



Intel[®] High Level Synthesis Compiler

Reference Manual

Updated for Intel[®] Quartus[®] Prime Design Suite: **18.1.1**



MNL-1083 | 2019.01.03

Latest document on the web: [PDF](#) | [HTML](#)



Contents

1. Intel® HLS Compiler Reference Manual.....	4
2. Compiler.....	5
2.1. Intel HLS Compiler Command Options.....	5
2.2. Compiler Interoperability.....	8
3. C Language and Library Support.....	10
3.1. Supported C and C++ Subset for Component Synthesis.....	10
3.2. C and C++ Libraries.....	10
3.3. Compiler-Defined Preprocessor Macros.....	12
3.4. Arbitrary Precision Math Support.....	13
3.4.1. Declaring <code>ac_int</code> Datatypes in Your Component.....	14
3.4.2. Debugging Your Use of the <code>ac_int</code> Datatype.....	14
3.4.3. Declaring <code>ac_fixed</code> Datatypes in Your Component.....	15
3.4.4. AC Datatypes and Native Compilers.....	16
4. Component Interfaces.....	17
4.1. Component Invocation Interface.....	17
4.1.1. Scalar Parameters.....	18
4.1.2. Pointer and Reference Parameters.....	18
4.1.3. Interface Definition Example: Component with Both Scalar and Pointer Arguments.....	18
4.2. Avalon Streaming Interfaces.....	19
4.3. Avalon Memory-Mapped Master Interfaces.....	24
4.3.1. Memory-Mapped Master Testbench Constructor.....	26
4.3.2. Implicit and Explicit Examples of Describing a Memory Interface.....	27
4.4. Slave Interfaces.....	28
4.4.1. Control and Status Register (CSR) Slave.....	28
4.4.2. Slave Memories.....	31
4.5. Component Invocation Interface Arguments.....	32
4.6. Unstable and Stable Component Arguments.....	33
4.7. Global Variables.....	35
4.8. Structs in Component Interfaces.....	35
4.9. Reset Behavior.....	35
5. Local Variables in Components (Memory Attributes).....	37
5.1. Static Variables.....	39
6. Loops in Components.....	41
6.1. Loop Initiation Interval (<code>ii</code> Pragma).....	43
6.2. Loop-Carried Dependencies (<code>ivdep</code> Pragma).....	44
6.3. Loop Coalescing (<code>loop_coalesce</code> Pragma).....	46
6.4. Loop Unrolling (<code>unroll</code> Pragma).....	47
6.5. Loop Concurrency (<code>max_concurrency</code> Pragma).....	47
7. Component Concurrency.....	49
7.1. Serial Equivalence within a Memory Space or I/O.....	49
7.2. Concurrency Control (<code>hls_max_concurrency</code> Attribute).....	49





8. Intel HLS Compiler Libraries.....	51
8.1. Random Number Generator Library.....	51
8.2. Matrix Multiplication Library.....	52
9. Document Revision History of the Intel HLS Compiler Reference Manual.....	54
A. Intel High Level Synthesis Compiler Quick Reference.....	59
B. Supported Math Functions.....	75
B.1. Math Functions Provided by the <code>math.h</code> Header File	75
B.2. Math Functions Provided by the <code>extendedmath.h</code> Header File.....	79
B.3. Math Functions Provided by the <code>ac_fixed_math.h</code> Header File.....	79

1. Intel® HLS Compiler Reference Manual


The *Intel® HLS Compiler Reference Manual* provides reference information about the features supported by the Intel HLS Compiler. The Intel HLS Compiler is sometimes referred to as the i++ compiler, reflecting the name of the compiler command.


The features and devices supported by the Intel HLS Compiler depend on what edition of Intel Quartus® Prime you have. The following icons indicate content in this publication that applies only to the Intel HLS Compiler provided with a certain edition of Intel Quartus Prime:

-  Indicates that a feature or content applies only to the Intel HLS Compiler provided with Intel Quartus Prime Pro Edition.
-  Indicates that a feature or content applies only to the Intel HLS Compiler provided with Intel Quartus Prime Standard Edition.

In this publication, *<quartus_installdir>* refers to the location where you installed Intel Quartus Prime Design Suite. The Intel High Level Synthesis (HLS) Compiler is installed as part of your Intel Quartus Prime Design Suite installation.

The default Intel Quartus Prime Design Suite installation location depends on your operating system and your Intel Quartus Prime edition:

- 

<i>Windows</i>	<code>C:\intelFPGA_pro\18.1</code>
<i>Linux</i>	<code>/home/<username>/intelFPGA_pro/18.1</code>
- 

<i>Windows</i>	<code>C:\intelFPGA_standard\18.1</code>
<i>Linux</i>	<code>/home/<username>/intelFPGA_standard/18.1</code>

2. Compiler

2.1. Intel HLS Compiler Command Options

Use the Intel HLS Compiler command options to customize how the compiler performs general functions, customize file linking, or customize compilation.

Table 1. General Command Options

These i++ command options perform general compiler functions.






Command Option	Description
--debug-log	Instructs the compiler to generate a log file that contains diagnostic information. By default, the <code>debug.log</code> file is in the <code>a.prj</code> subdirectory within your current working directory. If you also include the <code>-o <result></code> command option, the <code>debug.log</code> file will be in the <code><result>.prj</code> subdirectory. If your compilation fails, the <code>debug.log</code> file is generated whether you set this option or not.
-h or --help	Instructs the compiler to list all the command options and their descriptions on screen.
-o <result>	Instructs the compiler to place its output into the <code><result></code> executable and the <code><result>.prj</code> directory. If you do not specify the <code>-o <result></code> option, the compiler outputs an <code>a.out</code> file for Linux and an <code>a.exe</code> file for Windows. Use the <code>-o <result></code> command option to specify the name of the compiler output. Example command: <code>i++ -o hlsoutput multiplier.c</code> Invoking this example command creates an <code>hlsoutput</code> executable for Linux and an <code>hlsoutput.exe</code> for Windows in your working directory.
-v	Verbose mode that instructs the compiler to display messages describing the progress of the compilation. Example command: <code>i++ -v hls/multiplier/multiplier.c</code> , where <code>multiplier.c</code> is the input file.
--version	Instructs the compiler to display its version information on screen. Command: <code>i++ --version</code>

Table 2. Command Options that Customize Compilation

These i++ command options perform compiler functions that impact the translation from source file to object file.





Option	Description
-c	Instructs the compiler to preprocess, parse, and generate object files (<code>.o/.obj</code>) in the current working directory. The linking stage is omitted. Example command: <code>i++ -march="Arria 10" -c multiplier.c</code>


continued...

Option	Description
	Invoking this example command creates a multiplier.o file and sets the name of the <result>.prj directory to multiplier.prj. When you later compile the .o file, the -o option affects only the name of the executable file. The name of the <result>.prj directory remains unchanged from when the directory name was set by i++ -c command invocation.
--component <components>	Allows you to specify a comma-separated list of function names that you want to the compiler to synthesize to RTL. Example command: i++ counter.cpp --component count To use this option, your component must be configured with C-linkage using the extern "C" specification. For example: <pre>extern "C" int myComponent(int a, int b)</pre> Using the component function attribute is preferred over using the --component command option to indicate functions that you want the compiler to synthesize.
-D<macro> [=<val>]	Allows you to pass a macro definition (<macro>) and its value (<val>) to the compiler. If you do not specify a value for <val>, its default value will be 1.
-g	Generate debug information (default).
-g0	Do not generate debug information.
 -gcc-toolchain	Specifies the path to a GCC installation that you want to use for compilation. This path should be the absolute path to the directory that contains the GCC lib, bin, and include folders. You should not need to use this if you configured your system as described in the Getting Started Guide.
-I<dir>	Adds a directory (<dir>) to the end of the include path list.
-march=[x86-64 <FPGA_family> <FPGA_part_number>	Instructs the compiler to compile the component to the specified architecture or FPGA family. The -march compiler option can take one of the following values: x86-64 Instructs the compiler to compile the code for an emulator flow. "<FPGA_family>" Instructs the compiler to compile the code for a target FPGA device family. The <FPGA_family> value can be any of the following device families: <ul style="list-style-type: none"> •  Arria10 or "Arria 10" •  Cyclone10GX or "Cyclone 10 GX" •  Stratix10 or "Stratix 10" •  ArriaV or "Arria V"

continued...



Option	Description
	<ul style="list-style-type: none"> •  CycloneV or "Cyclone V" •  MAX10 or "MAX 10"⁽¹⁾ •  StratixV or "Stratix V" <p>Quotation marks are required only if you specify a FPGA family name specifier that contains spaces</p> <p><code><FPGA_part_number></code> Instructs the compiler to compile the code for a target device. The compiler determines the FPGA device family from the FPGA part number that you specify here.</p> <p>If you do not specify this option, <code>-march=x86-64</code> is assumed.</p> <p>If the parameter value that you specify contains spaces, surround the parameter value in quotation marks.</p>
 <code>--promote-integers</code>	<p>Instructs the compiler to use additional FPGA resources to mimic g++ integer promotion. Integer promotion occurs when all integer operations are carried out in 32 bits even if the largest operand is smaller than 32 bits.</p> <p>The default behavior is to carry out integer operations in the size of the largest operand. Refer to the <code><path to i++ installation>/examples/tutorials/best_practices/integer_promotion</code> design example for usage information on the <code>--promote-integers</code> command option.</p> <p>In Pro Edition, the compiler always promotes integers for standard types. Use the <code>ac_int</code> datatypes if you want smaller (or larger) datatypes.</p>
<code>--quartus-compile</code>	<p>Compiles your HDL file with the Intel Quartus Prime compiler.</p> <p>Example command: <code>i++ --quartus-compile <input_files> -march="Arria 10"</code></p> <p>When you specify this options, the Intel Quartus Prime compiler is run after the HDL is generated. The compiled Intel Quartus Prime project is put in the <code><result>.prj/quartus</code> directory and a summary of the FPGA resource consumption and maximum clock frequency is added to the high level design reports in the <code><result>.prj/reports</code> directory.</p>
<code>--simulator</code> <code><simulator_name></code>	<p>Specifies the simulator you are using to perform verification.</p> <p>This command option can take the following values for <code><simulator_name></code>:</p> <ul style="list-style-type: none"> • <code>modelsim</code> • <code>none</code> <p>If you do not specify this option, <code>--simulator modelsim</code> is assumed.</p> <p>Important: The <code>--simulator</code> command option only works in conjunction with the <code>-march</code> command option.</p>

(1)  If you develop your component IP for Intel MAX[®] 10 devices and you want to integrate your component IP into a system that you are developing in Intel Quartus Prime, ensure that the Intel Quartus Prime settings file (.qsf) for your system contains one of the following lines:

```

-- set_global_assignment -name INTERNAL_FLASH_UPDATE_MODE "SINGLE IMAGE WITH ERAM"
-- set_global_assignment -name INTERNAL_FLASH_UPDATE_MODE "SINGLE COMP IMAGE WITH ERAM"

```

When you compile the component IP for an Intel MAX 10 devices with Intel HLS Compiler, the generated Intel Quartus Prime example project contains all of the required QSF settings for your component. However, the Intel Quartus Prime project for the system into which you integrate your component might not have the required QSF setting.

Option	Description
	<p>The <code>--simulator none</code> option instructs the HLS compiler to skip the verification flow and generate RTL for the components without generating the corresponding test bench. If you use this option, the high-level design report (<code>report.html</code>) omits verification statistics such as component reset latency.</p> <p>Example command: <code>i++ -march="<FPGA_family_or_part_number>" --simulator none multiplier.c</code></p>

Table 3. Command Options that Customize File Linking

These HLS command options specify compiler actions that impact the translation of the object file to the binary or RTL component.

Option	Description
<code>--clock <clock_spec></code>	Optimizes the RTL for the specified clock frequency or period.
<code>--fpc</code>	Removes intermediate rounding and conversion whenever possible. To see an example of when and how to use this option, review the tutorial in <code><quartus_installdir>/hls/examples/tutorials/best_practices/floating_point_ops</code> .
<code>--fp-relaxed</code>	Relaxes the order of arithmetic operations. To see an example of how to use this option, review the tutorial in <code><quartus_installdir>/hls/examples/tutorials/best_practices/floating_point_ops</code> .
<code>-ghdl</code>	Logs all signals when running the verification executable. After running the executable, the simulator logs waveforms to the <code>a.prj/verification/vsim.wlf</code> file. For details about the ModelSim* waveform, see Debugging during Verification in <i>Intel High Level Synthesis Compiler User Guide</i> .
<code>-L<dir></code>	(Linux only) Adds a directory (<code><dir></code>) to the end of the search path for the library files.
<code>-l<library></code>	(Linux only) Specifies the library file name when linking the object file to the binary. On Windows
<code>--x86-only</code>	Creates only the testbench executable. The compiler outputs an <code><result></code> file for Linux or a <code><result>.exe</code> file for Windows. The <code><result>.prj</code> directory and its contents are not created.
<code>--fpga-only</code>	Creates only the <code><result>.prj</code> directory and its contents. The testbench executable file (<code><result>/<result>.exe</code>) is not created.

2.2. Compiler Interoperability

The Intel High Level Synthesis Compiler is compatible with x86-64 object code compiled by supported versions of GCC or Microsoft Visual Studio. You can compile your testbench code with GCC or Microsoft Visual Studio, but generating RTL and cosimulation support for your component always requires the Intel HLS Compiler.

To see what versions of GCC and Microsoft Visual Studio the Intel HLS Compiler supports, see "[Intel High Level Synthesis Compiler Prerequisites](#)" in *Intel High Level Synthesis Compiler Getting Started Guide*.

The interoperability between GCC or Microsoft Visual Studio, and the Intel HLS Compiler lets you decouple your testbench development from your component development. Decoupling your testbench development can be useful for situations where you want to iterate your testbench quickly with platform-native compilers (GCC/Microsoft Visual Studio), without having to recompile the RTL generated for your component.



With Microsoft Visual Studio, you can compile only code that does not explicitly use the Avalon[®]-Streaming interface.

To create only your testbench executable with the `i++` command, specify the `--x86-only` option.

You can choose to only generate RTL and cosimulation support for your component by linking the object file or files for your component with the Intel High Level Synthesis Compiler.

To generate only your RTL and cosimulation support for your component, specify the `--fpga-only` option.

3. C Language and Library Support

3.1. Supported C and C++ Subset for Component Synthesis

The Intel HLS Compiler has several synthesis limitations regarding the supported subset of C99 and C++.

The compiler cannot synthesize code for dynamic memory allocation, virtual functions, function pointers, and C++ or C library functions except the supported math functions explicitly mentioned in the appendix of this document. In general, the compiler can synthesize functions that include classes, structs, functions, templates, and pointers.

While some C++ constructs are synthesizable, aim to create a component function in C99 whenever possible.

Important: These synthesis limitations do not apply to testbench code.

3.2. C and C++ Libraries

The Intel High Level Synthesis (HLS) Compiler provides a number of header files to provide FPGA implementations of certain C and C++ functions.

Table 4. Intel High Level Synthesis (HLS) Compiler Header Files

Feature	Description
<code>#include "HLS/hls.h"</code>	Required for component identification and component parameter interfaces.
<code>#include "HLS/math.h"</code>	Includes FPGA-specific definitions for the math functions from the <code>math.h</code> for your operating system. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/best_practices/single_vs_double_precision_math</code>
<code>#include "HLS/extendedmath.h"</code>	Includes additional FPGA-specific definitions of math functions not in <code>math.h</code> . To learn more, review the design: <code><quartus_installdir>/hls/examples/QRD</code>
<code>#include "HLS/ac_int.h"</code>	Intel HLS Compiler version of <code>ac_int</code> header file. Provides arbitrary width integer support. To learn more, review the following tutorials: <ul style="list-style-type: none"> <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_basic_ops</code> <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow</code> <code><quartus_installdir>/hls/examples/tutorials/best_practices/struct_interfaces</code>
<i>continued...</i>	

Intel Corporation. All rights reserved. Intel, the Intel logo, Altera, Arria, Cyclone, Enpirion, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

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



Feature	Description
<code>#include "HLS/ac_fixed.h"</code>	Intel HLS Compiler version of the <code>ac_fixed</code> header file. Provides arbitrary precision fixed point support. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_constructor</code>
<code>#include "HLS/ac_fixed_math.h"</code>	Intel HLS Compiler version of the <code>ac_fixed_math</code> header file. Provides arbitrary precision fixed point math functions. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_math_library</code>
<code>#include "HLS/stdio.h"</code>	Provides <code>printf</code> support for components so that <code>printf</code> statements work in x86 emulations, but are disabled in component when compiling to an FPGA architecture.
<code>#include <iostream></code>	To use the C++ standard output streams (<code>cout</code> and <code>cerr</code>) provided by the standard <code><iostream></code> header, you must guard any standard output statements with the <code>HLS_SYNTHESIS</code> macro. This macro ensures that statements in a component work in x86 emulations but are disabled in the component when compiling to an FPGA architecture.

math.h

To access functions in `math.h` from a component to be synthesized, include the "HLS/math.h" file in your source code. The header ensures that the components call the hardware versions of the math functions.

For more information about supported `math.h` functions, see [Supported Math Functions](#) on page 75.

stdio.h

Synthesized component functions generally do not support C and C++ standard library functions such as FILE pointers.

A component can call `printf` by including the header file `HLS/stdio.h`. This header changes the behavior of `printf` depending on the compilation target:

- For compilation that targets the x86-64 architecture (that is, `-march=x86-64`), the `printf` call behaves as normal.
- For compilation that targets the FPGA architecture (that is, `-march="<FPGA_family_or_part_number>"`), the compiler removes the `printf` call.

