</code>	Documentation for the Intel IPP examples (<code>ipp-examples.html</code>)

Directory	Contents
ipp_custom_dispatcher	Custom dispatcher usage example
ipp_fft	Fast Fourier transformation example
ipp_morphology	Morphological reconstruction example
ipp_resize_mt	Image resizing example
ipp_thread	Example of external threading of Intel IPP functions

NOTE

Intel® IPP samples are no longer in active development and available as a separate download.

See Also

[Finding Intel® IPP on Your System](#)

Building Intel® IPP Examples

For building instructions refer to `examples_core/documentation/ipp-examples.html` provided with the `<ipp_directory>/components/components_and_examples_<os>.zip` archive.

See Also

[Intel® IPP Examples Directory Structure](#)

Finding the Intel® IPP Documentation

You can find getting started instructions and a listing of all the available online documents with links in the `get_started.htm` file available in the following directory:

- On Windows* OS: `C:\Program files (x86)\Intel\oneAPI\ipp\<version>\documentation` (on certain systems, instead of `Program Files (x86)`, the directory name is `Program Files`)
- On Linux* OS and macOS* : `/opt/intel/oneapi/ipp/<version>/documentation`

Additional documentation on the Intel IPP examples (`documentation/ipp-examples.html`) is available in the `components_and_examples_<os>.zip` archive at the `<ipp_directory>/components/` subdirectory.

The Intel IPP forum and knowledge base can be useful locations to search for questions not answered by the documents above. Please see: <http://software.intel.com/en-us/forums/intel-integrated-performance-primitives/> .

See Also

[Finding Intel® IPP on Your System](#)

Intel® Integrated Performance Primitives Theory of Operation

2

This section discusses dispatching of the Intel® Integrated Performance Primitives (Intel® IPP) libraries to specific processors, provides functions and parameters naming conventions, and explains the data types on which Intel IPP performs operations. This section also provides Intel IPP domain details, including existing library dependencies by domain.

Dispatching

Intel® IPP uses multiple function implementations optimized for various CPUs. Dispatching refers to detection of your CPU and selecting the corresponding Intel IPP binary path. For example, the `ippie9` library in the `/redist/intel64` directory contains the image processing libraries optimized for 64-bit applications on processors with Intel® Advanced Vector Extensions (Intel® AVX) enabled such as the 2nd Generation Intel® Core™ processor family.

A single Intel IPP function, for example `ippscopy_8u()`, may have many versions, each one optimized to run on a specific Intel® processor with specific architecture, for example, the 64-bit version of this function optimized for the 2nd Generation Intel® Core™ processor is `e9_ippscopy_8u()`, and version optimized for 64-bit applications on processors with Intel® Streaming SIMD Extensions 4.2 (Intel® SSE 4.2) is `y8_ippscopy_8u()`. This means that a prefix before the function name determines CPU model. However, during normal operation the dispatcher determines the best version and you can call a generic function (`ippscopy_8u` in this example).

Intel® IPP is designed to support application development on various Intel® architectures. This means that the API definition is common for all processors, while the underlying function implementation takes into account the strengths of each hardware generation.

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

The following table shows processor-specific codes that Intel IPP uses:

Description of Codes Associated with Processor-Specific Libraries

IA-32 Intel® architecture	Intel® 64 architecture	Windows*	Linux* OS	Description
w7		+	+	Optimized for processors with Intel SSE2
	m7	+	+	Optimized for processors with Intel SSE3
s8	n8	+	+	Optimized for processors with Supplemental Streaming SIMD Extensions 3 (SSSE3)
p8	y8	+	+	Optimized for processors with Intel SSE4.2
g9	e9	+	+	Optimized for processors with Intel® Advanced Vector Extensions (Intel® AVX) and Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)
h9	l9	+	+	Optimized for processors with Intel® Advanced Vector Extensions 2 (Intel® AVX2)
	n0	+	+	Optimized for 2nd Generation Intel® Xeon Phi™ Processor

IA-32 Intel® archit ectur e	Intel® 64 architecture	Windo ws*	Linux* OS	Description
	k0	+	+	Optimized for processors with Intel® Advanced Vector Extensions 512 (Intel® AVX-512)
Product and Performance Information				
Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex .				
Notice revision #20201201				

Function Naming Conventions

Intel IPP functions have the same naming conventions for all domains.

Function names in Intel IPP have the following general format:

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

NOTE

The core functions in Intel IPP do not need an input data type. These functions have `ipp` as a prefix without the data-domain field. For example, `ippGetStatusString`.

See Also

[Core and Support Functions](#)

Data-domain

The *data-domain* element is a single character indicating type of input data. Intel IPP supports the following data-domains:

s	one-dimensional operations on signals, vectors, buffers
i	two-dimensional operations on images, video frames

Primitive vs. Variant Name

The *name* element identifies the algorithm or operation of the function. The low-level algorithm that function implements is a *primitive*. This algorithm often has several *variants* for different data types and implementation variations.

For example, the `CToC` modifier in the `ippFFTInv_CToC_32fc` function signifies that the inverse fast Fourier transform operates on complex floating point data, performing the complex-to-complex (CToC) transform.