If you use `printf` in a component function without first including the `#include "HLS/stdio.h"` line in your code, you get an error message similar to the following error when you compile hardware to the FPGA architecture:

```
$ i++ -march="<FPGA_family_or_part_number>" --component dut test.cpp
Error: HLS gen_qsys FAILED.
See ./a.prj/dut.log for details.
```

You can use C and C++ standard library functions such as `fopen` and `printf` as normal in all testbench functions.

iostream

A component can use C++ standard output streams (`cout` or `cerr`) provided by the standard C++ header but you must guard any `cout` or `cerr` statements with the `HLS_SYNTHESIS` macro. This macro ensures that statements in a component work in x86 emulations (that is, `-march=x86-64`), but are disabled in the component when compiling it to an FPGA architecture (that is, `-march="<FPGA_family_or_part_number>"`). For example:

```
#include "HLS/hls.h"
#include <iostream>

component int debug_component (int a){
#ifdef HLS_SYNTHESIS
    std::cout << "input value: " << a << std::endl;
#endif
    return a;
}
```

If you attempt to use `cout` or `cerr` in a component function without guarding the line in your code with the `HLS_SYNTHESIS` macro, you get an error message similar to the following error when you compile hardware to the FPGA architecture:

```
$ i++ -march="<FPGA_family_or_part_number>" run.cpp
run.cpp:5: Compiler Error: Cannot synthesize std::cout used inside of a
component.
HLS Main Optimizer FAILED.
```

Related Information

Supported Math Functions on page 75

3.3. Compiler-Defined Preprocessor Macros

The Intel HLS Compiler has a built-in macros that you can use to tailor your code to create flow-dependent behaviors.

Table 5. Macro Definition for `__INTELFPGA_COMPILER__`

Tool Invocation	<code>__INTELFPGA_COMPILER__</code>
<code>g++</code>	Undefined
<code>i++ -march=x86-64</code>	1811
<code>i++ -march="<FPGA_family_or_part_number>"</code>	1811

Table 6. Macro Definition for `HLS_SYNTHESIS`

Tool Invocation	<code>HLS_SYNTHESIS</code>	
	Testbench Code	HLS Component Code
<code>g++</code>	Undefined	Undefined
<code>i++ -march=x86-64</code>	Undefined	Undefined
<code>i++ -march="<FPGA_family_or_part_number>"</code>	Undefined	Defined



3.4. Arbitrary Precision Math Support

The Algorithmic C (AC) datatypes are a collection of header files that Mentor Graphics* provides under the Apache license. Intel developed optimized versions of the AC datatypes to allow the Intel HLS Compiler to generate efficient hardware on Intel FPGAs for these datatypes. For more information on Algorithmic C datatypes, refer to *Mentor Graphics Algorithmic C (AC) Datatypes*, which is available as `<quartus_installdir>/hls/include/ref/ac_datatypes_ref.pdf`.

The Intel HLS Compiler supports the following AC datatypes:

Table 7. AC Datatypes Supported by the HLS Compiler

AC Datatype	Intel Header File	Description
ac_int	HLS/ac_int.h	Arbitrary width integer support To learn more, review the following tutorials: <ul style="list-style-type: none"> <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_basic_ops</code> <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow</code> <code><quartus_installdir>/hls/examples/tutorials/best_practices/struct_interfaces</code>
ac_fixed	HLS/ac_fixed.h	Arbitrary precision fixed-point support To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_constructor</code>
	HLS/ac_fixed_math.h	Support for some nonstandard math functions for arbitrary precision fixed-point datatypes To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_math_library</code>

Using the `ac_int` and `ac_fixed` datatypes has the following advantages over using standard C/C++ datatypes in your components:

- You can achieve smaller datapaths and processing elements for various operations in the circuit.
- The datatypes ensure that all operations are carried out in a size guaranteed not to lose any data. However, you can still lose data if you store data into a location where the datatype is too small.

The `ac_int` and `ac_fixed` datatypes have the following limitations:

- Multipliers are limited to generating 512-bit results.
- Dividers are limited to a maximum of 64 bits.
- The Intel header files are not compatible with GCC or MSVC. When you use the Intel header files, you cannot use GCC or MSVC to compile your testbench. Both your component and testbench must be compiled with the Intel HLS Compiler.

The Intel HLS Compiler also supports some nonstandard math functions for the `ac_fixed` datatype when you include the `HLS/ac_fixed_math.h` header file.

Related Information

[AC Datatypes Download page on the Mentor Graphics website](#)

3.4.1. Declaring `ac_int` Datatypes in Your Component

The HLS compiler package includes an `ac_int.h` header file for you to include in your component to use arbitrary precision integers in your component.

1. Include the `ac_int.h` header file in your component in the following manner:

```
#include "HLS/hls.h"
#include "HLS/ac_int.h"
```

2. After you include the header file, declare your `ac_int` variables in one of the following ways:

- Template-based declaration
 - `ac_int<N, true> var_name; //Signed N bit integer`
 - `ac_int<N, false> var_name; //Unsigned N bit integer`
- Predefined types up to 63 bits
 - `intN var_name; //Signed N bit integer`
 - `uintN var_name; //Unsigned N bit integer`

Where *N* is the total length of the integer in bits.

Restriction: If you want to initialize an `ac_int` variable to a value larger than 64 bits, you must initialize the variable with the `ac::init_array` constructor.

3.4.1.1. Important Usage Information on the `ac_int` Datatype

The `ac_int` datatype has a large number of API calls that are documented in the `ac_int` documentation included in the Intel HLS Compiler installation package. For more information on AC datatypes, refer to *Mentor Graphics Algorithmic C (AC) Datatypes*, which is available as `<quartus_installdir>/hls/include/ref/ac_datatypes_ref.pdf`.

The `ac_int` datatype automatically increases the size of the result of the operation to guarantee that the intermediate operations never overflow. However, the HLS compiler automatically truncates or extends the result to the size of the specified destination container, so ensure that your storage variable for the computation is large enough.

The HLS compiler installation package includes a number of examples in the tutorials. Refer to the tutorials in `<quartus_installdir>/hls/example/tutorials/ac_datatypes` for some of the recommended practices.

3.4.2. Debugging Your Use of the `ac_int` Datatype

The "HLS/ac_int.h" header file provides you with tools to help check `ac_int` operations and assignments for overflow in your component when you run an x86 emulation of your component: `DEBUG_AC_INT_WARNING` and `DEBUG_AC_INT_ERROR`.

When you use the `DEBUG_AC_INT_WARNING` and `DEBUG_AC_INT_ERROR` macros, you cannot declare `constexpr ac_int` variables or `constexpr ac_int` arrays.



Table 8. Intel HLS Compiler ac_int Debugging Tools

Feature		Description
<p><i>Macro:</i></p> <pre>#define DEBUG_AC_INT_WARNING If you use this macro, declare it in your code before you declare #include HLS/ac_int.h.</pre> <p><i>Compiler command line option:</i></p> <pre>-D DEBUG_AC_INT_WARNING</pre>	<p>Enables runtime tracking of ac_int datatypes during x86 emulation (the -march=x86-64 option, which the default option, of the i++ command).</p> <p>This tool uses additional resources for tracking the overflow and empty constructors, and emits a warning for each detected overflow. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow</p>	
<p><i>Macro:</i></p> <pre>#define DEBUG_AC_INT_ERROR If you use this macro, declare it in your code before you declare #include HLS/ac_int.h.</pre> <p><i>Compiler command line option:</i></p> <pre>-D DEBUG_AC_INT_ERROR</pre>	<p>Enables runtime tracking of ac_int datatypes during x86 emulation of your component (the -march=x86-64 option, which the default option, of the i++ command).</p> <p>This tool uses additional resources to track the overflow and empty constructors, and emits a message for the first overflow that is detected and then exits the component with an error. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow</p>	

After you use these tools to determine that your component has overflows, run the gdb debugger on your component to run the program again and step through the program to see where the overflows happen.

Review the ac_int_overflow tutorial in <quartus_installdir>/hls/example/tutorials/ac_datatypes to learn more.

3.4.3. Declaring ac_fixed Datatypes in Your Component

The HLS compiler package includes an ac_fixed.h header file for arbitrary precision fixed-point support.

1. Include the ac_fixed.h header file in your component in the following manner:

```
#include "HLS/hls.h"
#include "HLS/ac_fixed.h"
```

2. After you include the header file, declare your ac_fixed variables as follows:

```
– ac_fixed<N, I, true, Q, O> var_name; //Signed fixed-point
number
– ac_fixed<N, I, false, Q, O> var_name; //Unsigned fixed-
point number
```

Where the template attributes are defined as follows:

- N* The total length of the fixed-point number in bits.
- I* The number of bits used to represent the integer value of the fixed-point number.
The difference of $N-I$ determines how many bits represent the fractional part of the fixed-point number.



- Q The quantization mode that determines how to handle values where the generated precision (number of decimal places) exceeds the number of bits available in the variable to represent the fractional part of the number.
For a list of quantization modes and their descriptions, , see "2.1. Quantization and Overflow" in *Mentor Graphics Algorithmic C (AC) Datatypes*, which is available as
`<quartus_installdir>/hls/include/ref/ac_datatypes_ref.pdf.`

- O The overflow mode that determines how to handle values where the generated value has more bits than the number of bits available in the variable.
For a list of overflow modes and their descriptions, , see "2.1. Quantization and Overflow" in *Mentor Graphics Algorithmic C (AC) Datatypes*, which is available as
`<quartus_installdir>/hls/include/ref/ac_datatypes_ref.pdf.`

3.4.4. AC Datatypes and Native Compilers

The reference version of the Mentor Graphics Algorithmic C (AC) datatypes is also provided with the Intel HLS Compiler. Do not use these reference header files in your component if you want to compile your component with an FPGA target.

Use the reference header files for AC datatypes to confirm functional correctness in your component when you are compiling your component with native compilers (g++ or MSVC).

If you use the reference header files and compile your component to an FPGA target, your component can compile successfully but your component QoR will be poor.

All of your code must use the same header files (either the reference header files or the FPGA-optimized header files). For example, your code cannot use the reference header files in your testbench and, at the same time, use the FPGA-optimized header file in your component code.

The following reference header files are provided with the Intel HLS Compiler:

AC Datatype	Reference Header File	Description
ac_int	ref/ac_int.h	Arbitrary width integer support
ac_fixed	ref/ac_fixed.h	Arbitrary precision fixed-point support

4. Component Interfaces

Intel HLS Compiler generates a component interface for integrating your RTL component into a larger system. A component has two basic interface types: the component invocation interface and the parameter interface.

The *component invocation interface* is common to all HLS components and contains the return data (for nonvoid functions) and handshake signals for passing control to the component, and for receiving control back when the component finishes executing.

The *parameter interface* is the protocol you use to transfer data in and out of your component function. The parameter interface for your component is based on the parameters that you define in your component function signature.

4.1. Component Invocation Interface

For each function that you label as a component, the Intel HLS Compiler creates a corresponding RTL module. This RTL module must have top-level ports, or interfaces, that allow your overall system to interact with your HLS component.

By default, the RTL module for a component includes the following interfaces and data:

- A call interface that consists of `start` and `busy` signals
- A return interface that consists of `done` and `stall` signals
- Return data if the component function has a return type that is not `void`

See [Figure 1](#) on page 18 for an example component showing these interfaces.

Your component function parameters generate different RTL depending on their type. For details see the following sections:

- [Scalar Parameters](#) on page 18
- [Pointer and Reference Parameters](#) on page 18

You can also explicitly declare Avalon Streaming interfaces (`stream_in<>` and `stream_out<>` classes) and Avalon Memory-Mapped Master (`mm_master<>` classes) interfaces on component interfaces. For details see the following sections:

- [Avalon Streaming Interfaces](#) on page 19
- [Avalon Memory-Mapped Master Interfaces](#) on page 24

In addition, you can indicate the control signals that correspond to the actions of calling your component by using the component invocation interface arguments. For details, see [.Component Invocation Interface Arguments](#) on page 32.

4.1.1. Scalar Parameters

Each scalar argument in your component results in an input conduit that is associated with the component `start` and `busy` signals.

The inputs are read into the component when the external system pulls the `start` signal high and the component keeps the `busy` signal low.

For an example of how to specify a scalar parameters and how it is read in by a component, see the `a` argument in [Figure 1](#) on page 18 and [Figure 2](#) on page 19.

4.1.2. Pointer and Reference Parameters

Each pointer or reference argument of a component results in an input conduit, associated with the component `start` and `busy` signals, for the address. In addition to this input conduit, all pointers share a single Avalon Memory-Mapped (MM) master interface that the component uses to access system memory.

You can customize these pointer interfaces using the `mm_master<>` class.

Note: Explicitly-declared Avalon Memory-Mapped Master interfaces and Avalon Streaming interfaces are passed by reference.

For details about Avalon (MM) Master interfaces, see [Avalon Memory-Mapped Master Interfaces](#) on page 24.

4.1.3. Interface Definition Example: Component with Both Scalar and Pointer Arguments

The following design example illustrates the interactions between a component's interfaces and signals, and the waveform of the corresponding RTL module.

```
component int dut(int a, int* b, int i) {
    return a*b[i];
}
```

Figure 1. Block Diagram of the Interfaces and Signals for the Component dut


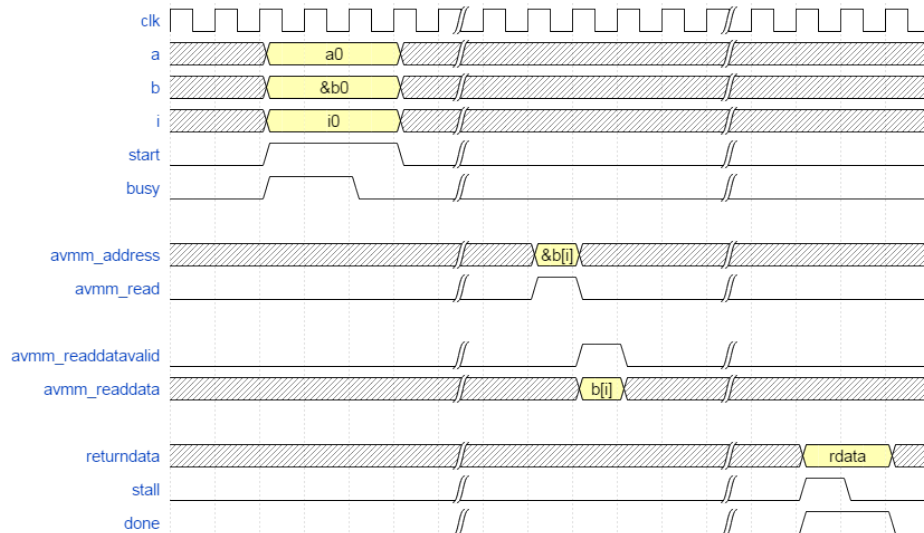




Figure 2. Waveform Diagram of the Signals for the Component dut

This diagram shows that the Avalon-MM read signal reads from a memory interface that has a read latency of one cycle and is non-blocking.



If the `dut` component raises the `busy` signal, the caller needs to keep the `start` signal high and continue asserting the input arguments. Similarly, if the component downstream of `dut` raises the `stall` signal, then `dut` holds the `done` signal high until the `stall` signal is de-asserted.

4.2. Avalon Streaming Interfaces

A component can have input and output streams that conform to the Avalon-ST interface specifications. These input and output streams are represented in the C source by passing references to `ihc::stream_in<>` and `ihc::stream_out<>` objects as function arguments to the component.

When you use an Avalon-ST interface, you can serialize the data over several clock cycles. That is, one component invocation can read from a stream multiple times.

You cannot derive new classes from the stream classes or encapsulate them in other formats such as structs or arrays. However, you may use these classes as references inside other classes, meaning that you can create a class that has a reference to a stream as a data member.

A component can have multiple read sites for a stream. Similarly, a component can have multiple write sites for the same stream. However, try to restrict each stream in your design to a single read site, a single write site, or one of each.

Note: Within the component, there is no guarantee on the order of execution of different streams unless a data dependency exists between streams.

For more information about streaming interfaces, refer to "[Avalon Streaming Interfaces](#)" in *Avalon Interface Specifications*.

Streaming Input Interfaces

Table 9. Intel HLS Compiler Streaming Input Interface `stream_in` Declarations

Feature	Valid Values	Default Value	Description
<code>ihc::stream_in<datatype, template arguments></code>	Any valid C++ datatype		<p>Streaming input interface to the component. The width of the stream data bus is equal to a width of <code>sizeof(datatype)</code>.</p> <p>The testbench must populate this buffer (stream) fully before the component can start to read from the buffer.</p> <p>To learn more, review the following tutorials:</p> <ul style="list-style-type: none"> <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_buffer</code> PRO <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packets_empty</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packet_ready_valid</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_ready_latency</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/multiple_stream_call_sites</code>
<code>ihc::buffer<value></code>	Non-negative integer value	0	<p>The capacity, in words, of the FIFO buffer on the input data that associates with the stream.</p> <p>This parameter is available only on input streams.</p>
<code>ihc::readylatency<value></code>	Non-negative integer value (between 0-8)	0	<p>The number of cycles between when the <code>ready</code> signal is deasserted and when the input stream can no longer accept new inputs.</p>
<code>ihc::bitsPerSymbol<value></code>	A positive integer value that evenly divides by the data type size	Datatype size	<p>Describes how the data is broken into symbols on the data bus.</p> <p>PRO Data is broken down according to how you set the <code>ihc::firstSymbolInHighOrderBits</code> declaration. By default, data is broken down in little endian order.</p> <p>STD Data is always broken down in little endian order.</p>
PRO <code>ihc::firstSymbolInHighOrderBits<value></code>	true or false	false	<p>Specifies whether the data symbols in the stream are in big endian order.</p>

continued...



Feature	Valid Values	Default Value	Description
<code>ihc::usesPackets<value></code>	true or false	false	Exposes the <code>startofpacket</code> and <code>endofpacket</code> sideband signals on the stream interface, which can be accessed by the packet based reads/writes
PRO <code>ihc::usesEmpty<value></code>	true or false	false	Exposes the empty out-of-band signal on the stream interface. Use this declaration only with streams that read more than one data symbol per clock cycle. The <code>empty</code> signal indicates the number of symbols on the data bus that do not represent valid data during the final stream read of a packet. You can control whether the <code>empty</code> symbols are in the low-order bits or high-order bits with the <code>ihc::firstSymbolInHighOrderBits</code> declaration.
<code>ihc::usesValid<value></code>	true or false	true	Controls whether a <code>valid</code> signal is present on the stream interface. If <code>false</code> , the upstream source must provide valid data on every cycle that <code>ready</code> is asserted. This is equivalent to changing the stream read calls to <code>tryRead</code> and assuming that <code>success</code> is always <code>true</code> . If set to <code>false</code> , <code>buffer</code> and <code>readyLatency</code> must be 0.
<p>The following code example illustrates both <code>stream_in</code> declarations and <code>stream_in</code> function APIs.</p> <pre> // Blocking read void foo (ihc::stream_in<int> &a) { int x = a.read(); } // Non-blocking read void foo_nb (ihc::stream_in<int> &a) { bool success = false; int x = a.tryRead(success); if (success) { // x is valid } } int main() { ihc::stream_in<int> a; ihc::stream_in<int> b; for (int i = 0; i < 10; i++) { a.write(i); b.write(i); } foo(a); foo_nb(b); } </pre>			

Table 10. Intel HLS Compiler Streaming Input Interface `stream_in` Function APIs

Feature	Description
<code>T read()</code>	Blocking read call to be used from within the component
<code>T read(bool& sop, bool& eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking read with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>T read(bool& sop, bool& eop, int& empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking read with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.
<code>T tryRead(bool &success)</code>	Non-blocking read call to be used from within the component. The <code>success</code> bool is set to true if the read was valid. If you use <code>tryRead</code> , your x86-64 results for your component might not match your FPGA results because emulation does not model the hardware behavior of blocking and non-blocking reads.
<i>continued...</i>	

Feature	Description
<code>T tryRead(bool& success, bool& sop, bool& eop)</code>	Available only if <code>usesPackets<true></code> is set. Non-blocking read with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>T tryRead(bool& success, bool& sop, bool& eop, int& empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Non-blocking read with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>emptysignals</code> .
<code>void write(T data)</code>	Blocking write call to be used from the testbench to populate the FIFO to be send to the component
<code>void write(T data, bool sop, bool eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking write call with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>void write(T data, bool sop, bool eop, int empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking write call with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.

Streaming Output Interfaces

Table 11. Intel HLS Compiler Streaming Output Interfaces `stream_out` Declaration

Feature	Valid Values	Default Value	Description
<code>ihc::stream_out<datatype, template arguments></code>	Any valid POD (plain old data) C++ datatype		Streaming output interface from the component. The testbench can read from this buffer once the component returns. To learn more, review the following tutorials: <ul style="list-style-type: none"> <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_buffer</code> PRO <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packets_empty</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packet_ready_valid</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_ready_latency</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/multiple_stream_call_sites</code>
<code>ihc::readylatency<value></code>	Non-negative integer value (between 0-8)	0	The number of cycles between when the <code>ready</code> signal is deasserted and when the input stream can no longer accept new inputs. Conceptually, you can view this parameter as an almost ready latency on the input FIFO buffer for the data that associates with the stream.
<code>ihc::bitsPerSymbol<value></code>	Positive integer value that evenly divides the data type size	Datatype size	Describes how the data is broken into symbols on the data bus. PRO Data is broken down according to how you set the <code>ihc::firstSymbolInHighOrderBits</code> declaration. By default, data is broken down in little endian order.

continued...



Feature	Valid Values	Default Value	Description
			STD Data is always broken down in little endian order.
PRO ihc::firstSymbolInHighOrderBits<value>	true or false	false	Specifies whether the data symbols in the stream are in big endian order.
ihc::usesPackets<value>	true or false	false	Exposes the startofpacket and endofpacket sideband signals on the stream interface, which can be accessed by the packet based reads/writes.
PRO ihc::usesEmpty<value>	true or false	false	Exposes the empty out-of-band signal on the stream interface. Use this declaration only with streams that write more than one data symbol per clock cycle. The empty signal indicates the number of symbols on the data bus that do not represent valid data during the final stream write of a packet. You can control whether the empty symbols are in the low-order bits or high-order bits with the ihc::firstSymbolInHighOrderBits declaration.
ihc::usesReady<value>	true or false	true	Controls whether a ready signal is present. If false, the downstream sink must be able to accept data on every cycle that valid is asserted. This is equivalent to changing the stream read calls to tryWrite and assuming that success is always true. If set to false, readyLatency must be 0.

The following code example illustrates both stream_out declarations and stream_out function APIs.

```
// Blocking write
void foo (ihc::stream_out<int> &a) {
    static int count = 0;
    for(int idx = 0; idx < 5; idx++){
        a.write(count++); // Blocking write
    }
}

// Non-blocking write
void foo_nb (ihc::stream_out<int> &a) {
    static int count = 0;
    for(int idx = 0; idx < 5; idx++){
        bool success = a.tryWrite(count++); // Non-blocking write
        if (success) {
            // write was successful
        }
    }
}

int main() {
    ihc::stream_out<int> a;
    foo(a); // or foo_nb(a);

    // copy output to an array
    int outputData[5];
    for (int i = 0; i < 5; i++) {
        outputData[idx] = a.read();
    }
}
```

Table 12. Intel HLS Compiler Streaming Output Interfaces `stream_out` Function Call APIs

Feature	Description
<code>void write(T data)</code>	Blocking write call from the component
<code>void write(T data, bool sop, bool eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking write with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>void write(T data, bool sop, bool eop, int empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking write with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.
<code>bool tryWrite(T data)</code>	Non-blocking write call from the component. The return value represents whether the write was successful.
<code>bool tryWrite(T data, bool sop, bool eop)</code>	Available only if <code>usesPackets<true></code> is set. Non-blocking write with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals. The return value represents whether the write was successful.
PRO <code>bool tryWrite(T data, bool sop, bool eop, int empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Non-blocking write with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals. The return value represents whether the write was successful.
<code>T read()</code>	Blocking read call to be used from the testbench to read back the data from the component
<code>T read(bool &sop, bool &eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking read call to be used from the testbench to read back the data from the component with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>T read(bool &sop, bool &eop, int &empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking read call to be used from the testbench to read back the data from the component with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.

Related Information

[Avalon Interface Specifications](#)

4.3. Avalon Memory-Mapped Master Interfaces

A component can interface with an external memory over an Avalon Memory-Mapped (MM) Master interface. You can specify the Avalon MM Master interface implicitly using a function pointer argument or reference argument, or explicitly using the `mm_master<>` class defined in the "HLS/hls.h" header file. Describe a customized Avalon MM Master interface in your code by including a reference to an `mm_master<>` object in your component function signature.

Each `mm_master` argument of a component results in an input conduit for the address. That input conduit is associated with the component start and busy signals. In addition to this input conduit, a unique Avalon MM Master interface is created for each address space. Master interfaces that share the same address space are arbitrated on the same interface.

For more information about Avalon MM Master interfaces, refer to "[Avalon Memory-Mapped Interfaces](#)" in *Avalon Interface Specifications*.



Table 13. Intel HLS Compiler Memory-Mapped Interfaces

Feature	Valid Values	Default Value	Description
<code>ihc::mm_master<datatype, template arguments></code>	Any valid C++ datatype	Default interface for pointer arguments	<p>The underlying pointer type. Pointer arithmetic performed on the master object conforms to this type. Dereferences of the master results in a load-store site with a width of <code>sizeof(datatype)</code>. The default alignment is aligned to the size of the datatype.</p> <p>Avalon Memory-Mapped (MM) Master interface argument: Multiple template arguments are supported. The template arguments are listed below. Any combination can be used as long as it describes a valid hardware configuration.</p> <p>Example:</p> <pre>component int dut(ihc::mm_master<int, ihc::aspace<2>, ihc::latency<3>, ihc::awidth<10>, ihc::dwidth<32> > &a)</pre> <p>To learn more, review the following tutorials:</p> <ul style="list-style-type: none"> <code><quartus_installdir>/hls/examples/tutorials/interfaces/pointer_mm_master</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/mm_master_testbench_operators</code>
<code>ihc::dwidth<value></code>	8, 16, 32, 64, 128, 256, 512, or 1024	64	The width of the memory-mapped data bus in bits
<code>ihc::awidth<value></code>	Integer value in the range 1 – 64	64	The width of the memory-mapped address bus in bits. This value affects only the width of the Avalon MM Master interface. The size of the conduit of the base address pointer is always set to 64-bits.
<code>ihc::aspace<value></code>	Integer value greater than 0	1	The address space of the interface that associates with the master. Each unique <code>value</code> results in a separate Avalon MM Master interface on your component. All masters with the same address space are arbitrated within the component to a single interface. As such, these masters must share the same template parameters that describe the interface.
<code>ihc::latency<value></code>	Non-negative integer value	1	The guaranteed latency from when a read command exits the component when the external memory returns valid read data. If this latency is variable (such as when accessing DRAM), set it to 0.
<code>ihc::maxburst<value></code>	Integer value in the range 1 – 1024	1	The maximum number of data transfers that can associate with a read or write transaction. This value controls the width of the <code>burstcount</code> signal. For fixed latency interfaces, this value must be set to 1. For more details, review information about burst signals and the <code>burstcount</code> signal role in "Avalon Memory-Mapped Interface Signal Roles" in <i>Avalon Interface Specifications</i> .
<code>ihc::align<value></code>	Integer value greater than the alignment of the datatype	Alignment of the datatype	<p>The alignment of the base pointer address in bytes. The Intel HLS Compiler uses this information to determine how many simultaneous loads and stores this pointer can permit.</p> <p>For example, if you have a bus with 4 32-bit integers on it, you should use <code>ihc::dwidth<128></code> (bits) and <code>ihc::align<16></code> (bytes). This means that up to 16 contiguous bytes (or 4 32-bit integers) can be loaded or stored as a coalesced memory word per clock cycle.</p>

continued...

Feature	Valid Values	Default Value	Description
			<i>Important:</i> The caller is responsible for aligning the data to the set value for the align argument; otherwise, functional failures might occur.
<code>ihc::readwrite_mode<value></code>	readwrite, readonly, or writeonly	readwrite	Port direction of the interface. Only the relevant Avalon master signals are generated
<code>ihc::waitrequest<value></code>	true or false	false	Adds the waitrequest signal that is asserted by the slave when it is unable to respond to a read or write request. For more information about the waitrequest signal, see "Avalon Memory-Mapped Interface Signal Roles" in <i>Avalon Interface Specifications</i> .
<code>getInterfaceAtIndex(int index)</code>			This testbench function is used to index into an <code>mm_master</code> object. It can be useful when iterating over an array and invoking a component on different indices of the array. This function is supported only in the testbench. Example: <pre>int main() { // for(int idx = 0; idx < N; idx++) { dut(src_mm.getInterfaceAtIndex(idx)); } // }</pre>

Related Information

[Avalon Interface Specifications](#)

4.3.1. Memory-Mapped Master Testbench Constructor

For components that use an instance of the Avalon Memory-Mapped (MM) Master class (`mm_master<>`) to describe their memory interfaces, you must create an `mm_master<>` object in the testbench for each `mm_master` argument.

To create an `mm_master<>` object, add the following constructor in your code:

```
ihc::mm_master<int, ... > mm(void* ptr, int size, bool use_socket=false);
```

where the constructor arguments are as follows:

- `ptr` is the underlying pointer to the memory in the testbench
- `size` is the total size of the buffer in bytes
- `use_socket` is the option you use to override the copying of the memory buffer and have all the memory accesses pass back to the testbench memory

By default, the Intel HLS Compiler copies the memory buffer over to the simulator and then copies it back after the component has run. In some cases, such as pointer-chasing in linked lists, copying the memory buffer back and forth is undesirable. You can override this behavior by setting `use_socket` to `true`.

Note: When you set `use_socket` to `true`, only Avalon MM Master interfaces with 64-bit wide addresses are supported. In addition, setting this option increases the run time of the simulation.



4.3.2. Implicit and Explicit Examples of Describing a Memory Interface

Optimize component code that describes a memory interface by specifying an explicit `mm_master` object.

Implicit Example

The following code example arbitrates the load and store instructions from both pointer dereferences to a single interface on the component's top-level module. This interface will have a data bus width of 64 bits, an address width of 64 bits, and a fixed latency of 1.

```
#include "HLS/hls.h"
component void dut(int *ptr1, int *ptr2) {
    *ptr1 += *ptr2;
    *ptr2 += ptr1[1];
}

int main(void) {
    int x[2] = {0, 1};
    int y = 2;

    dut(x, &y);

    return 0;
}
```

Explicit Example

This example demonstrates how to optimize the previous code snippet for a specific memory interface using the explicit `mm_master` class. The `mm_master` class has a defined template, and it has the following characteristics:

- Each interface is given a unique ID that infers two independent interfaces and reduces the amount of arbitration within the component.
- The data bus width is larger than the default width of 64 bits.
- The address bus width is smaller than the default width of 64 bits.
- The interfaces have a fixed latency of 2.

By defining these characteristics, you state that your system returns valid read data after exactly two clock cycles and that the interface never stalls for both reads and writes, but the system must be able to provide two different memories. A unique `aspace` is expected to correspond to a unique physical memory. If you connect to multiple Avalon-MM Master interface with the same `aspace` to the same physical memory, the Intel HLS Compiler cannot ensure functional correctness for any memory dependencies.

```
#include "HLS/hls.h"

typedef ihc::mm_master<int, ihc::dwidth<256>,
                    ihc::awidth<32>,
                    ihc::aspace<1>,
                    ihc::latency<2> > Master1;

typedef ihc::mm_master<int, ihc::dwidth<256>,
                    ihc::awidth<32>,
                    ihc::aspace<4>,
                    ihc::latency<2> > Master2;

component void dut(Master1 &mm1, Master2 &mm2) {
    *mm1 += *mm2;
    *mm2 += mm1[1];
}
```

```

}
int main(void) {
    int x[2] = {0, 1};
    int y = 2;

    Master1 mm_x(x, 2*sizeof(int), false);
    Master2 mm_y(&y, sizeof(int), false);

    dut(mm_x, mm_y);

    return 0;
}

```

4.4. Slave Interfaces

The Intel HLS Compiler provides two different types of slave interfaces that you can use with a component. In general, smaller scalar inputs should use slave registers. Large arrays should use slave memories if your intention is to copy these arrays into or out of the component.

Slave interfaces are implemented as Avalon Memory Mapped (Avalon-MM) Slave interfaces. For details about the Avalon-MM Slave interfaces, see "[Avalon Memory-Mapped Interfaces](#) in *Avalon Interface Specifications*.

Table 14. Types of Slave Interfaces

Slave Type	Associated Slave Interface	Read/Write Behavior	Synchronization	Read Latency	Controlling Interface Data Width
Register	The component control and status register (CSR) slave.	The component cannot update these registers from the datapath, so you can read back only data that you wrote in.	Synchronized with the component <code>start</code> signal.	Fixed value of 1.	Always 64 bits
Memory (M20K)	Dedicated slave interface on the component.	Updates from the component's datapath are visible in memory.	Reads and writes to slave memories from outside of the component should occur only when your component is not executing. You might experience undefined component behavior if outside slave memory accesses occur when your component is executing. The undefined behavior can occur even if a slave memory access is to a memory address that the component does not access.	Fixed value that is dependent on the component memory access pattern and any attributes or pragmas that you set. See the component <code>.qsys</code> file for more information.	The data width is a multiple of the slave data type, where the multiple is determined by coalescing the internal accesses.

4.4.1. Control and Status Register (CSR) Slave

A component can have a maximum of one CSR slave interface, but more than one argument can be mapped into this interface.



Any arguments that are labeled as `hls_avalon_slave_register_argument` are located in this memory space. The resulting memory map is described in the automatically generated header file `<results>.prj/components/<component_name>_csr.h`. This file also provides the C macros for a master to interact with the slave.

The control and status registers (that is, function call and return) of an `hls_avalon_slave_component` attribute are implemented in this interface.

You do not need to use the `hls_avalon_slave_component` attribute to use the `hls_avalon_slave_register_argument` attribute.