Data Types

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

```
<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* element contains only one value.

If a function operates on source and destination signals 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 more information about supported data types see the *Intel® IPP Reference Manual* available in the Intel® Software Documentation Library.

See Also

[Intel® Software Documentation Library](#)

Descriptor

The optional *descriptor* element describes the data associated with the operation. Descriptors are individual characters that indicate additional details of the operation.

The Intel IPP functions use the following descriptors:

Descriptor	Description	Example
A	Image data contains an alpha channel as the last channel, requires C4, alpha-channel is not processed.	ippiFilterMax_8u_AC4R
AXX	Advanced arithmetic operations with <i>xx</i> bits of accuracy.	ippsPowx_32f_A11
C	The function operates on a specified channel of interest (COI) for each source image.	ippiSet_8u_C3CR
C <i>n</i>	Image data consists of <i>n</i> channels. Possible values for <i>n</i> : 1, 2, 3, 4.	ippiFilterBorder_32f_C1R
DX	Signal is <i>x</i> -dimensional (default is D1).	ippsConcat_8u_D2
D		
I	Operation is in-place (default is not-in-place).	ippsAdd_16s_I
L	Platform-aware functions (see Using Intel IPP Platform-Aware Functions for details).	ippiAdd_8u_C1RSfs_L
TL	Threading layer functions (see Using Intel IPP Threading Layer Functions for details).	ippiAdd_8u_C1RSfs_LT
M	Operation uses a mask to determine pixels to be processed.	ippiCopy_8u_C1MR
P <i>n</i>	Image data consists of <i>n</i> discrete planar (not-interleaved) channels with a separate pointer to each plane. Possible values for <i>n</i> : 1, 2, 3, 4.	ippiAlphaPremul_8u_AP4R
R	Function operates on a defined region of interest (ROI) for each source image.	ippiMean_8u_C4R

Descriptor	Description	Example
s	Saturation and no scaling (default).	ippiConvert_16s16u_C1Rs
Sfs	Saturation and fixed scaling mode (default is saturation and no scaling).	ippsConvert_16s8s_Sfs

The descriptors in function names are presented in the function name in alphabetical order.

Some data descriptors are default for certain operations and 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* element specifies the function parameters (arguments).

The order of parameters is as follows:

- All source operands. Constants follow vectors.
- All destination operands. Constants follow vectors.
- Other, operation-specific parameters.

A parameter name has the following conventions:

- All parameters defined as pointers start with *p*, for example, *pPhase*, *pSrc*; parameters defined as double pointers start with *pp*, for example, *ppState*. All parameters defined as values start with a lowercase letter, for example, *val*, *src*, *srcLen*.
- Each new part of a parameter name starts with an uppercase character, without underscore; for example, *pSrc*, *lenSrc*, *pDlyLine*.
- Each parameter name specifies its functionality. Source parameters are named *pSrc* or *src*, in some cases followed by names or numbers, for example, *pSrc2*, *srcLen*. Output parameters are named *pDst* or *dst* followed by names or numbers, for example, *pDst2*, *dstLen*. For in-place operations, the input/output parameter contains the name *pSrcDst* or *srcDst*.

Intel® Integrated Performance Primitives Domain Details

Intel IPP is divided into groups of related functions. Each subdivision is called *domain*, and has its own header file, static libraries, dynamic libraries, and tests. The table below lists each domain's code, header and functional area.

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, include `ipp.h` in your application for forward compatibility. If you want to use cryptography functions, you must directly include `ippcp.h` in your application.

* available only within the Intel® System Studio suite

Library Dependencies by Domain

When you link to a certain Intel® IPP domain library, you must also link to the libraries on which it depends. The following table lists library dependencies by domain.

Library Dependencies by Domain

Domain	Domain Code	Depends on
Color Conversion	CC	Core, VM, S, I

Domain	Domain Code	Depends on
String Operations	CH	Core, VM, S
Computer Vision	CV	Core, VM, S, I
Data Compression	DC	Core, VM, S
Image Processing	I	Core, VM, S
Signal Processing	S	Core, VM
Vector Math	VM	Core

To find which domain your function belongs to, refer to the *Intel® IPP Developer Reference* available in the Intel® Software Documentation Library.

See Also

[Intel® Software Documentation Library](#)

Linking Your Application with Intel® Integrated Performance Primitives

3

This section discusses linking options available in Intel® Integrated Performance Primitives (Intel® IPP).

The Intel IPP library supports the following linking options:

- Single-threaded dynamic
- Single-threaded static
- Threading Layer static
- Threading Layer dynamic

Product and Performance Information

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Notice revision #20201201

Linking Options

Intel® Integrated Performance Primitives (Intel® IPP) is distributed as:

- **Static library:** static linking results in a standalone executable
- **Dynamic/shared library:** dynamic linking defers function resolution until runtime and requires that you bundle the redistributable libraries with your application

The following table provides description of libraries available for linking.

	Single-threaded (non-threaded)	Threading Layer (externally threaded)
Description	Suitable for application-level threading	Implementation of application-level threading depends on single-threaded libraries
Found in	Main package After installation: <code><ipp directory>/lib/<arch> (static)</code> and <code><ipp directory>/redist/<arch> (dynamic)</code>	Main package After installation: <code><ipp directory>/lib/<arch>/tl/<threading_type></code> , where <code><threading_type></code> is one of {tbb, openmp}
Static linking	Windows* OS: mt suffix in a library name (<code>ipp<domain>mt.lib</code>) Linux* OS: no suffix in a library name (<code>libipp<domain>.a</code>)	Windows* OS: <code>_mt_tl_<threading_sfx></code> suffix in a library name (<code>ipp<domain>_mt_tl_<threading_sfx>.lib</code>) Linux* OS: <code>_tl_<threading_sfx></code> suffix in a library name (<code>libipp<domain>_tl_<threading_sfx>.a</code>)

Dynamic Linking

Default (no suffix)

Windows* OS: `ipp<domain>.dll`Linux* OS: `libipp<domain>.so`

+ single-threaded libraries dependency, where `<threading_sfx>` is one of {tbb, omp}

`_tl_<threading_sfx> suffix`

Windows* OS:

`ipp<domain>_tl_<threading_sfx>.dll`

Linux* OS:

`libipp<domain>_tl_<threading_sfx>.so`

+ single-threaded library dependency, where `<threading_sfx>` is one of {tbb, omp}

To switch between Intel IPP libraries, set the path to the preferred library in system variables or in your project, for example:

- **Windows* OS:**

Single-threaded: `SET LIB=<ipp directory>/lib/<arch>`