To learn more, review the tutorial: `<quartus_installdir>/hls/examples/tutorials/interfaces/mm_slaves`

Example code of a component with a CSR slave:

```
#include "HLS/hls.h"

struct MyStruct {
    int f;
    double j;
    short k;
};

hls_avalon_slave_component
component MyStruct mycomp_xyz (hls_avalon_slave_register_argument int y,
                               hls_avalon_slave_register_argument MyStruct struct_argument,
                               hls_avalon_slave_register_argument unsigned long long mylong,
                               hls_avalon_slave_register_argument char char_arg
                               ) {
    return struct_argument;
}
```

Generated C header file for the component `mycomp_xyz`:

```
/* This header file describes the CSR Slave for the mycomp_xyz component */
#ifndef __MYCOMP_XYZ_CSR_REGS_H__
#define __MYCOMP_XYZ_CSR_REGS_H__

/
*****
/
/* Memory Map Summary
*/
/
*****
/

/*


| Register Address | Access | Register Contents (64-bits) | Description                                                                                                                              |
|------------------|--------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| 0x0              | R      | {reserved[62:0], busy[0:0]} | Read the busy status of the component<br>0 - the component is ready to accept a new start<br>1 - the component cannot accept a new start |


*/
```



0x8	W	{reserved[62:0], start[0:0]}	Write 1 to signal start to the component
0x10	R/W	{reserved[62:0], interrupt_enable[0:0]}	0 - Disable interrupt, 1 - Enable interrupt
0x18 completion	R/Wclr	{reserved[61:0], done[0:0], interrupt_status[0:0]}	Signals component done is read-only and interrupt_status is write 1 to clear
0x20	R	{returndata[63:0]}	Return data (0 of 3)
0x28	R	{returndata[127:64]}	Return data (1 of 3)
0x30	R	{returndata[191:128]}	Return data (2 of 3)
0x38	R/W	{reserved[31:0], y[31:0]}	Argument y
0x40 (0 of 3)	R/W	{struct_argument[63:0]}	Argument struct_argument
0x48 (1 of 3)	R/W	{struct_argument[127:64]}	Argument struct_argument
0x50 (2 of 3)	R/W	{struct_argument[191:128]}	Argument struct_argument
0x58	R/W	{mylong[63:0]}	Argument mylong
0x60	R/W	{reserved[55:0], char_arg[7:0]}	Argument char_arg

NOTE: Writes to reserved bits will be ignored and reads from reserved bits will return undefined values.

```

*/

/
*****
/
/* Register Address Macros
*/
/
*****
/

/* Byte Addresses */
#define MYCOMP_XYZ_CSR_BUSY_REG (0x0)
#define MYCOMP_XYZ_CSR_START_REG (0x8)
#define MYCOMP_XYZ_CSR_INTERRUPT_ENABLE_REG (0x10)
#define MYCOMP_XYZ_CSR_INTERRUPT_STATUS_REG (0x18)
#define MYCOMP_XYZ_CSR_RETURNDATA_0_REG (0x20)
#define MYCOMP_XYZ_CSR_RETURNDATA_1_REG (0x28)
#define MYCOMP_XYZ_CSR_RETURNDATA_2_REG (0x30)
#define MYCOMP_XYZ_CSR_ARG_Y_REG (0x38)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_0_REG (0x40)

```



```

#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_1_REG (0x48)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_2_REG (0x50)
#define MYCOMP_XYZ_CSR_ARG_MYLONG_REG (0x58)
#define MYCOMP_XYZ_CSR_ARG_CHAR_ARG_REG (0x60)

/* Argument Sizes (bytes) */
#define MYCOMP_XYZ_CSR_RETURNDATA_0_SIZE (8)
#define MYCOMP_XYZ_CSR_RETURNDATA_1_SIZE (8)
#define MYCOMP_XYZ_CSR_RETURNDATA_2_SIZE (8)
#define MYCOMP_XYZ_CSR_ARG_Y_SIZE (4)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_0_SIZE (8)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_1_SIZE (8)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_2_SIZE (8)
#define MYCOMP_XYZ_CSR_ARG_MYLONG_SIZE (8)
#define MYCOMP_XYZ_CSR_ARG_CHAR_ARG_SIZE (1)

/* Argument Masks */
#define MYCOMP_XYZ_CSR_RETURNDATA_0_MASK (0xffffffffffffffffULL)
#define MYCOMP_XYZ_CSR_RETURNDATA_1_MASK (0xffffffffffffffffULL)
#define MYCOMP_XYZ_CSR_RETURNDATA_2_MASK (0xffffffffffffffffULL)
#define MYCOMP_XYZ_CSR_ARG_Y_MASK (0xffffffff)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_0_MASK (0xffffffffffffffffULL)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_1_MASK (0xffffffffffffffffULL)
#define MYCOMP_XYZ_CSR_ARG_STRUCT_ARGUMENT_2_MASK (0xffffffffffffffffULL)
#define MYCOMP_XYZ_CSR_ARG_MYLONG_MASK (0xffffffffffffffffULL)
#define MYCOMP_XYZ_CSR_ARG_CHAR_ARG_MASK (0xff)

/* Status/Control Masks */
#define MYCOMP_XYZ_CSR_BUSY_MASK (1<<0)
#define MYCOMP_XYZ_CSR_BUSY_OFFSET (0)

#define MYCOMP_XYZ_CSR_START_MASK (1<<0)
#define MYCOMP_XYZ_CSR_START_OFFSET (0)

#define MYCOMP_XYZ_CSR_INTERRUPT_ENABLE_MASK (1<<0)
#define MYCOMP_XYZ_CSR_INTERRUPT_ENABLE_OFFSET (0)

#define MYCOMP_XYZ_CSR_INTERRUPT_STATUS_MASK (1<<0)
#define MYCOMP_XYZ_CSR_INTERRUPT_STATUS_OFFSET (0)
#define MYCOMP_XYZ_CSR_DONE_MASK (1<<1)
#define MYCOMP_XYZ_CSR_DONE_OFFSET (1)

#endif /* __MYCOMP_XYZ_CSR_REGS_H__ */

```

4.4.2. Slave Memories

By default, component functions access parameters that are passed by reference through an Avalon Memory-Mapped (MM) Master interface. An alternative way to pass parameters by reference is to use an Avalon MM Slave interface, which exists inside the component.

Having a pointer argument generate an Avalon MM Master interface on the component has two potential disadvantages:

- The master interface has a single port. If the component has multiple load-store sites, arbitration on that port might create stalling logic.
- Depending on the system in which the component is instantiated, other masters might use the memory bus while the component is running and create undesirable stalls on the bus.

Because a slave memory is internal to the component, the HLS compiler can create a memory architecture that is optimized for the access pattern of the component such as creating banked memories or coalescing memories.

Slave memories differ from local memories because they can be accessed from an Avalon MM Master outside of the component. Local memories are by definition local to the component and cannot be accessed outside the component. Unlike local memory components, you cannot explicitly configure slave memory arguments (for example, banking or coalescing). You must rely on the automatic configurations generated by the compiler. You can control the structure of your slave memories only by restructuring your load and store operations.

Important: Reads and writes to slave memories from outside of the component should occur only when your component is not executing. You might experience undefined component behavior if outside slave memory accesses occur when your component is executing. The undefined behavior can occur even if a slave memory access is to a memory address that the component does not access.

A component can have many slave memory interfaces. Unlike slave register arguments that are grouped together in the CSR slave interface, each slave memory has a separate interface with separate data buses. The slave memory interface data bus width is determined by the width of the slave type. If the internal accesses to the memory have been coalesced, the slave memory interface data bus width might be a multiple of the width of the slave type.

Argument Label	Description
<code>hls_avalon_slave_memory_argument(N)</code>	<p>The compiler implements the argument, where <i>N</i> specifies the size of the memory in bytes, in on-chip memory blocks, which can be read from or written to over a dedicated slave interface. The generated memory has the same architectural optimizations as all other internal component memories (that is, banking, coalescing, etc.).</p> <p>If the compiler performs static coalescing optimizations, the slave interface's data width will be the coalesced width. This attribute applies only to a pointer argument.</p> <p>Example:</p> <pre>component void foo(hls_avalon_slave_memory_argument(128*sizeof(int)) int *a)</pre> <p>To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/interfaces/mm_slaves</code></p>

4.5. Component Invocation Interface Arguments

The component invocation interface refers to the control signals that correspond to actions of calling the function. All unstable component argument inputs are synchronized according to this component invocation protocol. A component argument is unstable if it changes while there is live data in the component (that is, between pipelined function invocations).

Table 15. Intel HLS Compiler Component Invocation Interface Arguments

Feature	Description
<code>hls_avalon_streaming_component</code> This is the default component invocation interface.	This attribute follows the Avalon-ST protocol for both the function call and the return streams. The component consumes the unstable arguments when the <code>start</code> signal is asserted and the <code>busy</code> signal is deasserted. The component produces the return data when the <code>done</code> signal is asserted.

continued...



Feature	Description
	<p>Top-level module ports: Function call—start, busy Function return—done, stall Example:</p> <pre data-bbox="667 436 1409 485">component hls_avalon_streaming_component void foo(/*component arguments*/)</pre>
hls_avalon_slave_component	<p>The start, done, and returndata (if applicable) signals are registered in the component slave memory map. These component must take either slave, stream, or stable arguments. If you do not specify these types of arguments, the compiler generates an error message when you compile this component. Top-level module ports: Avalon-MM slave interface and irq_done signal Example:</p> <pre data-bbox="667 705 1409 730">component hls_avalon_slave_component void foo(/*component arguments*/)</pre> <p>To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/interfaces/mm_slaves</code></p>
hls_always_run_component	<p>The start signal is tied to 1 internally in the component. There is no done signal output. The control logic is optimized away when Intel Quartus Prime compiles the generated RTL for your FPGA. Use this protocol when the component datapath relies only on explicit streams for data input and output. IP verification does not support components with this component invocation protocol. Top-level module ports: None Example:</p> <pre data-bbox="667 1052 1409 1077">component hls_always_run_component void foo(/*component arguments*/)</pre>
hls_stall_free_return	<p>If the downstream component never stalls, the stall signal is removed by internally setting it to 0. This feature can be used with the hls_avalon_streaming_component, hls_avalon_slave_component, and hls_always_run_component arguments. This attribute can be used to specify that the downstream component is stall free. Example:</p> <pre data-bbox="667 1293 1409 1341">component hls_stall_free_return component int dut(int a, int b) { return a * b;}</pre>

Related Information

[Control and Status Register \(CSR\) Slave](#) on page 28

4.6. Unstable and Stable Component Arguments

If you do not specify the intended behavior for an argument, the default behavior of an argument is unstable. An unstable argument can change while there is live data in the component (that is, between pipelined function invocations).

You can declare an interface argument to be stable with the `hls_stable_argument` attribute. A stable interface argument is an argument that does not change while your component executes, but the argument might change between component executions.

You can mark the following the interface arguments as stable:

- Scalar (conduit) arguments
- Pointer interface arguments
The address input is stable. The associated Avalon MM Master interface is not affected.
- Pass-by-reference arguments
The address input is stable. The associated Avalon MM Master interface is not affected.
- Avalon Memory-Mapped (MM) Master interface arguments
The address input is stable. The associated Avalon MM Master interface is not affected.
- Avalon Memory-Mapped (MM) Slave register interface arguments


The following interface arguments cannot be marked as stable:

- Avalon Memory-Mapped (MM) Slave memory interface arguments
- Avalon Streaming interface arguments

You might save some FPGA area in your component design when you declare an interface argument as stable because there is no need to carry the data with the pipeline.

You cannot have two component invocations in flight with different stable arguments between the two component invocations.

Argument Label	Description
hls_stable_argument	<p>A stable argument is an argument that does not change while there is live data in the component (that is, between pipelined function invocations).</p> <p>Changing a stable argument during component execution results in undefined behavior; each use of the stable argument might be the old value or the new value, but with no guarantee of consistency. The same variable in the same invocation can appear with multiple values.</p> <p>Using stable arguments, where appropriate, might save a significant number of registers in a design.</p> <p>Stable arguments can be used with conduits, mm_master interfaces, and slave_registers.</p> <p>Example:</p> <pre data-bbox="852 1386 1388 1480">component int dut(hls_stable_argument int a, hls_stable_argument int b) { return a * b;}</pre> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/interfaces/stable_arguments</p>



4.7. Global Variables

Components can use and update C++ global variables. If you access a global variable in your component function, it is implemented as an Avalon Memory-Mapped (MM) Master interfaces, like a pointer parameter.

If you access more than one global variable, each global variable uses the same Avalon- MM Master interface, which might result in stallable arbitration. If you use pointers and non-constant global memory accesses, then the pointers and global memory accesses all share the same Avalon- MM Master interface.

In addition to the Avalon- MM Master interface, each global variable that the component uses has an input conduit that must be supplied with the address of the global variable in system memory. The input conduit arguments that are generated in the RTL are named `@<global variable name>`. Input conduits generated for pointer arguments omit the `@` are named for the corresponding pointer argument.

If your global variable is declared as `const`, then no Avalon- MM Master interface and no additional input conduit is generated. Therefore, global variables declared as `const` use significantly less FPGA area than modifiable global variable.

4.8. Structs in Component Interfaces

Review the `interface_structs.sv` file in your `<a.prj>/components/<component_name>` folder to see information about the padding and packed-ness of the implementation interfaces for the structs in your component.

The `interface_structs.sv` file contains the Verilog-style definitions of the structs found on your component interface.

4.9. Reset Behavior

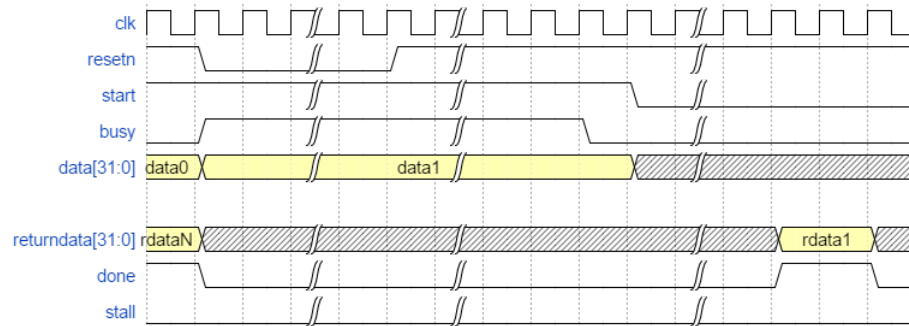
For your HLS component, the reset assertion can be asynchronous but the reset deassertion must be synchronous.

The reset assertion and deassertion behavior can be generated from an asynchronous reset input by using a reset synchronizer, as described in the following example Verilog code:

```
reg [2:0] sync_resetn;
always @(posedge clock or negedge resetn) begin
    if (!resetn) begin
        sync_resetn <= 3'b0;
    end else begin
        sync_resetn <= {sync_resetn[1:0], 1'b1};
    end
end
```

This synchronizer code is used in the example Intel Quartus Prime project that is generated for your components included in an i++ compile.

When the reset is asserted, the component holds its `busy` signal high and its `done` signal low. After the reset is deasserted, the component holds its `busy` signal high until the component is ready to accept the next invocation. All component interfaces (slaves, masters, and streams) are valid only after the component `busy` signal is low.



Simulation Component Reset

You can check the reset behavior of your component during simulation by using the `ihc_hls_sim_reset` API. This API returns 1 if the reset was exercised (that is, if the reset is called during hardware simulation of the component). Otherwise, the API returns 0.

Call the API as follows:

```
int ihc_hls_sim_reset(void);
```

During x86 emulation of your component, the `ihc_hls_sim_reset` API always returns 0. You cannot reset a component during x86 emulation.

5. Local Variables in Components (Memory Attributes)

The Intel High Level Synthesis (HLS) Compiler tries to provide the maximum throughput whenever possible. In certain cases, particularly when the Intel HLS Compiler optimizes local memory configurations for throughput, it might be beneficial to trade some throughput for a smaller area. Apply the component memory attributes to local variables in your component to customize the on-chip memory architecture of the local memory system and lower the FPGA area utilization of your component. These component memory attributes are defined in the "HLS/hls.h" header file, which you can include in your code.

Restriction: You can apply these attributes to primitives and objects, but not to class members. That is, memory attributes cannot be applied to member variables in a struct or a class.

Table 16. Intel HLS Compiler Component Memory Attributes


Attribute	Default Value	Description
hls_register	Based on the memory access pattern inferred by the compiler.	Forces a variable or array inside component to be implemented as registers. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/best_practices/swap_vs_copy
hls_memory	Based on the memory access pattern inferred by the compiler.	Forces a variable or array inside component to be implemented as embedded memory. To learn more, review the design: <quartus_installdir>/hls/examples/QRD
hls_singlepump	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must be single pumped. That is, the memory is clocked at the same operating frequency as the operating frequency of your component. To learn more, review the design: <quartus_installdir>/hls/examples/QRD
hls_doublepump	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must be double pumped. That is, the memory is clocked at twice the operating frequency of your component.
hls_numbanks(<i>N</i>)*	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must have <i>N</i> banks, where <i>N</i> is a power-of-two constant number.
hls_bankwidth(<i>N</i>)*	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must have banks that are <i>N</i> bytes wide, where <i>N</i> is a power-of-two constant number.

continued...

(*) This attribute is subject to constraints outlined in [Constraints on Attributes for Memory Banks](#) on page 39.



Attribute	Default Value	Description
		To learn more, review the design: <quartus_installdir>/hls/tutorials/ component_memories/bank_bits
hls_bankbits(b_0, b_1, \dots, b_n)(*)	Lowest bits of the address based on number of banks.	Forces the memory system to split into 2^{n+1} banks, with $\{b_0, b_1, \dots, b_n\}$ forming the bank-select bits. <i>Important:</i> b_0, b_1, \dots, b_n must be consecutive, positive integers. You can specify the consecutive, positive integers in ascending or descending order. If you do not specify the hls_bankwidth(N) attribute along with this attribute, then b_0, b_1, \dots, b_n are mapped to array index bits 0 to $n-1$ in the memory bank implementation. To learn more, review the design: <quartus_installdir>/hls/tutorials/ component_memories/bank_bits
hls_numports_readonly_writeonly(M, N)	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must have M read ports and N write ports, where M and N are constant numbers greater than zero.
hls_simple_dual_port_memory		Specifies the configuration that is defined by the presence of both the hls_singlepump and the hls_numports_readonly_writeonly(1,1) macros.
hls_merge("mem_name", "depth")		Allows merging two or more local variables to be implemented in component memory as a single merged memory system in a depth-wise manner. All variables with same <mem_name> label specified in their hls_merge attribute are merged into the same memory system. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ best_practices/depth_wise_merge
hls_merge("mem_name", "width")		Allows merging two or more local variables to be implemented in component memory as a single merged memory system in a width-wise manner. All variables with same <mem_name> label specified in their hls_merge attribute are merged into the same memory system. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ best_practices/width_wise_merge
hls_init_on_reset	Default behavior for static variables.	Forces the static variables inside the component to be reset when the component reset signal is asserted. This requires the an additional write port to the component memory implemented and can increase the power-up latency when the component is reset. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ component_memories/static_var_init
hls_init_on_powerup		Sets the component memory implementing the static variable to set on power-up when the FPGA is programmed. When the component is reset, the component memory is not reset back to the initialized value of the static. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ component_memories/ static_var_init



Constraints on Attributes for Memory Banks

The properties of memory banks constrain how you can divide local memory into banks with the memory bank attributes.

The relationship between the following properties is constrained:

- The size (in bytes) of the local variable (S). If you are accessing an array, this value represents the number of bytes that you want to access at one time.
- The number of memory banks specified by `hls_numbanks` attribute (N_{banks}).
- The width (in bytes) of the memory banks specified by `hls_bankwidth` attribute (W).
- The number of memory bank-select bits specified by `hls_bankbits` attribute. That is, $n+1$ when you specify b_0, b_1, \dots, b_n as the bank-select bits (N_{bits}).

These attributes are subject to the following constraints:

- $N_{\text{banks}} \times W = S$
The size of a local variable is equal to the number of memory banks it uses times the width of the memory banks.
- N_{banks} must be a power of 2 value.
- $N_{\text{banks}} = 2^{N_{\text{bits}}}$
 N_{bits} bank-selection bits that are required to address N_{banks} number of memory banks.

Values that you specify for the `hls_numbanks`, `hls_bankwidth`, and `hls_bankbits` attributes must meet these constraints. For attributes that you do not specify, the Intel HLS Compiler infers values for the attributes following these constraints.