Threading Layer: `SET LIB=<ipp directory>/lib/<arch>/tl/<threading_type>`. Additionally, set path to single-threaded libraries: `SET LIB=<ipp directory>/lib/<arch>`

- **Linux* OS:**

Single-threaded: `gcc <options> -L <ipp directory>/lib/<arch>`

Threading Layer: `gcc <options> -L <ipp directory>/lib/<arch>/tl/<threading_type>`.

Additionally, set path to single-threaded libraries: `gcc <options> -L <ipp directory>/lib/<arch>`

NOTE

On Linux* OS, Intel IPP library depends on the following Intel® C++ Compiler runtime libraries: `libirc.a`, `libsvml.a`, and `libimf.a`. You should add a link to these libraries into your project. You can find these libraries in `<intel compiler directory>/lib` folders.

Threading Layer depends on the OpenMP* or Intel® Threading Building Blocks (Intel® TBB) library according to the selected threading type. You can find these libraries in `<intel compiler directory>/lib` or `<tbb directory>/lib` folders.

See Also

[Automatically Linking Your Microsoft* Visual Studio* Project with Intel IPP](#)

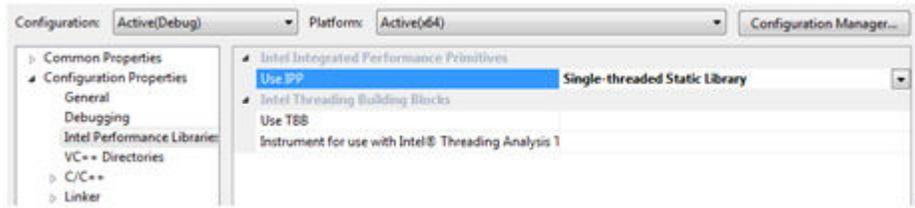
[Configuring the Microsoft Visual Studio* IDE to Link with Intel® IPP](#)

[Library Dependencies by Domain](#)

Automatically Linking Your Microsoft* Visual Studio* Project with Intel IPP

After a default installation of the Intel® IPP, you can easily configure your project to automatically link with Intel IPP. Configure your Microsoft* Visual Studio* project for automatic linking with Intel IPP as follows:

1. Go to **Project>Properties>Configuration Properties>Intel Performance Libraries**.
2. Change the **Use IPP** property setting by selecting one of the options to set the include directories and the linking model, as shown on the screen shot below.



Using Intel® Integrated Performance Primitives Platform-Aware Functions



Intel® Integrated Performance Primitives (Intel® IPP) library provides so-called platform-aware functions for signal and image processing. While the rest of Intel IPP functions support only signals or images of 32-bit integer size, Intel IPP platform-aware functions work with 64-bit object sizes if it is supported by the target platform.

The API of platform-aware functions is similar to the API of other Intel IPP functions and has only slight differences. You can distinguish Intel IPP platform-aware functions by the `L` suffix in the function name, for example, `ippiAdd_8u_C1RSfs_L`. With Intel IPP platform-aware functions you can overcome 32-bit size limitations.

Intel IPP platform-aware functions are declared in separate header files with the `_l` suffix, for example, `ippi_l.h`. However, you do not have to additionally include these headers in your application because they are already included in standard Intel IPP headers (without the `_l` suffix). Platform-aware functions cover only the functionality that is implemented in standard Intel IPP functions, and can be considered as additional flavors to the existing functions declared in standard Intel IPP headers.

Using Intel® Integrated Performance Primitives Threading Layer (TL) Functions

5

Intel® Integrated Performance Primitives (Intel® IPP) library provides threading layer (TL) functions for image processing. Intel IPP TL functions are visual examples of external threading for Intel IPP functions. Taking advantage of multithreaded execution and tile processing, Intel IPP TL functions enable you to overcome 32-bit size limitations.

TL functions are provided as:

- **Pre-built binaries:**
 - Header files have the `_tl` suffix and can be found in: `<ipp_directory>/include`
 - Library files for use with Intel® OpenMP have the `_tl_omp` suffix and can be found in: `<ipp_directory>/lib/<arch>/tl/openmp`
 - Library files for use with Intel® oneTBB have the `_tl_tbb` suffix and can be found in: `<ipp_directory>/lib/<arch>/tl/tbb`
- **Source code samples:** the source code and corresponding header files are available in the `components_and_examples_<os>.zip` archive inside the `<ipp_directory>/components` subdirectory. For more information about the archive contents and source code building instructions, refer to [Finding Intel® IPP TL Source Code Files](#) and [Building Intel® IPP TL Libraries from Source Code](#), respectively.

The API of TL functions is similar to the API of other Intel IPP functions and has only slight differences. You can distinguish Intel IPP TL functions by the `_LT` or `_T` suffix in the function name, for example, `ippiAdd_8u_C1RSfs_LT`. Intel IPP TL functions are implemented as wrappers over Intel IPP functions by using tiling and multithreading with OpenMP* or the Intel® Threading Building Blocks. For implementation details, please see the corresponding source code files.

Product and Performance Information

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Notice revision #20201201

See Also

[Finding Intel® IPP TL Source Code Files](#)

[Building Intel® IPP TL Libraries from Source Code](#)

[Using Intel® Integrated Performance Primitives Platform-Aware Functions](#)

Finding Intel® IPP TL Source Code Files

You can find the Intel IPP TL source code files in the `components_and_examples_<os>.zip` archive available in the `<ipp_directory>/components` subdirectory. The library source code and header files are located in the `interfaces/tl` subdirectory.

Building Intel® IPP TL Libraries from Source Code

You can find the TL libraries source code and the `tl_resize` example in the `/components/interfaces/tl` directory inside the `components_and_examples_<os>` archive available in `<ipp directory>/components/`. Before building an application that uses TL, make sure that the `IPPROOT` environment variable is set correctly and points to the Intel IPP library location, for more information see [Setting Environment Variables](#).

To build Intel IPP TL libraries and the `tl_resize` example, do the following:

Windows* OS

Prerequisites: The `tl_resize` example uses OpenGL rendering to display results. This requires Windows* SDK to be installed on your system. Usually Windows* SDK is provided with the Microsoft* Visual Studio* distribution. Alternatively, you can download Windows* SDK for your version of Windows* OS from <https://www.microsoft.com>. To disable the rendering part of `tl_resize`, remove the `ENABLE_RENDERING` macro from the preprocessors definitions.

1. Open the `tl.sln` file in Microsoft* Visual Studio*.
2. Choose the required configuration in the solution and build the solution using the **Build** command. The example will be linked with the newly built TL libraries from the same solution.

To build TL libraries on the Intel® Threading Building Blocks (Intel® TBB) library, you need to install the Intel TBB library (for the Intel IPP standalone package).

Linux* OS

Prerequisites: The `tl_resize` example uses OpenGL rendering to display results. This requires the following packages to be installed:

- `libx11-dev`
- `libgl1-mesa-dev`

Execute the following commands using `gcc4` or higher:

- To build TL libraries:

```
make libs [ARCH=ia32|intel64] [CONF=release|debug] [TBBROOT=]
```

- To build the `tl_resize` example and TL libraries:

```
make all [ARCH=ia32|intel64] [CONF=release|debug] [RENDERER=0|1] [TBBROOT=]
```

If `TBBROOT` is set to the Intel® TBB installation root, TL libraries will be built with the TBB support. In this case, you need to install Intel TBB library (for the Intel IPP standalone package).

If `TBBROOT` is set to nothing, the OpenMP* support will be used.

See Also

[Setting Environment Variables](#)

Using Custom Library Tool for Intel® Integrated Performance Primitives



With the Intel® Integrated Performance Primitives (Intel® IPP) Custom Library Tool, you can build your own dynamic library containing only the Intel IPP/Intel IPP Cryptography functionality that is necessary for your application.

The use of custom libraries built with the Custom Library Tool provides the following advantages:

- **Package size.** Your package may have much smaller size if linked with a custom library because standard dynamic libraries additionally contain all optimized versions of Intel IPP/Intel IPP Cryptography functions and a dispatcher. The following table compares the contents and size of packages for an end-user application linked with a custom dynamic library and an application linked with the standard Intel IPP dynamic libraries:

Application linked with custom DLL	Application linked with Intel IPP dynamic libraries
ipp_test_app.exe (for Windows*) or ipp_test_app (for Linux* OS and macOS*) ipp_custom_{dll so}.{dll so dylib}	ipp_test_app.exe (for Windows*) or ipp_test_app (for Linux* OS and macOS*) ippi.{dll so dylib} ippig9.{dll so dylib} ippih9.{dll so dylib} ippip8.{dll so dylib} ippipx.{dll so dylib} ippis8.{dll so dylib} ippiw7.{dll so dylib} ippes.{dll so dylib} ippsg9.{dll so dylib} ippsh9.{dll so dylib} ippsp8.{dll so dylib} ippspx.{dll so dylib} ippss8.{dll so dylib} ippsw7.{dll so dylib} ippcore.{dll so dylib}
Package size: 0.1 Mb	Package size: 121.5 Mb

- **Smooth transition to a higher version of Intel IPP/Intel IPP Cryptography.** You can easily build the same custom dynamic library from a higher version of Intel IPP/Intel IPP Cryptography and substitute the libraries in your application without relinking.

NOTE The current Python* version of the Intel IPP Custom Library Tool supports the host-host configuration only, the host-target configuration is currently not supported.

System Requirements for Custom Library Tool

Recommended hardware:

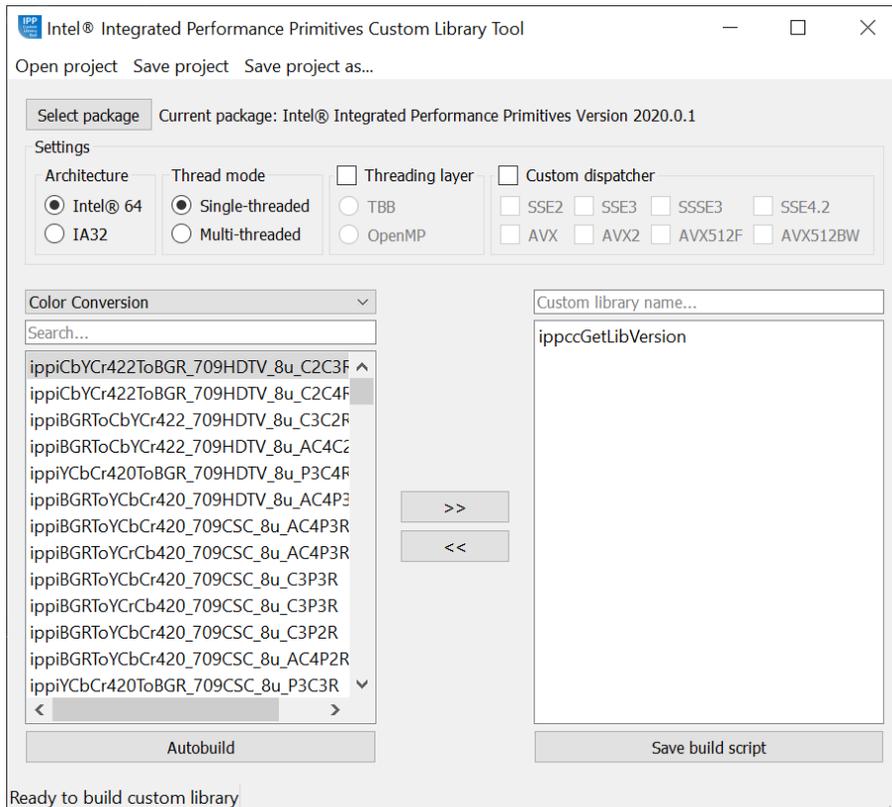
- System based on the 2nd Generation Intel® Core i3, i5, i7 or higher processor

Software requirements:

- Visual Studio* 2013 (or higher) Redistributable Packages
- Python* 3.7
- PyQt5 (required only for the GUI version of the Intel IPP Custom Library Tool)

Operation Modes

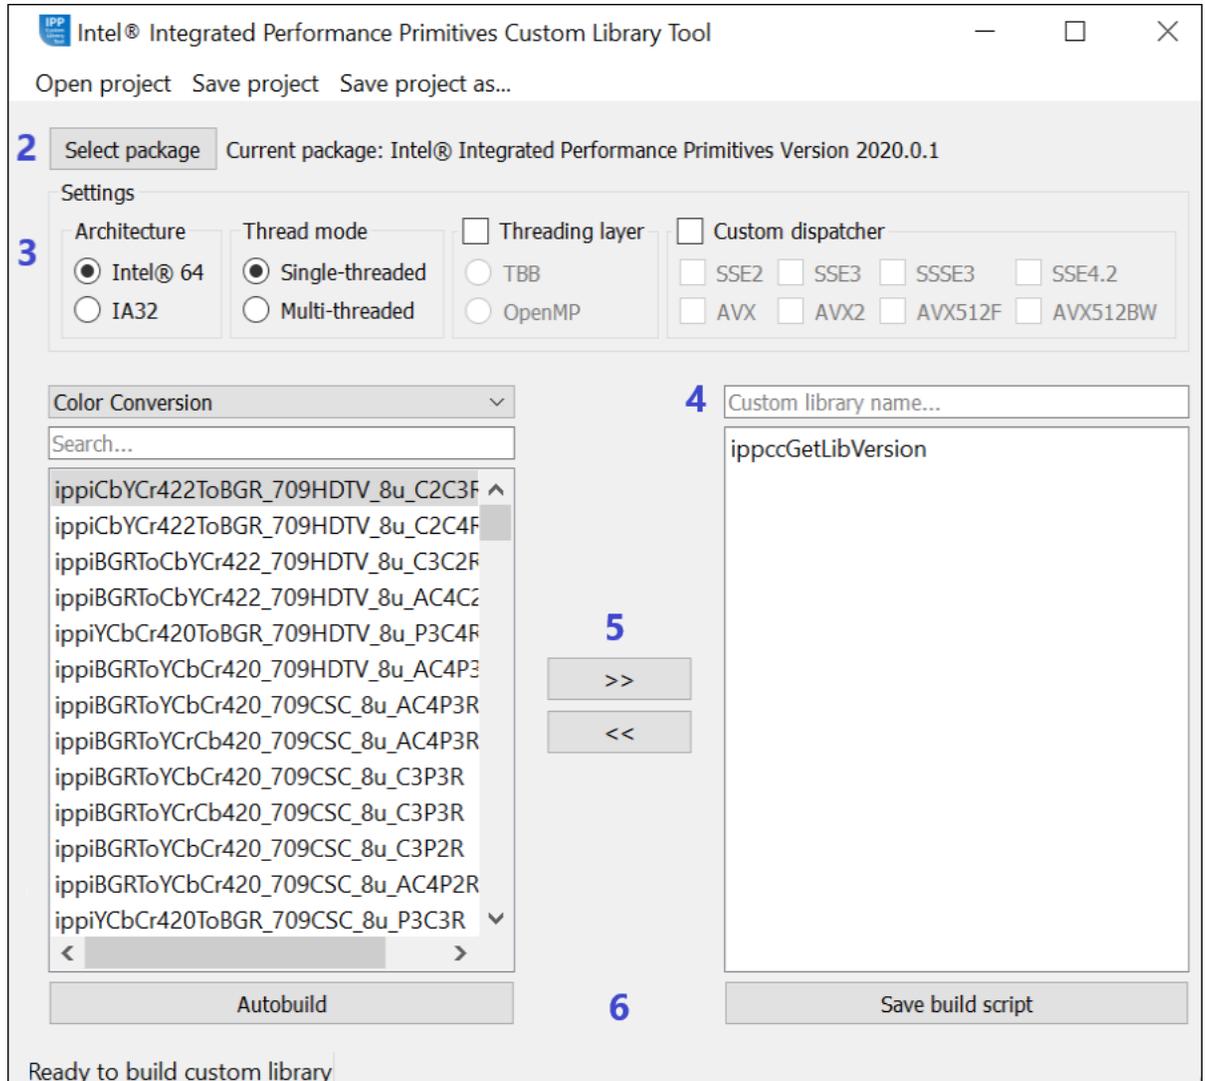
You can choose one of two tool operation modes, as shown at the screen shot below:



- **Auto build.** The tool automatically sets the environment and builds a dynamic library.
- **Save script.** The tool generates and saves a custom build script.

Building a Custom DLL with Custom Library Tool

Follow the steps below to build a custom dynamic library using the Intel IPP Custom Library Tool:



1. Run `python main.py` to launch the GUI version of the tool.
2. Select the Intel IPP or Intel IPP Cryptography package (optional). If you run the tool inside the Intel IPP or Intel IPP Cryptography package, the current one will be used as default. Otherwise, you need to provide the path to the package.
3. Configure your custom library.
4. Set the library name.
5. Select functions from the list. You can build a dynamic library containing Intel IPP or Intel IPP Cryptography functionality, but not both. If you need to add threaded functions to the custom list, select **Threading layer** checkbox to show the list of threaded functions.
6. Build the library automatically (if available) or save a build script.

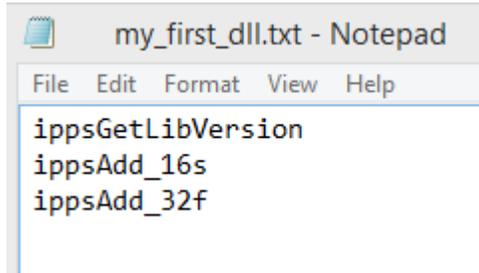
NOTE

You can save the configuration and the list of custom functions as a project by clicking **Save project** or **Save project as....** The project is saved as a file with the `.cltproj` extension. Then you can open this project by clicking **Open project** button.

Using Console Version of Custom Library Tool

Follow the steps below to build a custom dynamic library using console version of the Custom Library Tool:

1. Define a list of Intel IPP functions that the Intel IPP Custom Library Tools should export to your custom dynamic library. See the example text file below:



```

ippsGetLibVersion
ippsAdd_16s
ippsAdd_32f

```

2. Run `python main.py` with the following parameters:

<code>-c, --console</code>	Launches the console version of the tool (the GUI version is used by default).
<code>-g, --generate</code>	Enables the script generation mode (the build mode is used by default).
<code>-n <name>, --name <name></code>	Output library name.
<code>-p <path>, --path <path></code>	Path to the output directory.
<code>-root <root_path></code>	Path to Intel IPP or Intel IPP Cryptography package root directory
<code>-f <function>, --function <function></code>	Name of a function to be included into your custom dynamic library.
<code>-ff <functions_file>, --functions_file <functions_file></code>	Path to a file with a list of functions to be included into your final dynamic library (the <code>-f</code> or <code>--function</code> flag can be used to add functions on the command line).
<code>-arch={ia32 intel64}</code>	Enables all actions for the IA-32 or the Intel® 64 architecture (Intel® 64 architecture is used by default).
<code>-mt, --multi-threaded</code>	Enables multi-threaded libraries (single-threaded libraries are used by default).
<code>-tl={tbb openmp}</code>	Sets Intel TBB or OpenMP* as the threading layer.
<code>-d , --custom_dispatcher <cpu_set></code>	Sets the exact list of CPUs that must be supported by custom dynamic library and generates a C-file with the custom dispatcher.
<code>-h, --help</code>	Prints command help.

For example:

```

# Generate build scripts in console mode
# with the output dynamic library name "my_custom_dll.dll"
# with functions defined in the "functions.txt" file
# optimized only for processors with

```

```
# Intel® Advanced Vector Extensions 512 (Intel® AVX-512)
# using multi-threaded IA-32 Intel IPP libraries

python main.py -c -g
-n my_custom_dll
-p "C:\my_project"
-ff "C:\my_project\functions.txt"
-d avx512bw
-arch=ia32 -mt
```

Using Integration Wrappers for Intel® Integrated Performance Primitives

7

Intel® Integrated Performance Primitives (Intel® IPP) Integration Wrappers aggregate Intel IPP functionality in easy-to-use functions and help to reduce effort required to integrate Intel IPP into your code.

Integration Wrappers consist of C and C++ interfaces:

- **C interface** aggregates Intel IPP functions of similar functionality with various data types and channels into one function. Initialization steps required by several Intel IPP functions are implemented in one initialization function for each functionality. To reduce the size of your code and save time required for integration, the wrappers handle all memory management and Intel IPP function selection routines.
- **C++ interface** wraps around the C interface to provide default parameters, easily initialized objects as parameters, exception handling, and objects for complex Intel IPP functions with automatic memory management for specification structures.

In general, Integration Wrappers are designed to improve user experience with threading of Intel IPP functions and tiling.

Integration Wrappers are provided as a separate download. For more information about the main concepts, usage, and implementation details, refer to the *Developer Guide and Reference for Intel IPP Integration Wrappers* document available with the Integration Wrappers package.

Programming Considerations

Core and Support Functions

There are several general purpose functions that simplify using the library and report information on how it is working:

- `Init/GetCpuFeatures/ SetCpuFeatures/GetEnabledCpuFeatures`
- `GetStatusString`
- `GetLibVersion`
- `Malloc/Free`

Init/GetCpuFeatures/ SetCpuFeatures/GetEnabledCpuFeatures

The `ippInit` function detects the processor type and sets the dispatcher to use the processor-specific code of the Intel® IPP library corresponding to the instruction set capabilities available. If your application does not call the `ippInit` function, initialization of the library to the available instruction set capabilities is performed automatically with the first call of any Intel IPP function from the domain different from `ippCore`.

In some cases like debugging and performance analysis, you may want to get the data on the difference between various processor-specific codes on the same machine. Use the `ippSetCpuFeatures` function for this. This function sets the dispatcher to use the processor-specific code according to the specified set of CPU features. You can obtain features supported by CPU using `ippGetCpuFeatures` and obtain features supported by the currently dispatched Intel IPP code using `ippGetEnabledCpuFeatures`. If you need to enable support of some CPU features without querying the system (without `CPUID` instruction call), you must set the `ippCPUID_NOCHECK` bit for `ippSetCpuFeatures`, otherwise, only supported by the current CPU features are set.

The `ippInit`, `ippGetCpuFeatures`, `ippGetEnabledCpuFeatures`, and `ippSetCpuFeatures` functions are a part of the `ippCore` library.

GetStatusString

The `ippGetStatusString` function decodes the numeric status return value of Intel® IPP functions and converts them to a human readable text:

```
status= ippInit();
if( status != ippStsNoErr ) {
    printf("IppInit() Error:\n");
    printf("%s\n", ippGetStatusString(status) );
    return -1;
}
```

The `ippGetStatusString` function is a part of the `ippCore` library.

GetLibVersion

Each domain has its own `GetLibVersion` function that returns information about the library layer in use from the dispatcher. The code snippet below demonstrates the usage of the `ippiGetLibVersion` from the image processing domain:

```
const IppLibraryVersion* lib = ippiGetLibVersion();
printf("%s %s %d.%d.%d.%d\n", lib->Name, lib->Version,
lib->major, lib->minor, lib->majorBuild, lib->build);
```

Use this function in combination with `ippInitCpu` to compare the output of different implementations on the same machine.

Malloc/Free

Intel IPP functions provide better performance if they process data with aligned pointers. Intel IPP provides the following functions to ensure that data is aligned appropriately - 16-byte for CPU that does not support Intel® Advanced Vector Extensions (Intel® AVX) instruction set, 32-byte for Intel AVX and Intel® Advanced Vector Extensions 2 (Intel® AVX2), and 64-byte for Intel® Many Integrated Core instructions.

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

The `ippMalloc` function provides appropriately aligned buffer, and the `ippFree` function frees it.

The signal and image processing libraries provide `ippsMalloc` and `ippiMalloc` functions, respectively, to allocate appropriately aligned buffer that can be freed by the `ippsFree` and `ippiFree` functions.

NOTE

- When using buffers allocated with routines different from Intel IPP, you may get better performance if the starting address is aligned. If the buffer is created without alignment, use the `ippAlignPtr` function.
-

For more information about the Intel IPP functions see the *Intel® Integrated Performance Primitives for Intel® Architecture Developer Reference* available in Intel® Software Documentation Library.

See Also

Cache Optimizations

Intel® Software Documentation Library

Channel and Planar Image Data Layouts

Intel® IPP functions operate on two fundamental data layouts: channel and planar.

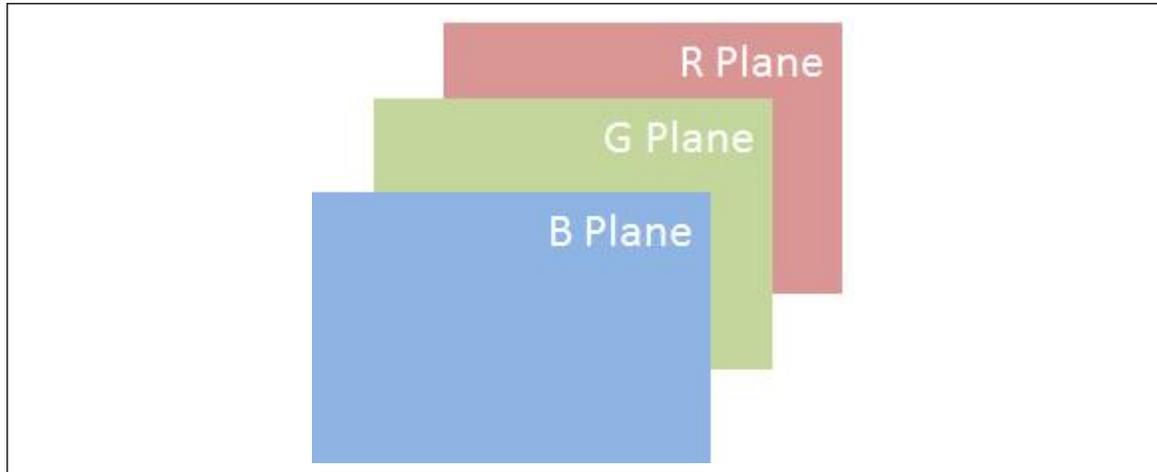
In channel format, all values share the same buffer and all values for the same pixel position are interleaved together. Functions working with channel data have a `_Cn` descriptor, where `n` can take one of the following values: 1, 2, 3, or 4. The figure below shows 24 bit per pixel RGB data, which is represented as `_C3`.

RGB data in `_c3` layout

RGB									
RGB									
RGB									
RGB									
RGB									
RGB									
RGB									
RGB									

For planar format, there is one value per pixel but potentially several related planes. Functions working with planar data have a `_Pn` descriptor, where n can take one of the following values: 1, 2, 3, or 4. The figure below shows 24 bit per pixel RGB data represented as `_P3`.

RGB data in `_P3` layout



NOTE

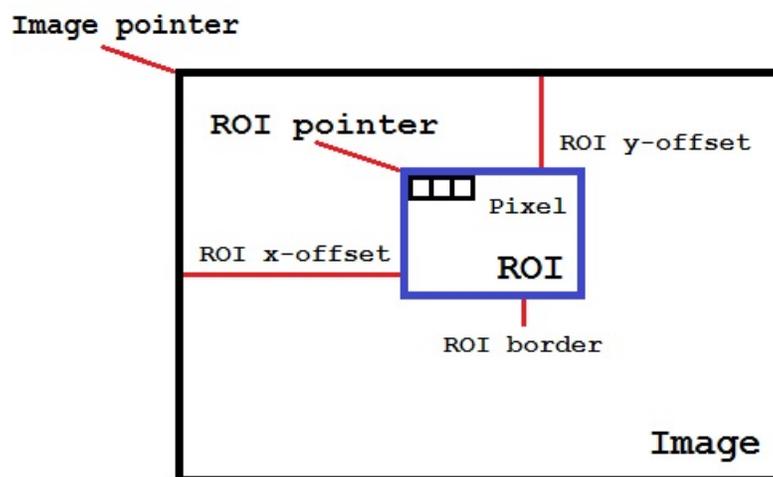
For many video and image processing formats planes may have different sizes.

Regions of Interest

Many Intel® IPP image processing functions operate with a region of interest (ROI). These functions include an R descriptor in their names.

A ROI can be the full image or a subset. This can simplify thread or cache blocking.

Many functions sample a neighborhood and cannot provide values for an entire image. In this case a ROI must be defined for the subset of the destination image that can be computed.



Managing Memory Allocations

In Intel® Integrated Performance Primitives (Intel® IPP) functions, the areas in memory allocated for the source and destination data must not overlap, except for functions that have the descriptor `I` in their name. Only the functions that have the descriptor `I` (see [Descriptors](#)) in their name can have the same area in memory allocated for both the source and destination data. Intel IPP does not guarantee correct behavior and results for not-in-place functions that are used in in-place mode.

Depending on the implementation layer and the specific operation parameters, some Intel IPP functions need varying amounts of memory for internal structures and working buffers. To address this, follow the steps below:

1. Compute the size of the required buffer using the `<function base name>GetSize` function (some functions have `GetBufSize` or `GetBufferSize` in their name instead of `GetSize`).
2. Set up any buffers needed for initialization. For more information, see the section [Setting up Buffers](#) below.