5.1. Static Variables

The HLS compiler supports function-scope static variables with the same semantics as in C and C++.

Function-scope static variables are initialized to the specified values on reset. In addition, changes to these variables are visible across component invocations, making function-scope static variables ideal for storing states in a component.

To initialize static variables, the component requires extra logic, and the component might take some time to exit the reset state while this logic is active.

Static Variable Initialization

Unlike a typical program, you can control when the static variables in your component are initialized, if they are implemented as memories. A static variable can be initialized either when a component is powered up or reset.

Initializing a static variable when a component is powered up resembles a traditional programming model where you cannot reinitialize the static variable value after the program starts to run.

Initializing a static variable when a component is reset initializes the static variable each time each time your component receives a `reset` signal, including on power up. However, this type of static variable initialization requires extra logic. This extra logic can affect the start-up latency and the FPGA area needed for your component.

You can explicitly set the static variable initialization by adding one of the following attributes to your static variable declaration:

hls_init_on_reset
(default behavior)

The static variable value is initialized after the component is reset.

Add this attribute to your static variable declaration as shown in the following example:

```
static char arr[128] hls_init_on_reset;
```

This is the default behavior for initializing static variables. You do not need to specify the `hls_init_on_reset` keyword with your static variable declaration to get this behavior.

For example, the static variable in the following example is also initialized when the component is reset:

```
static int arr[64];
```

hls_init_on_powerup

The static variable is initialized only on power up. This initialization uses a memory initialization file (`.mif`) to initialize the memory, which reduces the resource utilization and start-up latency of the component.

Add this keyword to your static variable declaration as shown in the following example:

```
static char arr[128] hls_init_on_powerup;
```

Some static variables might not be able to take advantage of this initialization because of the complexity of the static variables (for example, an array of structs). In these cases, the compiler returns an error.

For a demonstration of initializing static variables, review the tutorial in `<quartus_installdir>/hls/examples/tutorials/component_memories/static_var_init`.

For information about resetting your component, see [Reset Behavior](#) on page 35.

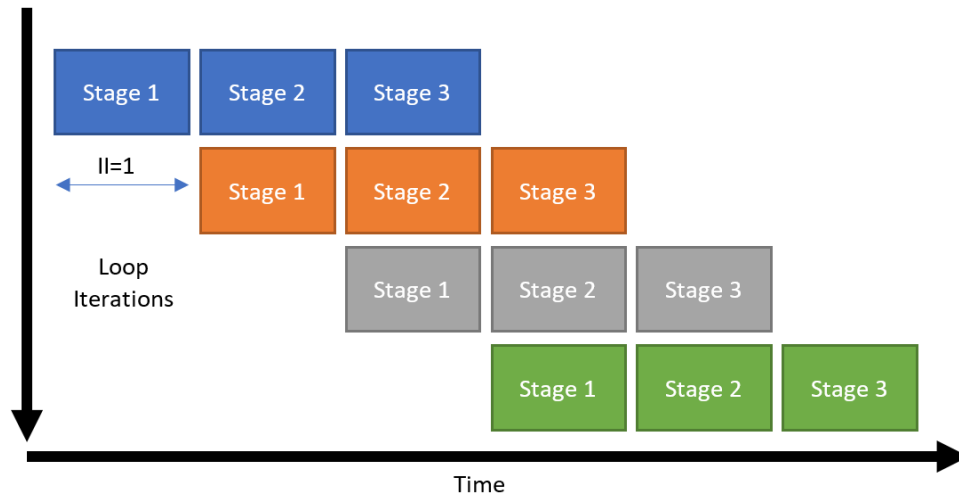
6. Loops in Components

The Intel HLS Compiler attempts to pipeline loops to maximize throughput of the various components that you define.

Loop Pipelining

Pipelining loops enables the Intel HLS Compiler to execute subsequent iterations of a loop in a pipeline-parallel fashion. Pipeline-parallel execution means that multiple iterations of the loop, at different points in their executions, are executing at the same time. Because all stages of the loop are always active, pipelining loops helps maximize usage of the generated hardware.

Figure 3. Pipelined loop with three stages and four iterations



There are some cases where pipelining is not possible at all. In other cases, a new iteration of the loop cannot start until N cycles after the previous iteration.

The number of cycles for which a loop iteration must wait before it can start is called the initiation interval (II) of the loop. This loop pipelining status is captured in the high level design report (`report.html`). In general, an II of 1 is desirable.

A common case where $II > 1$ is when a part of the loop depends in some way on the results of the previous iteration of the same loop. The circuit must wait for these loop-carried dependencies to be resolved before starting a new iteration of the loop. These loop-carried dependencies are indicated in the optimization report.

In the case of nested loops, $II > 1$ for an outer loop is not considered a significant performance limiter if a critical inner loop carries out the majority of the work. One common performance limiter is if the HLS compiler cannot statically compute the trip count of an inner loop (for example, a variable inner loop trip count). Without a known trip count, the compiler cannot pipeline the outer loop.

For more information about loop pipelining, see [Pipeline Loops](#) in *Intel High Level Synthesis Compiler Best Practices Guide*.

Compiler Pragas Controlling Loop Pipelining

The Intel HLS Compiler has several pragmas that you can specify in your code to control how the compiler pipelines your loops.

Loop pragmas must immediately precede the loop that the pragma applies to. You cannot have a loop pragma before elements such as labels on loops. The following table shows examples of how to apply loop pragmas correctly.

Incorrect	Correct
<pre>#pragma ivdep TEST_LOOP: for(int idx = 0; idx < counter; idx ++) {...}</pre>	<pre>TEST_LOOP: #pragma ivdep for(int idx = 0; idx < counter; idx++) {...}</pre>

Table 17. Intel HLS Compiler Loop Pragas

Pragma	Description
<pre>#pragma ii N</pre>	<p>Forces the loop that this is applied on to have a loop initiation interval (II) of $<N>$, where $<N>$ is a positive integer value.</p> <p>This can have an adverse effect on the f_{MAX} of your component because using this pragma to get a lower loop II combines pipeline stages together and creates logic with a long propagation delay.</p> <p>Using this pragma with a larger loop II inserts more pipeline stages and can give you a better component f_{MAX} value.</p> <p>Example:</p> <pre style="background-color: #f0f0f0;">#pragma ii 2 for (int i = 0; i < 8; i++) { // Loop body }</pre>
<pre>#pragma ivdep safelen(N) array(array_name)</pre>	<p>Tells the compiler to ignore memory dependencies between iterations of this loop.</p> <p>It can accept an optional argument that specifies the name of the array. If <code>array</code> is not specified, all component memory dependencies are ignored. If there are loop-carried dependencies, your generated RTL produces incorrect results.</p> <p>The <code>safelen</code> parameter specifies the dependency distance. The dependency distance is the number of iterations between successive load/stores that depend on each other. It is safe to not include <code>safelen</code> is only when the dependence distance is infinite (that is, there are no real dependencies).</p> <p>Example:</p> <pre style="background-color: #f0f0f0;">#pragma ivdep safelen(2) for (int i = 0; i < 8; i++) { // Loop body }</pre> <p>To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/best_practices/loop_memory_dependency</code></p>

continued...



Pragma	Description
#pragma loop_coalesce <i>N</i>	<p>The compiler tries to fuse all loops nested within this loop into a single loop. This pragma accepts an optional value <i>N</i> which indicates the number of levels of loops to coalesce together.</p> <pre>#pragma loop_coalesce 2 for (int i = 0; i < 8; i++) { for (int j = 0; j < 8; j++) { // Loop body } }</pre>
#pragma unroll <i>N</i>	<p>This pragma unrolls the loop completely or by <i><N></i> times, where <i><N></i> is optional and is a positive integer value.</p> <p>Example:</p> <pre>#pragma unroll 8 for (int i = 0; i < 8; i++) { // Loop body }</pre> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/best_practices/resource_sharing_filter</p>
#pragma max_concurrency <i>N</i>	<p>This pragma limits the number of iterations of a loop that can simultaneously execute at any time.</p> <p>This pragma is useful mainly when component memory is duplicated to improve the throughput of the loop. This is mentioned in the details pane for the loop in the Loop Analysis pane of the high level design report (<code>report.html</code>).</p> <p>This can occur only when the scope of a component memory (through its declaration or access pattern) is limited to this loop. Adding this pragma can be used to reduce the area that the loop consumes at the cost of some throughput.</p> <p>Example:</p> <pre>// Without this pragma, // multiple copies // of the array "arr" #pragma max_concurrency 1 for (int i = 0; i < 8; i++) { int arr[1024]; // Loop body }</pre>

6.1. Loop Initiation Interval (`ii` Pragma)

The initiation interval, or II, is the number of clock cycles between the launch of successive loop iterations. Use the `ii` pragma to direct the Intel High Level Synthesis (HLS) Compiler to attempt to set the initiation interval (II) for the loop that follows the pragma declaration. If the compiler cannot achieve the specified II for the loop, then the compilation errors out.

For some loops in your component, specifying a higher II value with the `ii` pragma than the value the compiler chooses by default can increase the maximum operating frequency (f_{max}) of your component without a decrease in throughput.

A loop is a good candidate to have the `ii` pragma applied to it if the loop meets the following conditions:

- The loop is pipelined because the component is single-threaded.
- The loop is not critical to the throughput of your component.
- The running time of the loop is small compared to other loops it might contain.

To specify a loop initiation interval for a loop, specify the pragma before the loop as follows:

```
#pragma ii <desired_initiation_interval>
```

The *<desired_initiation_interval>* parameter is required and is an integer that specifies the number of clock cycles to wait between the beginning of execution of successive loop iterations.

Example

Consider a case where your component has two distinct, pipelineable loops: a short-running initialization loop that has a loop-carried dependence and a long-running loop that does the bulk of your processing. In this case, the compiler does not know that the initialization loop has a much smaller impact on the overall throughput of your design. If possible, the compiler attempts to pipeline both loops with an II of 1.

Because the initialization loop has a loop-carried dependence, it will have a feedback path in the generated hardware. To achieve an II with such a feedback path, some clock frequency might be sacrificed. Depending on the feedback path in the main loop, the rest of your design could have run at a higher operating frequency.

If you specify `#pragma ii 2` on the initialization loop, you tell the compiler that it can be less aggressive in optimizing II for this loop. Less aggressive optimization allows the compiler to pipeline the path limiting the f_{\max} and could allow your overall component design to achieve a higher f_{\max} .

The initialization loop takes longer to run with its new II. However, the decrease in the running time of the long-running loop due to higher f_{\max} compensates for the increased length in running time of the initialization loop.

6.2. Loop-Carried Dependencies (`ivdep` Pragma)

When compiling your components, the HLS compiler generates hardware to avoid any data hazards between load and store instructions. In particular, read-write dependencies can limit performance when they exist across loop iterations because they prevent the compiler from beginning a new loop iteration before the current iteration finishes executing its load and store instructions. You have the option to guarantee to the HLS compiler that there are no implicit dependencies across loop iterations in your component by adding the `ivdep` pragma in your code.

The `ivdep` pragma tells the compiler that a dependency between loop iterations can be ignored. Ignoring the dependency saves area and lowers the loop initiation interval (II) of the affected loop because the hardware required for avoiding data hazards is no longer required.

You can provide more information about loop dependencies by adding the `safelen(N)` clause to the `ivdep` pragma. The `safelen(N)` clause specifies the maximum number of consecutive loop iterations without loop-carried dependencies. For example, `#pragma ivdep safelen(32)` indicates to the compiler that there are a maximum of 32 iterations of the loop before loop-carried dependencies might be introduced. That is, while `#pragma ivdep` promises that there are no implicit memory dependency between any iteration of this loop, `#pragma safelen(32)` promises that the iteration that is 32 iterations away is the closest iteration that could be dependent on this iteration.



To specify that accesses to a particular memory array inside a loop will not cause loop-carried dependencies, add the line `#pragma ivdep array (array_name)` before the loop in your component code. The array specified by the `ivdep` pragma must be a local or private memory array, or a pointer variable that points to a global, local, or private memory storage. If the specified array is a pointer, the `ivdep` pragma also applies to all arrays that may alias with specified pointer. The array specified by the `ivdep` pragma can also be an array or a pointer member of a struct.

Caution: Incorrect usage of the `ivdep` pragma might introduce functional errors in hardware.

Use Case 1:

If all accesses to memory arrays inside a loop do not cause loop-carried dependencies, add `#pragma ivdep` before the loop.

```
1 // no loop-carried dependencies for A and B array accesses
2 #pragma ivdep
3 for(int i = 0; i < N; i++) {
4     A[i] = A[i + N];
5     B[i] = B[i + N];
6 }
```

Use Case 2:

You may specify `#pragma ivdep array (array_name)` on particular memory arrays instead of all array accesses. This pragma is applicable to arrays, pointers, or pointer members of structs. If the specified array is a pointer, the `ivdep` pragma applies to all arrays that may alias with the specified pointer.

```
1 // No loop-carried dependencies for A array accesses
2 // Compiler inserts hardware that reinforces dependency constraints for B
3 #pragma ivdep array(A)
4 for(int i = 0; i < N; i++) {
5     A[i] = A[i - X[i]];
6     B[i] = B[i - Y[i]];
7 }
8
9 // No loop-carried dependencies for array A inside struct
10 #pragma ivdep array(S.A)
11 for(int i = 0; i < N; i++) {
12     S.A[i] = S.A[i - X[i]];
13 }
14
15 // No loop-carried dependencies for array A inside the struct pointed by S
16 #pragma ivdep array(S->X[2][3].A)
17 for(int i = 0; i < N; i++) {
18     S->X[2][3].A[i] = S.A[i - X[i]];
19 }
20
21 // No loop-carried dependencies for A and B because ptr aliases
22 // with both arrays
23 int *ptr = select ? A : B;
24 #pragma ivdep array(ptr)
25 for(int i = 0; i < N; i++) {
26     A[i] = A[i - X[i]];
27     B[i] = B[i - Y[i]];
28 }
29
30 // No loop-carried dependencies for A because ptr only aliases with A
31 int *ptr = &A[10];
32 #pragma ivdep array(ptr)
33 for(int i = 0; i < N; i++) {
```

```
34     A[i] = A[i - X[i]];
35     B[i] = B[i - Y[i]];
36 }
```

6.3. Loop Coalescing (`loop_coalesce` Pragma)

Use the `loop_coalesce` pragma to direct the Intel High Level Synthesis (HLS) Compiler to coalesce nested loops into a single loop without affecting the loop functionality. Coalescing loops can help reduce your component area usage by directing the compiler to reduce the overhead needed for loop control.

Coalescing nested loops also reduces the latency of the component, which could further reduce your component area usage. However, in some cases, coalescing loops might lengthen the critical loop initiation interval path, so coalescing loops might not be suitable for all components.

To coalesce nested loops, specify the pragma as follows:

```
#pragma loop_coalesce <loop_nesting_level>
```

The `<loop_nesting_level>` parameter is optional and is an integer that specifies how many nested loop levels that you want the compiler to attempt to coalesce. If you do not specify the `<loop_nesting_level>` parameter, the compiler attempts to coalesce all of the nested loops.

For example, consider the following set of nested loops:

```
for (A)
  for (B)
    for (C)
      for (D)
        for (E)
```

If you place the pragma before loop (A), then the loop nesting level for these loops is defined as:

- Loop (A) has a loop nesting level of 1.
- Loop (B) has a loop nesting level of 2.
- Loop (C) has a loop nesting level of 3.
- Loop (D) has a loop nesting level of 4.
- Loop (E) has a loop nesting level of 3.

Depending on the loop nesting level that you specify, the compiler attempts to coalesce loops differently:

- If you specify `#pragma loop_coalesce 1` on loop (A), the compiler does not attempt to coalesce any of the nested loops.
- If you specify `#pragma loop_coalesce 2` on loop (A), the compiler attempts to coalesce loops (A) and (B).
- If you specify `#pragma loop_coalesce 3` on loop (A), the compiler attempts to coalesce loops (A), (B), (C), and (E).
- If you specify `#pragma loop_coalesce 4` on loop (A), the compiler attempts to coalesce all of the loops [loop (A) - loop (E)].



Example

The following simple example shows how the compiler coalesces two loops into a single loop.

Consider a simple nested loop written as follows:

```
#pragma loop_coalesce
for (int i = 0; i < N; i++)
  for (int j = 0; j < M; j++)
    sum[i][j] += i+j;
```

The compiler coalesces the two loops together so that they run as if they were a single loop written as follows:

```
int i = 0;
int j = 0;
while(i < N){

  sum[i][j] += i+j;
  j++;

  if (j == M){
    j = 0;
    i++;
  }
}
```

6.4. Loop Unrolling (unroll Pragma)

The Intel HLS Compiler supports the `unroll` pragma for unrolling multiple copies of a loop.

Example code:

```
1 #pragma unroll <N>
2 for (int i = 0; i < M; ++i) {
3     // Some useful work
4 }
```

In this example, N specifies the unroll factor, that is, the number of copies of the loop that the HLS compiler generates. If you do not specify an unroll factor, the HLS compiler unrolls the loop fully. You can find the unroll status of each loop in the high level design report (`report.html`).

6.5. Loop Concurrency (max_concurrency Pragma)

You can use the `max_concurrency` pragma to increase or limit the concurrency of a loop in your component. The concurrency of a loop is how many iterations of that loop can be in progress at one time. By default, the Intel HLS Compiler tries to maximize the concurrency of loops so that your component runs at peak throughput.

To achieve maximum concurrency in loops, sometimes local memory has to be duplicated to break dependencies on the underlying hardware that prevents the loop from being fully pipelined. You can see this in the Details pane Loop analysis report in your component HLD report (`report.html`) as a message that says that the maximum number of simultaneous executions has been limited to N . Duplicating local memory in this case is not the same as replicating memory in order to increase the number of ports.



If you want to exchange some performance for local memory savings, apply `#pragma max_concurrency <N>` to the loop. When you apply this pragma, the duplication factor changes and controls the number of threads entering the loop, as shown in the following example:

```
#pragma max_concurrency 1
for (int i = 0; i < N; i++) {
    int arr[M];
    // Doing work on arr
}
```

You can also control the concurrency of your component by using the `hls_max_concurrency(N)` component attribute. For more information about the `hls_max_concurrency(N)` component attribute, see [Concurrency Control \(hls_max_concurrency Attribute\)](#) on page 49.

7. Component Concurrency

The Intel HLS Compiler assumes that you want a fully pipelined datapath in your component. In the C++ implementation, you may think of a fully pipelined datapath as calling a function multiple times (for example, by multiple threads) before the first call has returned. The behavior of threads within the synthesized datapath is subject to the concurrency model, so the Intel HLS Compiler might not be able to deliver a component with a fully-pipelined datapath.

The Intel HLS Compiler provides you with the `hls_max_concurrency` component attribute to help you control the maximum concurrency of your component.

7.1. Serial Equivalence within a Memory Space or I/O

Within a single memory space or I/O, every call to the component (that is, every cycle where the `start` signal is asserted and the component holds the `busy` signal low) on the function protocol interface behaves as though the previous function call was fully executed.

When visualizing a single shared memory space, think of multiple function calls as executing sequentially, one after another. This way, when the component asserts the `done` signal, the results of a component invocation in hardware are guaranteed to be visible to both the next component invocation and the external system.

The HLS compiler leverages pipeline parallelism to execute component invocations and loop iterations in parallel if the associated dependencies allow for parallel execution. Because the HLS compiler generates hardware that keeps track of dependencies across component invocations, it can support pipeline parallelism while guaranteeing serial equivalence across memory spaces. Ordering between independent I/O instructions is not guaranteed.

7.2. Concurrency Control (`hls_max_concurrency` Attribute)

You can use the `hls_max_concurrency` component attribute to increase or limit the maximum concurrency of your component. The concurrency of a component is the number of invocations of the component that can be in progress at one time. By default, the Intel HLS Compiler tries to maximize concurrency so that the component runs at peak throughput.

You can control the maximum concurrency of your component by adding the `hls_max_concurrency` attribute immediately before you declare your component, as shown in the following example:

```
#include "HLS/hls.h"
hls_max_concurrency(3)
```

```
component void foo ( /* arguments */ ){  
    // Component code  
}
```

The Intel HLS Compiler sets the component concurrency to one in the following cases:

- The Intel HLS compiler does not automatically duplicate local memory to increase the throughput at the component level. If your component invocation uses a (non-static) local memory system that is used by a component invocation, the next invocation cannot start until the previous invocation has finished all of its accesses to and from that local memory. This limitation is shown in the Loop analysis report as load-store dependencies on the array. Adding the `hls_max_concurrency(N)` attribute on the component duplicates the local memory so that you can have multiple invocations of your component in progress at the same time.
- In some cases, the compiler reduces concurrency to save a great deal of area. In these cases, the `hls_max_concurrency(N)` attribute can increase the concurrency from 1.
- This attribute can also accept a value of 0. When this attribute is set to 0, the component should be able to accept new invocations as soon as the downstream datapath frees up. Only use this value when you see loop initiation interval (II) issues (such as extra bubbles) in your component, because using this attribute can increase the component area.

You can also control the concurrency of loops with the `max_concurrency(N)` pragma. For more information about the `max_concurrency(N)` pragma, see [Loop Concurrency \(max_concurrency Pragma\)](#) on page 47.

8. Intel HLS Compiler Libraries

The Intel HLS Compiler comes with templated libraries that help speed the development of your components by providing you with FPGA-optimized code for some commonly-used algorithms.

The Intel HLS Compiler provides the following libraries:

Library	Description	Header file
Random number generator	Generate random integers or floating point numbers that follow a uniform distribution, or random floating point numbers that follow a Gaussian distribution	HLS/rand_lib.h
Matrix multiplication	Multiply two 2-D matrices.	HLS/matrix_mult.h

8.1. Random Number Generator Library

Random number generators are used in applications like security algorithms, cryptography, and gaming. The random number generator library provided with the Intel HLS Compiler gives you FPGA-optimized random number generator template classes that you can add to your component without needing to write your own.

Table 18. Properties of Values That Can Be Generated by the Intel HLS Compiler Random Number Generator Library

Value distribution	Value type	Value range	Generation method
Uniform	Integer	$[-2^{31}, 2^{31}-1]$	Tausworthe Generator
	Floating point	$[0, 1)$ (non-inclusive)	Tausworthe Generator
Gaussian	Floating point	$[0, 1)$	Central limit theorem (CLT) (Default)
			Box-Muller

Header File

To include the random number generator library in your component, add the following line to your component:

```
#include "HLS/rand_lib.h"
```

The header file is self-documented. You can review the header file to learn how to use the random number generator library in your component.

Random Number Object Declarations

Declare random number objects in your components as follows. In all cases, specifying **<seed_value>** is optional.

- Uniform distribution integer random number

```
static RNG_Uniform<int> <object_name>(<seed_value>)
```

- Uniform distribution floating point random number

```
static RNG_Uniform<float> <object_name>(<seed_value>)
```

- Gaussian distribution floating point random number (CLT method)

```
static RNG_Gaussian<float> <object_name>(<seed_value>)
```

OR

```
static RNG_Gaussian<float, ihc::GAUSSIAN_CLT> <object_name>(<seed_value>)
```

- Gaussian distribution floating point random number (Box-Muller method)

```
static RNG_Gaussian<float, ihc::GAUSSIAN_BOX_MULLER>  
<object_name>(<seed_value>)
```

8.2. Matrix Multiplication Library

The matrix multiplication library provided with the Intel HLS Compiler gives you an FPGA-optimized templated library to perform matrix multiplication of two matrices stored in a 2-D array.

When you use the matrix multiplication library, you can affect the number of DSP blocks and RAM blocks by controlling the dot product vector size and the number of matrix elements read at one time. Increasing the dot product vector size can achieve better latency, but at the cost of using more DSP blocks and other FPGA resources.

Header File

To include the matrix multiplication library in your component, add the following line to your component:


```
#include "HLS/matrix_mult.h"
```

The header file is self-documented. You can review the header file to learn how to use the matrix multiplication library in your component.

Template Arguments

The matrix multiplication library multiplies two 2-D matrices, A and B. The resulting product is returned in a third matrix, C. The matrix multiplication library has the following template arguments:

<i>T</i>	The data type of the matrix elements (For example, int, float, long, double).
<i>t_rowsA</i>	The number of rows in matrix A.
<i>t_colsA</i>	The number of columns in matrix A. This value also the number of rows in matrix B.
<i>t_colsB</i>	The number of columns in matrix B.



<i>DOT_VEC_SIZE</i>	<p>The number of DSP blocks to use in a single computation. This value must be a factor of <code>t_colsA</code>.</p> <p>You can achieve better component latency by increasing this value. However, you use more FPGA area to achieve this. Keeping this value low lowers your FPGA resource usage, but increases the latency.</p>
<i>BLOCK_SIZE</i>	<p>The number of elements to read at one time from matrix A. The default value of <code>BLOCK_SIZE</code> is the value of <code>DOT_VEC_SIZE</code>. You can reduce this number if the bandwidth needed by matrix A is lower than the value of <code>DOT_VEC_SIZE</code>, but it must remain a factor of <code>DOT_VEC_SIZE</code>.</p>
<i>RUNNING_SUM_MULT_L</i>	<p>This parameter can be adjusted to try and improve the f_{MAX} of a component that uses this library. Review the header file for a detailed description of this argument and its effects.</p>

Tip:

When you compile a component that uses the matrix multiplication library, the Intel HLS Compiler issues the following warning:

```
Warning: default template arguments for a function template are a C++11
extension
[-Wc++11-extensions].
```

This warning is thrown if the compiler encounters a feature is part of the C++11 extension because not all C++11 extensions are supported by the Intel HLS Compiler. The C++11 extension used in the matrix multiplication library (the ability to specify default template arguments) is fully supported by the Intel HLS Compiler. You can safely ignore this compiler warning when you use the matrix multiplication library.

9. Document Revision History of the Intel HLS Compiler Reference Manual

Document Version	Intel Quartus Prime Version	Changes
2019.01.03	18.1.1	<ul style="list-style-type: none"> Fixed typos in table headings in Compiler-Defined Preprocessor Macros on page 12.
2018.12.24	18.1.1	<ul style="list-style-type: none"> Removed information about the "HLS/iostream" header file. The function provided by this header file is replaced by using the standard C++ iostream header and the HLS_SYNTHESIS macro. Added description of the HLS_SYNTHESIS macro to C and C++ Libraries on page 10.
2018.12.24	18.1	<ul style="list-style-type: none"> Updated Slave Interfaces on page 28 and Slave Memories on page 31 with information about slave memory reads and writes that come from outside of the component. Added information about conduit creation and address spaces to Avalon Memory-Mapped Master Interfaces on page 24.
2018.09.24	18.1	<ul style="list-style-type: none"> PRO The Intel HLS Compiler has a new front end. For a summary of the changes introduced by this new front end, see <i>Improved Intel HLS Compiler Front End</i> in the Intel High Level Synthesis Compiler Release Notes. PRO The <code>--promote-integers</code> flag and the <code>best_practices/integer_promotion</code> tutorial are no longer supported in Pro Edition because integer promotion is now done by default. The flag and tutorial are still supported in Standard Edition. Components invoked with the <code>hls_avalon_slave_component</code> argument must take slave or stable arguments. If the component arguments are not slave or stable arguments, compiling the component generates an error message. The description of the <code>hls_avalon_slave_component</code> argument in Component Invocation Interface Arguments on page 32 and Intel High Level Synthesis Compiler Quick Reference on page 59 now reflects that requirement. In Loops in Components on page 41, clarified the pragma statements that apply to loops must immediately precede the loop that the pragma applies to. In Declaring ac_int Datatypes in Your Component on page 14, added initialization requirement for <code>ac_int</code> variables larger than 64 bits. You must use <code>ac::init_array</code> constructors to initialize <code>ac_int</code> variables larger than 64 bits. In Static Variables on page 39, removed the restriction on applying memory attributes to file-scoped static variables. Both file-scoped and function-scoped static variables can have memory attributes applied to them.

continued...




Document Version	Intel Quartus Prime Version	Changes
2018.07.08	18.0	<ul style="list-style-type: none"> In Static Variables on page 39, highlighted paragraph that says that memory attributes applied to static variables work only if the static variable is declared within the component function. In Control and Status Register (CSR) Slave on page 28, corrected a typo. The sentence " You do not need to use the <code>hls_avalon_slave_component</code> attribute to use the <code>hls_avalon_slave_component</code> attribute" was corrected to say "You do not need to use the <code>hls_avalon_slave_component</code> attribute to use the <code>hls_avalon_slave_register_argument</code> attribute".
2018.05.07	18.0	<ul style="list-style-type: none"> Starting with Intel Quartus Prime Version 18.0, the features and devices supported by the Intel HLS Compiler depend on what edition of Intel Quartus Prime you have. Intel HLS Compiler publications now use icons to indicate content and features that apply only to a specific edition as follows: <ul style="list-style-type: none"> PRO Indicates that a feature or content applies only to the Intel HLS Compiler provided with Intel Quartus Prime Pro Edition. STD Indicates that a feature or content applies only to the Intel HLS Compiler provided with Intel Quartus Prime Standard Edition. PRO Corrected the code example in Table 29 on page 67 and Table 30 on page 69. The corrected line is <code>int x = a.tryRead(success);</code> (was <code>int x = a.tryRead(&success);</code>). PRO Added <code><quartus_installdir>/hls/examples/tutorials/interfaces/ explicit_streams_packets_empty</code> to list of tutorials in Table 29 on page 67 and Table 31 on page 69. PRO Added <code>ihc::firstSymbolInHighOrderBits</code> and <code>ihc::usesEmpty</code> to the list of stream interface declarations in Table 29 on page 67 and Table 31 on page 69. Also, revised the description of the <code>ihc::bitsPerSymbol</code> declaration to include the effect of the <code>ihc::firstSymbolInHighOrderBits</code> declaration. STD Added a footnote to the <code>-march MAX10</code> option in #unique_3/unique_3_Connect_42_section_N10130_N1001B_N10001 on page 5 about a prerequisite required before you synthesize your component IP for Intel MAX 10 devices. Added new topic AC Datatypes and Native Compilers on page 16 describing use of reference AC datatype headers with the Intel HLS Compiler. Intel HLS Compiler Libraries on page 51 added to document Intel HLS Compiler libraries. The following Intel HLS Compiler libraries were added: <ul style="list-style-type: none"> PRO Random Number Generator Library on page 51 PRO Matrix Multiplication Library on page 52
2017.12.22	17.1.1	<ul style="list-style-type: none"> Updated <code>hls_avalon_slave_memory_argument(N)</code> description in Slave Memories on page 31 to include the description that the parameter value <i>N</i> is the size of the memory in bytes. Updated Table 8 on page 15 and Table 35 on page 73 to indicate that the <code>ac_int</code> debug macros have the following restrictions: <ul style="list-style-type: none"> You must declare the macros in your code before you declare <code>#include HLS/ac_int.h</code>. The <code>ac_int</code> debugging tools work only for x86 emulation of your component.

continued...



Document Version	Intel Quartus Prime Version	Changes
		<ul style="list-style-type: none"> Updated <code>-march "<FPGA_family>"</code> options in Intel HLS Compiler Command Options on page 5 to include FPGA family options without a space. Revised the description of the <code>ihc::align</code> argument in Table 33 on page 72 in Intel High Level Synthesis Compiler Quick Reference on page 59. The same information also appears in Avalon Memory-Mapped Master Interfaces on page 24.
2017.11.06	17.1	<ul style="list-style-type: none"> Updated Intel HLS Compiler Command Options on page 5 as follows: <ul style="list-style-type: none"> Revised description of <code>-c i++</code> command option. Added descriptions of the <code>--x86-only</code> and <code>--fpga-only i++</code> command options. Updated Supported Math Functions on page 75 as follows: <ul style="list-style-type: none"> Noted that the <code>HLS/extendedmath.h</code> header file is supported only by the Intel HLS Compiler, not by the GCC or MSVC compilers. Added <code>popcount</code> to the list functions supported by the <code>HLS/extendedmath.h</code> header file. Expanded list of functions provided by <code>HLS/extendedmath.h</code> to explicitly list double-precision and single-precision floating point versions of the functions. Added a list of <code>popcount</code> function variations available for different data types. Updated Arbitrary Precision Math Support on page 13 to include restriction that the Intel arbitrary precision header files cannot be compiled with GCC. Added the <code>ihc::readwrite_mode</code> Avalon-MM interface to Avalon Memory-Mapped Master Interfaces on page 24 and Intel High Level Synthesis Compiler Quick Reference on page 59. Added the <code>ihc::waitrequest</code> Avalon-MM interface to Avalon Memory-Mapped Master Interfaces on page 24 and Intel High Level Synthesis Compiler Quick Reference on page 59. Added the <code>hls_stall_free_return</code> macro and <code>stall_free_return</code> attribute to Unstable and Stable Component Arguments on page 33 and Intel High Level Synthesis Compiler Quick Reference on page 59. Reorganized the overall structure of the book, breaking up chapter 1 into smaller chapters and changing the order of the chapters. Updated mentions of the HLS or <code>i++</code> installation directory to use the Intel Quartus Prime Design Suite installation directory as the starting point. Moved the following content to Intel High Level Synthesis Compiler Best Practices Guide: <ul style="list-style-type: none"> Moved "Avoid Pointer Aliasing" section to "Avoid Pointer Aliasing".
2017.06.23	—	<ul style="list-style-type: none"> Updated Static Variables on page 39 to add information about static variable initialization and how to control it. Minor changes and corrections.
2017.06.09	—	<ul style="list-style-type: none"> Revised Declaring ac_int Datatypes in Your Component on page 14 for changes in how to include <code>ac_int.h</code>. Revised Arbitrary Precision Math Support on page 13 to clarify support for Algorithmic C datatypes. Removed all mentions of <code>--device</code> compiler option. This option has been replaced by the changed function of the <code>-march</code> compiler option. See #unique_3/unique_3_Connect_42_section_N10130_N1001B_N10001 on page 5 for details about the changed function of the <code>-march</code> compiler option. Updated the generated C header file for the component <code>mycomp_xyz</code> in Control and Status Register (CSR) Slave on page 28.
<i>continued...</i>		



Document Version	Intel Quartus Prime Version	Changes
		<ul style="list-style-type: none"> Added information about structs in component interfaces to Component Interfaces on page 17. Revised C and C++ Libraries on page 10 with updates to iostream behavior. Added information about math functions supported by <code>extendedmath.h</code> header file to Supported Math Functions on page 75.
2017.02.03	—	<ul style="list-style-type: none"> In <i>Scalar Parameters and Avalon Streaming Interfaces</i>, updated information in the <i>Available Scalar Parameters for Avalon-ST Interfaces</i> table. In <i>Pointer Parameters, Reference Parameters, and Avalon Memory-Mapped Master Interfaces</i>, updated information in the <i>Available Template Arguments for Configuration of the Avalon-MM Interface</i> table. Added new information to <i>Global Variables</i> about area usage and optimizing for global constants, pointers, and variables.
2016.11.30	—	<ul style="list-style-type: none"> In <i>HLS Compiler Command Options</i>, modified the table <i>Command Options that Customize Compilation</i> in the following manner: <ul style="list-style-type: none"> Removed the <code>--rtl-only</code> command option and its description because it is no longer in use. Added the <code>--simulator <name></code> command option and its description. Remove the <code>-g</code> command option because the HLS compiler now generates debug information in reports by default for both Windows and Linux. In addition, debug data is available by default in final binaries for Linux. In <i>Pointer Parameters, Reference Parameters, and Avalon Memory-Mapped Master Interfaces</i>, added information on the <code>altera::align<value></code> template argument in the table. Added the topics <i>Memory-Mapped Test Bench Constructor</i> and <i>Implicit and Explicit Examples of Creating a Memory-Mapped Master Test Bench</i>. In <i>Usage Examples of Component Invocation Protocol Macros</i>, replaced component invocation protocol attributes in the code examples with their corresponding macros. Added the line <code>#include "HLS/hls.h"</code> to the code snippets in the following sections: <ul style="list-style-type: none"> <i>Usage Examples of Interface Synthesis Macros</i> <i>Usage Examples of Component Invocation Protocol Macros</i> Added the topic <i>Arbitrary Precision Integer Support</i> to introduce the <code>ac_int</code> datatype and the Intel-provided <code>ac_int.h</code> header file. Included the following subtopics: <ul style="list-style-type: none"> <i>Defining the ac_int Datatype in Your Component for Arbitrary Precision Integer Support</i> <i>Important Usage Information on the ac_int Datatype</i> Updated the content in <i>Area Minimization and Control of On-Chip Memory Architecture</i>: <ul style="list-style-type: none"> Replaced the <code>numreadports(n)</code> and <code>numwriteports(n)</code> entries the <i>Attributes for Controlling On-Chip Memory Architecture</i> table with a single <code>numports_readonly_writeonly(m,n)</code> entry. Added information on the <code>hls_simple_dual_port_memory</code> macro. Added information on the <code>hls_merge ("label", "direction")</code> and the <code>hls_bankbits(b0, b1, ..., bn)</code> attributes. Added example use cases for the <code>hls_merge("label", "direction")</code> and the <code>hls_bankbits(b0, b1, ..., bn)</code> attributes.


continued...