3. Initialize the specification or state structure for the operation using `<function base name>Init` function. For more information about the specification and state structures, see the section [Specification and State Structures](#) below.
4. Free the buffers need for initialization only (the ones you set up in step 2).
5. Set up working buffers for the main operation. For more information, see the section [Setting up Buffers](#) below.
6. Do the main operation.
7. Free the specification or state buffers that you set up in step 3 and the working buffers that you set up in step 5.

If you use several Intel IPP functions with the `pBuffer` parameter (external memory buffer), for better efficiency and performance it is recommended to call all `<function base name>GetSize` functions in one single location within your application and allocate only one buffer that has the largest size. This approach ensures optimal use of system memory and all cache levels.

Setting up Buffers

In this document, "setting up a buffer" refers to allocating the required amount of memory and providing a pointer to this memory to the Intel IPP function you are calling. For better performance, you should allocate aligned memory buffers, where the alignment factor depends on the architecture and should be at least 16 bytes for Intel® Streaming SIMD Extensions, 32 bytes for Intel® Advanced Vector Extensions, and 64 bytes for Intel® Advanced Vector Extensions 512 Foundation instruction sets.

To set up aligned memory buffers, it is recommended to use the `ipp<domain letter>Malloc_<IPP data type>` functions; these functions always provide memory buffers with the required alignment.

NOTE Intel IPP functions do not allocate any memory internally. You must manually allocate and free previously allocated memory, that is required for your Intel IPP functions at the application level. `ipp<domain letter>Malloc_<IPP data type>` and `ipp<domain letter>Free` functions allocate and free a memory block aligned to 64-byte boundary for elements of different data types. Not aligned memory allocation could cause not reproducible performance and precision results.

Specification and State Structures

Specification, or spec, structures are `const`; an instance of a specification structure does not change between Intel IPP function calls. Therefore, you can use one instance of a specification structure simultaneously in different application threads for the same operation.

State structures are not `const`; they always contain the state of an intermediate computation stage of an Intel IPP function. Therefore, you can use a single instance of a state structure only for consecutive operations. In the case of a threaded application, each thread must have its own instance of the state structure.

Product and Performance Information

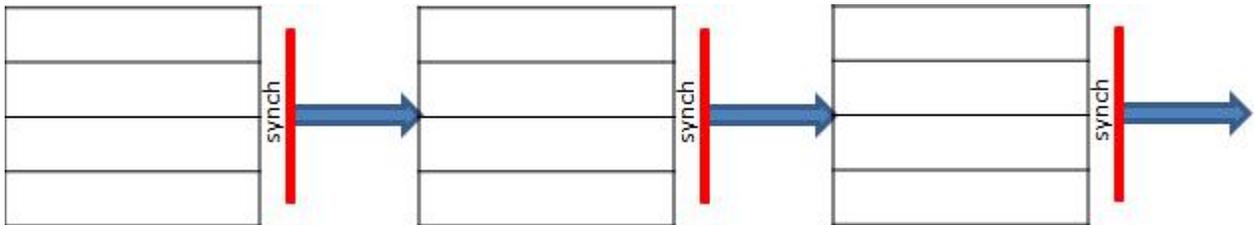
Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Notice revision #20201201

Cache Optimizations

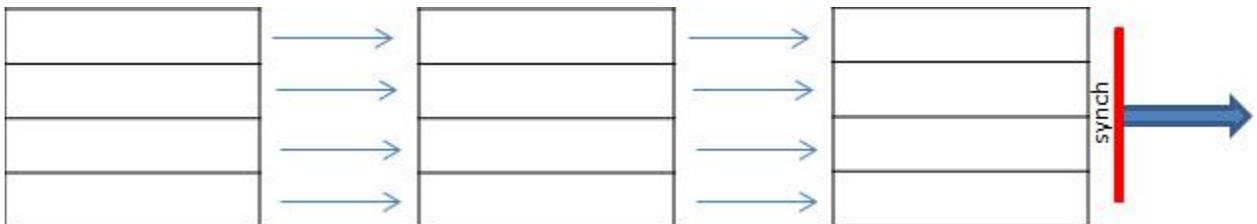
To get better performance, work should be grouped to take advantage of locality in the lowest/fastest level of cache possible. This is the same for threading or cache blocking optimizations.

For example, when operations on each pixels in an image processing pipeline are independent, the entire image is processed before moving to the next step. This may cause many inefficiencies, as shown in a figure below.



In this case cache may contain wrong data, requiring re-reading from memory. If threading is used, the number of synchronization point/barriers is more than the algorithm requires.

You can get better performance after combining steps on local data, as shown in a figure below. In this case each thread or cache-blocking iteration operates with ROIs, not full image.



NOTE

It is recommended to subdivide work into smaller regions considering cache sizes, especially for very large images/buffers.

Programming with Intel® Integrated Performance Primitives in the Microsoft* Visual Studio* IDE

9

This section provides instructions on how to configure your Microsoft* Visual Studio* IDE to link with the Intel® IPP, explains how to access Intel IPP documentation and use IntelliSense* Sense features.

Configuring the Microsoft* Visual Studio* IDE to Link with Intel® IPP

Steps for configuring Microsoft Visual C/C++* development system for linking with Intel® Integrated Performance Primitives (Intel® IPP) depend on whether you installed the C++ Integration(s) in Microsoft Visual Studio* component:

- If you installed the integration component, see [Automatically Linking Your Microsoft* Visual Studio* Project with Intel IPP](#)
- If you did not install the integration component or need more control over Intel IPP libraries to link, you can configure the Microsoft Visual Studio* by performing the following steps. Though some versions of the Visual Studio* development system may vary slightly in the menu items mentioned below, the fundamental configuring steps are applicable to all these versions.
 1. In Solution Explorer, right-click your project and click **Properties**.
 2. Select **Configuration Properties>VC++ Directories** and set the following from the **Select directories for** drop down menu:
 - **Include Files** menu item, and then type in the directory for the Intel IPP include files (default is `<ipp directory>\include`)
 - **Library Files** menu item, and then type in the directory for the Intel IPP library files (default is `<ipp directory>\lib`)
 - **Executable Files** menu item, and then type in the directory for the Intel IPP executable files (default is `<install_dir>\redist\)`

Using the IntelliSense* Features

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

NOTE

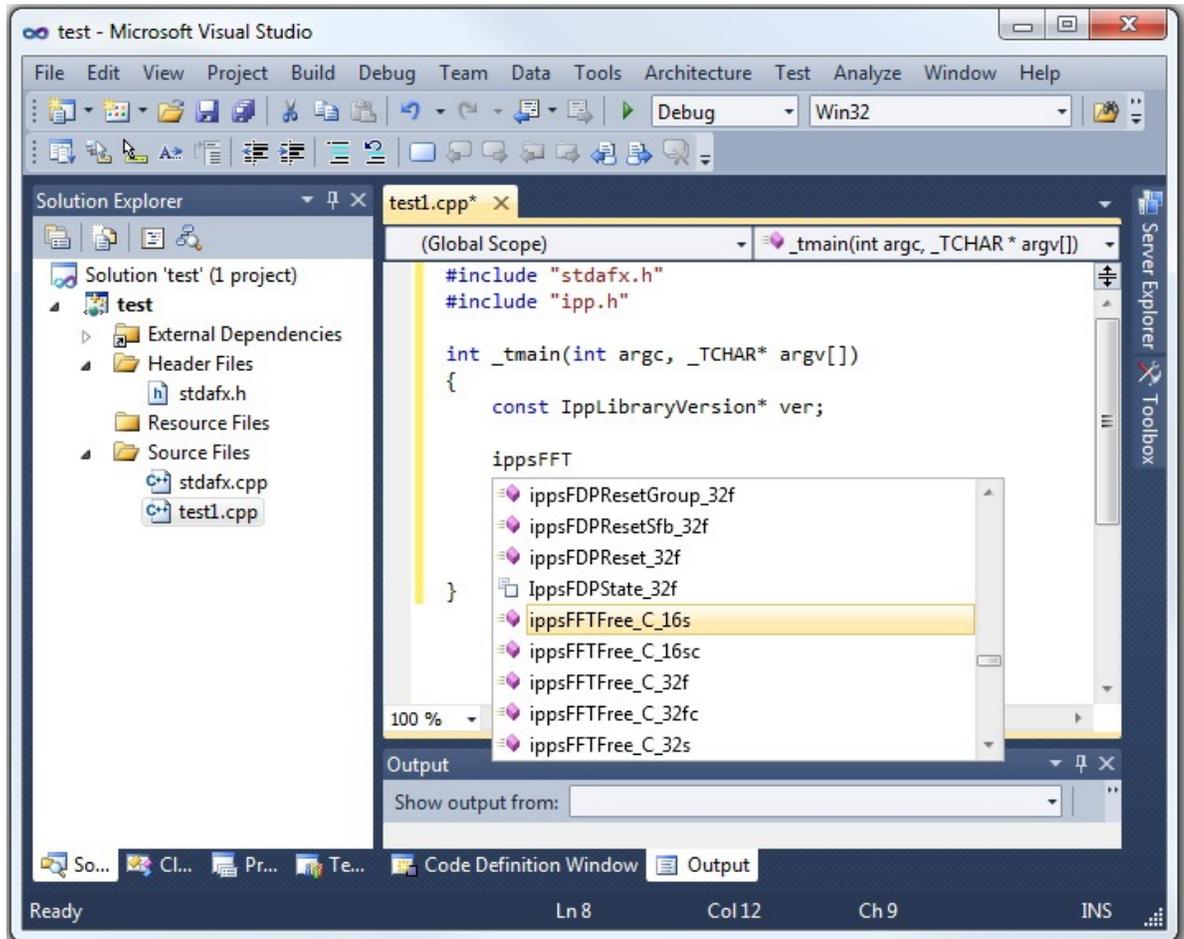
Both features require header files. Therefore, to benefit from IntelliSense, make sure the path to the include files is specified in the Visual Studio solution settings. On how to do this, see [Configuring the Microsoft Visual Studio* IDE to Link with Intel® IPP](#).

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, `ippsFFT`).
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.

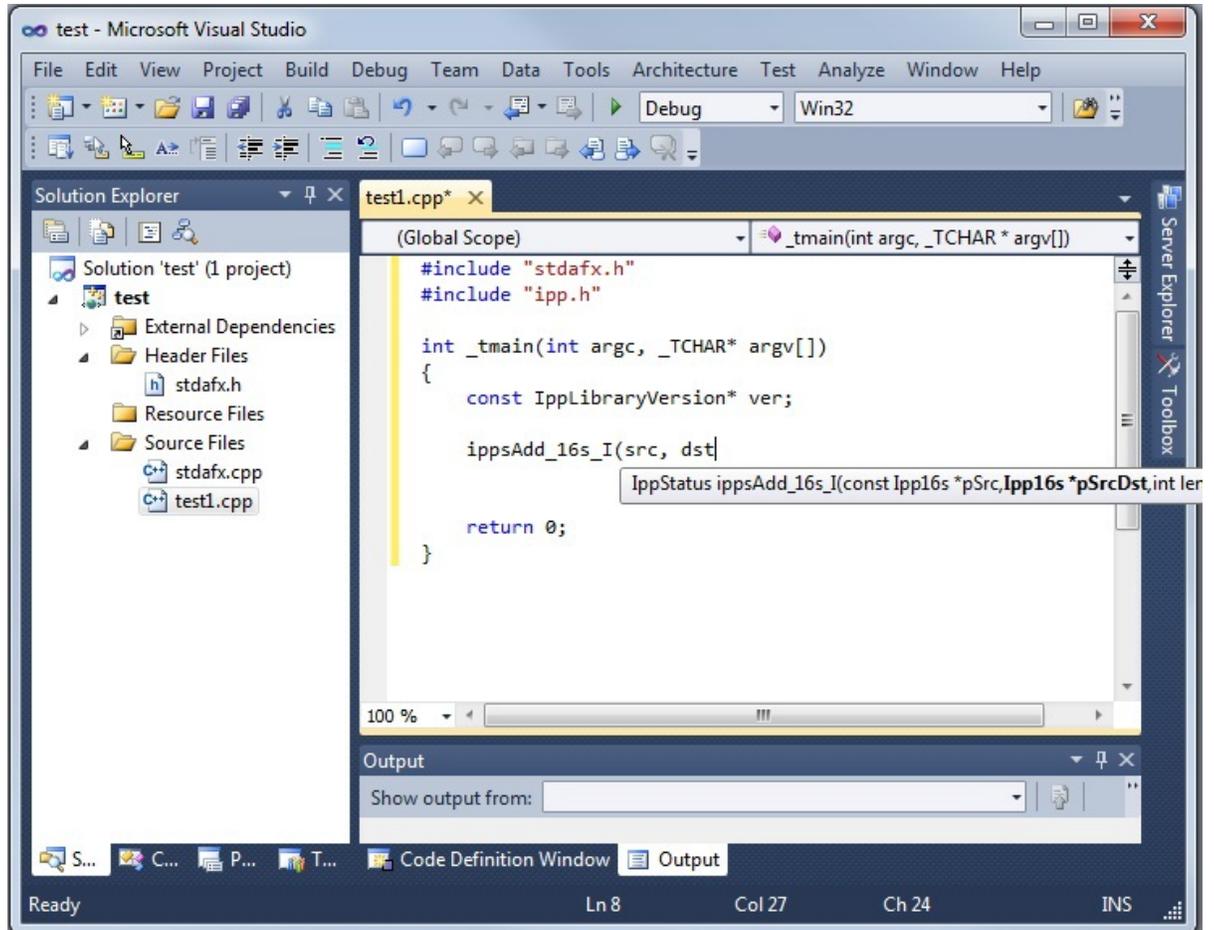
Parameter Info

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

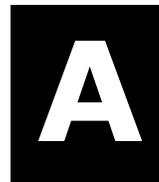
A tooltip appears with the function API prototype, and the current parameter in the API prototype is highlighted - see the figure below.



See Also

Configuring the Microsoft Visual Studio* IDE to Link with Intel® IPP

Appendix: Intel® IPP Threading and OpenMP* Support



All Intel® Integrated Performance Primitives functions are thread-safe. They support multithreading in both dynamic and static libraries and can be used in multi-threaded applications. However, if an application has its own threading model or if other threaded applications are expected to run at the same time on the system, it is strongly recommended to use non-threaded/single-threaded libraries.

Product and Performance Information

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Notice revision #20201201

Using a Shared L2 Cache

Several 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`, and `Sqrt`) achieve the maximum performance if both two threads are executed on the same die. In this case, the 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
```

Product and Performance Information

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Notice revision #20201201

Avoiding Nested Parallelization

Nested parallelization may occur if you use a threaded Intel IPP function in a multithreaded application. Nested parallelization may cause performance degradation because of thread oversubscription.

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 multi-threading in the threaded Intel IPP function to avoid this issue.

Disabling Multi-threading (Recommended)

The best option to disable multi-threading is to link your application with the Intel® IPP single-threaded (non-threaded) libraries included in the default package and discontinue use of the separately downloaded multi-threaded versions.

You may also call the `ippSetNumThreads` function with parameter 1, but this method may still incur some OpenMP* overhead.

Product and Performance Information

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Notice revision #20201201