Document Version	Intel Quartus Prime Version	Changes
		<ul style="list-style-type: none">Added the topic <i>Relationship between hls_bankbits Specifications and Memory Address Bits</i> to explain the derivation of a memory address in the presence of the <code>hls_bankbits</code> and <code>hls_bankwidth</code> attributes.
2016.09.12	—	Initial release.

A. Intel High Level Synthesis Compiler Quick Reference

Table 19. i++ Command Line Arguments

Feature	Default Value	Description
General i++ command options		
--debug-log		Generate the compiler diagnostics log.
-h, --help		List compiler command options along with brief descriptions.
-oresult		Place compiler output into the <result> executable and the <result>.prj directory.
-v		Display messages describing the progress of the compilation.
--version		Display compiler version information.
Command options affecting compilation		
-c		Preprocess, parse, and generate object files.
--component component_name		Comma-separated list of function names to synthesize to RTL. To use this option, your component must be configured with C-linkage using the extern "C" specification. For example: <pre>extern "C" int myComponent(int a, int b)</pre> Using the component function attribute is preferred over using the --component command option to indicate functions that you want the compiler to synthesize.
-Dmacro [=val]		Define a <macro> with <val/> as its value.
-g		Generate debug information (default option).
-g0		Do not generate debug information.
-Idir		Add directory <dir> to the end of the main include path.
-march=[x86-64 FPGA_family FPGA_part_number]	x86-64	Generate code for an emulator flow (x86-64) or for the specified FPGA family or FPGA part number.
 --promote-integers		Use extra FPGA resources to mimic g++ integer promotion. In Pro Edition, the compiler always promotes integers for standard types. Use the ac_int datatypes if you want smaller (or larger) datatypes. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/best_practices/integer_promotion
--quartus-compile		Run the HDL generated through Intel Quartus Prime.
--simulator simulator_name	modelsim	Specifies the simulator you are using to perform verification. This command option can take the following values for <simulator_name>:
continued...		

Intel Corporation. All rights reserved. Intel, the Intel logo, Altera, Arria, Cyclone, Enpirion, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

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



Feature	Default Value	Description
		<p><code>modelsim</code> Use ModelSim for component verification.</p> <p><code>none</code> Disable verification. That is, generate RTL for components without the test bench.</p> <p>If you do not specify this option, <code>--simulator modelsim</code> is assumed.</p>
Command options affecting linking		
<code>--clock <i>clock target</i></code>	240 MHz	<p>Optimize the RTL for the specified clock frequency or period.</p> <p>For example:</p> <pre>i++ -march="Arria 10" test.cpp --clock 100MHz i++ -march="Arria 10" test.cpp --clock 10ns</pre>
<code>--fp-relaxed</code>		<p>Relax the order of floating point arithmetic operations.</p> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/best_practices/floating_point_ops</p>
<code>--fpc</code>		<p>Remove intermediate rounding and conversion when possible.</p> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/best_practices/floating_point_ops</p>
<code>-ghdl</code>		Enable full debug visibility and logging of all HDL signals in simulation.
<code>-Ldir</code>		(Linux only) Add directory <code><dir></code> to the list of directories to be searched for <code>-l</code> .
<code>-llibrary</code>		(Linux only) Search the library name <code><library></code> when linking.
<code>--x86-only</code>		Create only the testbench executable (<code><result>.out/<result>.exe</code>).
<code>--fpga-only</code>		Create only the <code><result>.prj</code> directory and its contents.

Table 20. Intel High Level Synthesis (HLS) Compiler Header Files

Feature	Description
<code>#include "HLS/hls.h"</code>	Required for component identification and component parameter interfaces.
<code>#include "HLS/math.h"</code>	<p>Includes FPGA-specific definitions for the math functions from the <code>math.h</code> for your operating system.</p> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/best_practices/single_vs_double_precision_math</p>
<code>#include "HLS/extendedmath.h"</code>	<p>Includes additional FPGA-specific definitions of math functions not in <code>math.h</code>.</p> <p>To learn more, review the design: <quartus_installdir>/hls/examples/QRD</p>
<code>#include "HLS/ac_int.h"</code>	<p>Intel HLS Compiler version of <code>ac_int</code> header file.</p> <p>Provides arbitrary width integer support.</p> <p>To learn more, review the following tutorials:</p> <ul style="list-style-type: none"> <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_basic_ops <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow <quartus_installdir>/hls/examples/tutorials/best_practices/struct_interfaces
<code>#include "HLS/ac_fixed.h"</code>	<p>Intel HLS Compiler version of the <code>ac_fixed</code> header file.</p> <p>Provides arbitrary precision fixed point support.</p>

continued...



Feature	Description
	To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_constructor</code>
<code>#include "HLS/ac_fixed_math.h"</code>	Intel HLS Compiler version of the <code>ac_fixed_math</code> header file. Provides arbitrary precision fixed point math functions. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_math_library</code>
<code>#include "HLS/stdio.h"</code>	Provides <code>printf</code> support for components so that <code>printf</code> statements work in x86 emulations, but are disabled in component when compiling to an FPGA architecture.
<code>#include <iostream></code>	To use the C++ standard output streams (<code>cout</code> and <code>cerr</code>) provided by the standard <code><iostream></code> header, you must guard any standard output statements with the <code>HLS_SYNTHESIS</code> macro. This macro ensures that statements in a component work in x86 emulations but are disabled in the component when compiling to an FPGA architecture.

Table 21. Intel HLS Compiler Keywords

Feature	Description
<code>component</code>	Indicates that a function is a component. Example: <pre>component void foo()</pre>

Table 22. Intel HLS Compiler Simulation API (Testbench only)

Function	Description
<code>ihc_hls_enqueue(void* retptr, void* funcptr, /*function arguments*/)</code>	This function enqueues one invocation of an HLS component. The return value is stored in the first argument which should be a pointer to the return type. The component is not run until the <code>ihc_hls_component_run_all()</code> is invoked. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/usability/enqueue_call</code>
<code>ihc_hls_enqueue_noret(void* funcptr, /*function arguments*/)</code>	This function enqueues one invocation of an HLS component. This function should be used when the return type of the HLS component is void. The component is not run until the <code>ihc_hls_component_run_all()</code> is invoked. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/usability/enqueue_call</code>
<code>ihc_hls_component_run_all(void* funcptr)</code>	This function accepts a pointer to the HLS component function. When run, all enqueued invocations of the component will be pushed into the component in the HDL simulator as quickly as the component can accept new invocations. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/usability/enqueue_call</code>
<code>int ihc_hls_sim_reset(void)</code>	This function sends a reset signal to the component during automated simulation. It returns 1 if the reset was exercised or 0 otherwise. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/component_memories/static_var_init</code>
<pre>component int foo(int val) { // function definition } component void bar (int val) { // function definition } int main() { // int input = 0; int res[5];</pre>	
continued...	



Function	Description
<pre> ihc_hls_enqueue(&res, &foo, input); ihc_hls_enqueue_noret(&bar, input); input = 1; ihc_hls_enqueue(&res, &foo, input); ihc_hls_enqueue_noret(&bar, input); ihc_hls_component_run_all(&foo); ihc_hls_component_run_all(&bar); } </pre>	

Table 23. Intel HLS Compiler Component Memory Attributes

Attribute	Default Value	Description
hls_register	Based on the memory access pattern inferred by the compiler.	Forces a variable or array inside component to be implemented as registers. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/best_practices/swap_vs_copy
hls_memory	Based on the memory access pattern inferred by the compiler.	Forces a variable or array inside component to be implemented as embedded memory. To learn more, review the design: <quartus_installdir>/hls/examples/QRD
hls_singlepump	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must be single pumped. That is, the memory is clocked at the same operating frequency as the operating frequency of your component. To learn more, review the design: <quartus_installdir>/hls/examples/QRD
hls_doublepump	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must be double pumped. That is, the memory is clocked at twice the operating frequency of your component.
hls_numbanks(<i>N</i>)(*)	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must have <i>N</i> banks, where <i>N</i> is a power-of-two constant number.
hls_bankwidth(<i>N</i>)(*)	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must have banks that are <i>N</i> bytes wide, where <i>N</i> is a power-of-two constant number. To learn more, review the design: <quartus_installdir>/hls/tutorials/component_memories/bank_bits
hls_bankbits(<i>b</i> ₀ , <i>b</i> ₁ , ..., <i>b</i> _{<i>n</i>})(*)	Lowest bits of the address based on number of banks.	Forces the memory system to split into 2 ^{<i>n</i>+1} banks, with { <i>b</i> ₀ , <i>b</i> ₁ , ..., <i>b</i> _{<i>n</i>} } forming the bank-select bits. <i>Important:</i> <i>b</i> ₀ , <i>b</i> ₁ , ..., <i>b</i> _{<i>n</i>} must be consecutive, positive integers. You can specify the consecutive, positive integers in ascending or descending order. If you do not specify the hls_bankwidth(<i>N</i>) attribute along with this attribute, then <i>b</i> ₀ , <i>b</i> ₁ , ..., <i>b</i> _{<i>n</i>} are mapped to array index bits 0 to <i>n</i> -1 in the memory bank implementation. To learn more, review the design: <quartus_installdir>/hls/tutorials/component_memories/bank_bits

continued...

(*) This attribute is subject to constraints outlined in [Constraints on Attributes for Memory Banks](#) on page 39.



Attribute	Default Value	Description
<code>hls_numports_readonly_writeonly(M, N)</code>	Based on the memory access pattern inferred by the compiler.	Specifies that the memory implementing the local variable must have M read ports and N write ports, where M and N are constant numbers greater than zero.
<code>hls_simple_dual_port_memory</code>		Specifies the configuration that is defined by the presence of both the <code>hls_singlepump</code> and the <code>hls_numports_readonly_writeonly(1,1)</code> macros.
<code>hls_merge("mem_name", "depth")</code>		Allows merging two or more local variables to be implemented in component memory as a single merged memory system in a depth-wise manner. All variables with same <code><mem_name></code> label specified in their <code>hls_merge</code> attribute are merged into the same memory system. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/best_practices/depth_wise_merge</code>
<code>hls_merge("mem_name", "width")</code>		Allows merging two or more local variables to be implemented in component memory as a single merged memory system in a width-wise manner. All variables with same <code><mem_name></code> label specified in their <code>hls_merge</code> attribute are merged into the same memory system. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/best_practices/width_wise_merge</code>
<code>hls_init_on_reset</code>	Default behavior for static variables.	Forces the static variables inside the component to be reset when the component <code>reset</code> signal is asserted. This requires the an additional write port to the component memory implemented and can increase the power-up latency when the component is reset. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/component_memories/static_var_init</code>
<code>hls_init_on_powerup</code>		Sets the component memory implementing the static variable to set on power-up when the FPGA is programmed. When the component is reset, the component memory is not reset back to the initialized value of the static. To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/component_memories/ static_var_init</code>

Table 24. Intel HLS Compiler Loop Pragmas

Pragma	Description
<code>#pragma ii N</code>	<p>Forces the loop that this is applied on to have a loop initiation interval (II) of <code><N></code>, where <code><N></code> is a positive integer value.</p> <p>This can have an adverse effect on the f_{MAX} of your component because using this pragma to get a lower loop II combines pipeline stages together and creates logic with a long propagation delay.</p> <p>Using this pragma with a larger loop II inserts more pipeline stages and can give you a better component f_{MAX} value.</p> <p>Example:</p> <pre>#pragma ii 2 for (int i = 0; i < 8; i++) { // Loop body }</pre>
<i>continued...</i>	

Pragma	Description
<pre>#pragma ivdep safelen(N) array(array_name)</pre>	<p>Tells the compiler to ignore memory dependencies between iterations of this loop.</p> <p>It can accept an optional argument that specifies the name of the array. If <code>array</code> is not specified, all component memory dependencies are ignored. If there are loop-carried dependencies, your generated RTL produces incorrect results.</p> <p>The <code>safelen</code> parameter specifies the dependency distance. The dependency distance is the number of iterations between successive load/stores that depend on each other. It is safe to not include <code>safelen</code> is only when the dependence distance is infinite (that is, there are no real dependencies).</p> <p>Example:</p> <pre>#pragma ivdep safelen(2) for (int i = 0; i < 8; i++) { // Loop body }</pre> <p>To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/best_practices/loop_memory_dependency</code></p>
<pre>#pragma loop_coalesce N</pre>	<p>The compiler tries to fuse all loops nested within this loop into a single loop. This pragma accepts an optional value <code>N</code> which indicates the number of levels of loops to coalesce together.</p> <pre>#pragma loop_coalesce 2 for (int i = 0; i < 8; i++) { for (int j = 0; j < 8; j++) { // Loop body } }</pre>
<pre>#pragma unroll N</pre>	<p>This pragma unrolls the loop completely or by <code><N></code> times, where <code><N></code> is optional and is a positive integer value.</p> <p>Example:</p> <pre>#pragma unroll 8 for (int i = 0; i < 8; i++) { // Loop body }</pre> <p>To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/best_practices/resource_sharing_filter</code></p>
<pre>#pragma max_concurrency N</pre>	<p>This pragma limits the number of iterations of a loop that can simultaneously execute at any time.</p> <p>This pragma is useful mainly when component memory is duplicated to improve the throughput of the loop. This is mentioned in the details pane for the loop in the Loop Analysis pane of the high level design report (<code>report.html</code>).</p> <p>This can occur only when the scope of a component memory (through its declaration or access pattern) is limited to this loop. Adding this pragma can be used to reduce the area that the loop consumes at the cost of some throughput.</p> <p>Example:</p> <pre>// Without this pragma, // multiple copies // of the array "arr" #pragma max_concurrency 1 for (int i = 0; i < 8; i++) { int arr[1024]; // Loop body }</pre>

Table 25. Intel HLS Compiler Component Attributes

Feature	Description
<pre>hls_max_concurrency(N)</pre>	<p>In some cases, the concurrency of a component is limited to 1. This limit occurs when the generated hardware cannot be shared across component invocations. For example, when using local memories for a non-static variable.</p>



Feature	Description
	<p>You can use this attribute to request more copies of the local memory so that the component can run multiple invocations in parallel.</p> <p>This attribute can accept any non-negative whole number, including 0.</p> <p><i>Value greater than 0</i> A value greater than 0 indicates how many copies of the local memory to instantiate as well as how many component invocations can be in flight at once.</p> <p><i>Value equal to 0</i> Setting <code>hls_max_concurrency</code> to a value of 0 is useful in cases when there is no local memory but the component still has a poor dynamic loop initiation interval (II) even if you believe your component II should be 1. You can review the II for loops in your component in the high level design report.</p> <p>Example:</p> <pre> hls_max_concurrency(2) void foo(ihc::stream_in<int> &data_in, ihc::stream_out<int> &data_out) { int arr[N]; for (int i = 0; i < N; i++) { arr[i] = data_in.read(); } // Operate on the data and modify in place for (int i = 0; i < N; i++) { data_out.write(arr[i]); } } </pre> <p>To learn more, review the design example: <code><quartus_installdir>/hls/examples/inter_decim_filter</code></p>

Table 26. Intel HLS Compiler Default Interfaces

Feature	Description
Component invocation interface (component call and return)	The component call is implemented as an interface consisting of the component <code>start</code> and <code>busy</code> conduits. The component return is also implemented as an interface that includes the component <code>done</code> and <code>stall</code> signals.
Scalar parameter interface (passed by value)	Scalar parameters are implemented as input conduits that are synchronized with the component invocation interface.
Pointer parameter interface (passed by reference)	Pointer parameters are implemented as an implicit Avalon Memory-Mapped Master (<code>mm_master</code>) interface with the default parametrization. By default, the base address is treated as a scalar parameter so it is implemented as a conduit that is synchronized to the component invocation interface. A memory mapped interface is also exposed on the component.

Table 27. Intel HLS Compiler Component Invocation Interface Arguments

Feature	Description
<code>hls_avalon_streaming_component</code> This is the default component invocation interface.	<p>This attribute follows the Avalon-ST protocol for both the function call and the return streams. The component consumes the unstable arguments when the <code>start</code> signal is asserted and the <code>busy</code> signal is deasserted. The component produces the return data when the <code>done</code> signal is asserted.</p> <p>Top-level module ports:</p> <p>Function call—<code>start</code>, <code>busy</code></p> <p>Function return—<code>done</code>, <code>stall</code></p>

continued...



Feature	Description
	<p>Example:</p> <pre>component hls_avalon_streaming_component void foo(/*component arguments*/)</pre>
hls_avalon_slave_component	<p>The start, done, and returndata (if applicable) signals are registered in the component slave memory map. These component must take either slave, stream, or stable arguments. If you do not specify these types of arguments, the compiler generates an error message when you compile this component.</p> <p>Top-level module ports: Avalon-MM slave interface and irq_done signal</p> <p>Example:</p> <pre>component hls_avalon_slave_component void foo(/*component arguments*/)</pre> <p>To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/interfaces/mm_slaves</code></p>
hls_always_run_component	<p>The start signal is tied to 1 internally in the component. There is no done signal output. The control logic is optimized away when Intel Quartus Prime compiles the generated RTL for your FPGA.</p> <p>Use this protocol when the component datapath relies only on explicit streams for data input and output.</p> <p>IP verification does not support components with this component invocation protocol.</p> <p>Top-level module ports: None</p> <p>Example:</p> <pre>component hls_always_run_component void foo(/*component arguments*/)</pre>
hls_stall_free_return	<p>If the downstream component never stalls, the stall signal is removed by internally setting it to 0.</p> <p>This feature can be used with the hls_avalon_streaming_component, hls_avalon_slave_component, and hls_always_run_component arguments. This attribute can be used to specify that the downstream component is stall free.</p> <p>Example:</p> <pre>component hls_stall_free_return component int dut(int a, int b) { return a * b;}</pre>

Table 28. Intel HLS Compiler Component Argument Macros


Feature	Description
hls_conduit_argument This is the default interface for scalar arguments.	<p>The compiler implements the argument as an input conduit that is synchronous to the component's call (start and busy).</p> <p>Example:</p> <pre>component void foo(hls_conduit_argument int b)</pre>
hls_avalon_slave_register_argument	<p>The compiler implements the argument as a register that can be read from and written to over an Avalon-MM slave interface. The argument will be read into the component's pipeline, similar to the conduit implementation. The implementation is synchronous to the start and busy interface.</p> <p>Changes to the value of this argument made by the component data path will not be reflected on this register.</p> <p>Example:</p> <pre>component void foo(hls_avalon_slave_register_argument int b)</pre> <p>To learn more, review the tutorial: <code><quartus_installdir>/hls/examples/tutorials/interfaces/mm_slaves</code></p>

continued...




Feature	Description
hls_avalon_slave_memory_argument(<i>N</i>)	<p>The compiler implements the argument, where <i>N</i> specifies the size of the memory in bytes, in on-chip memory blocks, which can be read from or written to over a dedicated slave interface. The generated memory has the same architectural optimizations as all other internal component memories (that is, banking, coalescing, etc.).</p> <p>If the compiler performs static coalescing optimizations, the slave interface's data width will be the coalesced width. This attribute applies only to a pointer argument.</p> <p>Example:</p> <pre>component void foo(hls_avalon_slave_memory_argument(128*sizeof(int)) int *a)</pre> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/interfaces/mm_slaves</p>
hls_stable_argument	<p>A stable argument is an argument that does not change while there is live data in the component (that is, between pipelined function invocations).</p> <p>Changing a stable argument during component execution results in undefined behavior; each use of the stable argument might be the old value or the new value, but with no guarantee of consistency. The same variable in the same invocation can appear with multiple values.</p> <p>Using stable arguments, where appropriate, might save a significant number of registers in a design.</p> <p>Stable arguments can be used with conduits, mm_master interfaces, and slave_registers.</p> <p>Example:</p> <pre>component int dut(hls_stable_argument int a, hls_stable_argument int b) { return a * b;}</pre> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/interfaces/stable_arguments</p>

Table 29. Intel HLS Compiler Streaming Input Interface `stream_in` Declarations

Feature	Valid Values	Default Value	Description
<code>ihc::stream_in<datatype, template arguments></code>	Any valid C++ datatype		<p>Streaming input interface to the component.</p> <p>The width of the stream data bus is equal to a width of <code>sizeof(datatype)</code>.</p> <p>The testbench must populate this buffer (stream) fully before the component can start to read from the buffer.</p> <p>To learn more, review the following tutorials:</p> <ul style="list-style-type: none"> <quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_buffer  <quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packets_empty <quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packet_ready_valid <quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_ready_latency <quartus_installdir>/hls/examples/tutorials/interfaces/multiple_stream_call_sites
<code>ihc::buffer<value></code>	Non-negative integer value	0	<p>The capacity, in words, of the FIFO buffer on the input data that associates with the stream.</p> <p>This parameter is available only on input streams.</p>

continued...

Feature	Valid Values	Default Value	Description
<code>ihc::readyLatency<value></code>	Non-negative integer value (between 0-8)	0	The number of cycles between when the <code>ready</code> signal is deasserted and when the input stream can no longer accept new inputs.
<code>ihc::bitsPerSymbol<value></code>	A positive integer value that evenly divides by the data type size	Datatype size	Describes how the data is broken into symbols on the data bus. PRO Data is broken down according to how you set the <code>ihc::firstSymbolInHighOrderBits</code> declaration. By default, data is broken down in little endian order. STD Data is always broken down in little endian order.
PRO <code>ihc::firstSymbolInHighOrderBits<value></code>	true or false	false	Specifies whether the data symbols in the stream are in big endian order. 
<code>ihc::usesPackets<value></code>	true or false	false	Exposes the <code>startofpacket</code> and <code>endofpacket</code> sideband signals on the stream interface, which can be accessed by the packet based reads/writes
PRO <code>ihc::usesEmpty<value></code>	true or false	false	Exposes the <code>empty</code> out-of-band signal on the stream interface. Use this declaration only with streams that read more than one data symbol per clock cycle. The <code>empty</code> signal indicates the number of symbols on the data bus that do not represent valid data during the final stream read of a packet. You can control whether the <code>empty</code> symbols are in the low-order bits or high-order bits with the <code>ihc::firstSymbolInHighOrderBits</code> declaration.
<code>ihc::usesValid<value></code>	true or false	true	Controls whether a <code>valid</code> signal is present on the stream interface. If <code>false</code> , the upstream source must provide valid data on every cycle that <code>ready</code> is asserted. This is equivalent to changing the stream read calls to <code>tryRead</code> and assuming that success is always true. If set to <code>false</code> , <code>buffer</code> and <code>readyLatency</code> must be 0.
<p>The following code example illustrates both <code>stream_in</code> declarations and <code>stream_in</code> function APIs.</p> <pre> // Blocking read void foo (ihc::stream_in<int> &a) { int x = a.read(); } // Non-blocking read void foo_nb (ihc::stream_in<int> &a) { bool success = false; int x = a.tryRead(success); </pre>			




Feature	Valid Values	Default Value	Description
<pre> if (success) { // x is valid } } int main() { ihc::stream_in<int> a; ihc::stream_in<int> b; for (int i = 0; i < 10; i++) { a.write(i); b.write(i); } foo(a); foo_nb(b); } </pre>			

Table 30. Intel HLS Compiler Streaming Input Interface `stream_in` Function APIs

Feature	Description
<code>T read()</code>	Blocking read call to be used from within the component
<code>T read(bool& sop, bool& eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking read with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>T read(bool& sop, bool& eop, int& empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking read with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.
<code>T tryRead(bool &success)</code>	Non-blocking read call to be used from within the component. The <code>success</code> bool is set to true if the read was valid. If you use <code>tryRead</code> , your x86-64 results for your component might not match your FPGA results because emulation does not model the hardware behavior of blocking and non-blocking reads.
<code>T tryRead(bool& success, bool& sop, bool& eop)</code>	Available only if <code>usesPackets<true></code> is set. Non-blocking read with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>T tryRead(bool& success, bool& sop, bool& eop, int& empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Non-blocking read with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>emptysignals</code> .
<code>void write(T data)</code>	Blocking write call to be used from the testbench to populate the FIFO to be send to the component
<code>void write(T data, bool sop, bool eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking write call with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>void write(T data, bool sop, bool eop, int empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking write call with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.

Table 31. Intel HLS Compiler Streaming Output Interfaces `stream_out` Declaration

Feature	Valid Values	Default Value	Description
<code>ihc::stream_out<datatype, template arguments></code>	Any valid POD (plain old data) C++ datatype		Streaming output interface from the component. The testbench can read from this buffer once the component returns. To learn more, review the following tutorials: continued...



Feature	Valid Values	Default Value	Description
			<ul style="list-style-type: none"> • <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_buffer</code> • PRO <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packets_empty</code> • <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_packet_ready_valid</code> • <code><quartus_installdir>/hls/examples/tutorials/interfaces/explicit_streams_ready_latency</code> • <code><quartus_installdir>/hls/examples/tutorials/interfaces/multiple_stream_call_sites</code>
<code>ihc::readylatency<value></code>	Non-negative integer value (between 0-8)	0	<p>The number of cycles between when the <code>ready</code> signal is deasserted and when the input stream can no longer accept new inputs.</p> <p>Conceptually, you can view this parameter as an almost ready latency on the input FIFO buffer for the data that associates with the stream.</p>
<code>ihc::bitsPerSymbol<value></code>	Positive integer value that evenly divides the data type size	Datatype size	<p>Describes how the data is broken into symbols on the data bus.</p> <p>PRO Data is broken down according to how you set the <code>ihc::firstSymbolInHighOrderBits</code> declaration. By default, data is broken down in little endian order.</p> <p>STD Data is always broken down in little endian order.</p>
PRO <code>ihc::firstSymbolInHighOrderBits<value></code>	true or false	false	Specifies whether the data symbols in the stream are in big endian order.
<code>ihc::usesPackets<value></code>	true or false	false	Exposes the <code>startofpacket</code> and <code>endofpacket</code> sideband signals on the stream interface, which can be accessed by the packet based reads/writes.
PRO <code>ihc::usesEmpty<value></code>	true or false	false	<p>Exposes the <code>empty</code> out-of-band signal on the stream interface.</p> <p>Use this declaration only with streams that write more than one data symbol per clock cycle.</p> <p>The <code>empty</code> signal indicates the number of symbols on the data bus that do not represent valid data during the final stream write of a packet.</p> <p>You can control whether the <code>empty</code> symbols are in the low-order bits or high-order bits with the <code>ihc::firstSymbolInHighOrderBits</code> declaration.</p>
<code>ihc::usesReady<value></code>	true or false	true	<p>Controls whether a <code>ready</code> signal is present. If <code>false</code>, the downstream sink must be able to accept data on every cycle that valid is asserted. This is equivalent to changing the stream read calls to <code>tryWrite</code> and assuming that <code>success</code> is always <code>true</code>.</p> <p>If set to <code>false</code>, <code>readyLatency</code> must be 0.</p>
continued...			



Feature	Valid Values	Default Value	Description
<p>The following code example illustrates both <code>stream_out</code> declarations and <code>stream_out</code> function APIs.</p> <pre> // Blocking write void foo (ihc::stream_out<int> &a) { static int count = 0; for(int idx = 0; idx < 5; idx++){ a.write(count++); // Blocking write } } // Non-blocking write void foo_nb (ihc::stream_out<int> &a) { static int count = 0; for(int idx = 0; idx < 5; idx++){ bool success = a.tryWrite(count++); // Non-blocking write if (success) { // write was successful } } } int main() { ihc::stream_out<int> a; foo(a); // or foo_nb(a); // copy output to an array int outputData[5]; for (int i = 0; i < 5; i++) { outputData[idx] = a.read(); } } </pre>			

Table 32. Intel HLS Compiler Streaming Output Interfaces `stream_out` Function Call APIs

Feature	Description
<code>void write(T data)</code>	Blocking write call from the component
<code>void write(T data, bool sop, bool eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking write with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>void write(T data, bool sop, bool eop, int empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking write with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.
<code>bool tryWrite(T data)</code>	Non-blocking write call from the component. The return value represents whether the write was successful.
<code>bool tryWrite(T data, bool sop, bool eop)</code>	Available only if <code>usesPackets<true></code> is set. Non-blocking write with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals. The return value represents whether the write was successful.
PRO <code>bool tryWrite(T data, bool sop, bool eop, int empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Non-blocking write with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals. The return value represents whether the write was successful.
<code>T read()</code>	Blocking read call to be used from the testbench to read back the data from the component
<code>T read(bool &sop, bool &eop)</code>	Available only if <code>usesPackets<true></code> is set. Blocking read call to be used from the testbench to read back the data from the component with out-of-band <code>startofpacket</code> and <code>endofpacket</code> signals.
PRO <code>T read(bool &sop, bool &eop, int &empty)</code>	Available only if <code>usesPackets<true></code> and <code>usesEmpty<true></code> are set. Blocking read call to be used from the testbench to read back the data from the component with out-of-band <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.

Table 33. Intel HLS Compiler Memory-Mapped Interfaces

Feature	Valid Values	Default Value	Description
<code>ihc::mm_master<datatype, template arguments></code>	Any valid C++ datatype	Default interface for pointer arguments	<p>The underlying pointer type. Pointer arithmetic performed on the master object conforms to this type. Dereferences of the master results in a load-store site with a width of <code>sizeof(datatype)</code>. The default alignment is aligned to the size of the datatype.</p> <p>Avalon Memory-Mapped (MM) Master interface argument: Multiple template arguments are supported. The template arguments are listed below. Any combination can be used as long as it describes a valid hardware configuration.</p> <p>Example:</p> <pre>component int dut(ihc::mm_master<int, ihc::aspace<2>, ihc::latency<3>, ihc::awidth<10>, ihc::dwidth<32> > &a)</pre> <p>To learn more, review the following tutorials:</p> <ul style="list-style-type: none"> <code><quartus_installdir>/hls/examples/tutorials/interfaces/pointer_mm_master</code> <code><quartus_installdir>/hls/examples/tutorials/interfaces/mm_master_testbench_operators</code>
<code>ihc::dwidth<value></code>	8, 16, 32, 64, 128, 256, 512, or 1024	64	The width of the memory-mapped data bus in bits
<code>ihc::awidth<value></code>	Integer value in the range 1 - 64	64	The width of the memory-mapped address bus in bits. This value affects only the width of the Avalon MM Master interface. The size of the conduit of the base address pointer is always set to 64-bits.
<code>ihc::aspace<value></code>	Integer value greater than 0	1	The address space of the interface that associates with the master. Each unique <code>value</code> results in a separate Avalon MM Master interface on your component. All masters with the same address space are arbitrated within the component to a single interface. As such, these masters must share the same template parameters that describe the interface.
<code>ihc::latency<value></code>	Non-negative integer value	1	The guaranteed latency from when a read command exits the component when the external memory returns valid read data. If this latency is variable (such as when accessing DRAM), set it to 0.
<code>ihc::maxburst<value></code>	Integer value in the range 1 - 1024	1	The maximum number of data transfers that can associate with a read or write transaction. This value controls the width of the <code>burstcount</code> signal. For fixed latency interfaces, this value must be set to 1. For more details, review information about burst signals and the <code>burstcount</code> signal role in " Avalon Memory-Mapped Interface Signal Roles " in <i>Avalon Interface Specifications</i> .
<code>ihc::align<value></code>	Integer value greater than the alignment of the datatype	Alignment of the datatype	<p>The alignment of the base pointer address in bytes. The Intel HLS Compiler uses this information to determine how many simultaneous loads and stores this pointer can permit.</p> <p>For example, if you have a bus with 4 32-bit integers on it, you should use <code>ihc::dwidth<128></code> (bits) and <code>ihc::align<16></code> (bytes). This means that up to 16 contiguous bytes (or 4 32-bit integers) can be loaded or stored as a coalesced memory word per clock cycle.</p>

continued...



Feature	Valid Values	Default Value	Description
			<i>Important:</i> The caller is responsible for aligning the data to the set value for the align argument; otherwise, functional failures might occur.
ihc::readwrite_mode<value>	readwrite, readonly, or writeonly	readwrite	Port direction of the interface. Only the relevant Avalon master signals are generated
ihc::waitrequest<value>	true or false	false	Adds the waitrequest signal that is asserted by the slave when it is unable to respond to a read or write request. For more information about the waitrequest signal, see "Avalon Memory-Mapped Interface Signal Roles" in <i>Avalon Interface Specifications</i> .
getInterfaceAtIndex(int index)			This testbench function is used to index into an mm_master object. It can be useful when iterating over an array and invoking a component on different indices of the array. This function is supported only in the testbench. Example: <pre>int main() { // for(int idx = 0; idx < N; idx++) { dut(src_mm.getInterfaceAtIndex(idx)); } // }</pre>

Table 34. AC Datatypes Supported by the HLS Compiler

AC Datatype	Intel Header File	Description
ac_int	HLS/ac_int.h	Arbitrary width integer support To learn more, review the following tutorials: <ul style="list-style-type: none"> <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_basic_ops <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow <quartus_installdir>/hls/examples/tutorials/best_practices/struct_interfaces
ac_fixed	HLS/ac_fixed.h	Arbitrary precision fixed-point support To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_constructor
	HLS/ac_fixed_math.h	Support for some nonstandard math functions for arbitrary precision fixed-point datatypes To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_fixed_math_library

Table 35. Intel HLS Compiler ac_int Debugging Tools

Feature	Description
<p><i>Macro:</i></p> <pre>#define DEBUG_AC_INT_WARNING If you use this macro, declare it in your code before you declare #include HLS/ac_int.h.</pre>	<p>Enables runtime tracking of ac_int datatypes during x86 emulation (the -march=x86-64 option, which the default option, of the i++ command).</p> <p>This tool uses additional resources for tracking the overflow and empty constructors, and emits a warning for each detected overflow.</p> <p>To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow</p> <p style="text-align: right;">continued...</p>



Feature	Description
<i>Compiler command line option:</i> -D DEBUG_AC_INT_WARNING	
<i>Macro:</i> #define DEBUG_AC_INT_ERROR If you use this macro, declare it in your code before you declare #include HLS/ac_int.h. <i>Compiler command line option:</i> -D DEBUG_AC_INT_ERROR	Enables runtime tracking of ac_int datatypes during x86 emulation of your component (the -march=x86-64 option, which the default option, of the i++ command). This tool uses additional resources to track the overflow and empty constructors, and emits a message for the first overflow that is detected and then exits the component with an error. To learn more, review the tutorial: <quartus_installdir>/hls/examples/tutorials/ac_datatypes/ac_int_overflow

B. Supported Math Functions

The Intel HLS Compiler has built-in support for generating efficient IP out of standard math functions present in the `math.h` C header file. The compiler also has support for some math functions that are not supported by the `math.h` header file, and these functions are provided in `extendedmath.h` C header file.

To use the Intel implementation of `math.h` for Intel FPGAs, include `HLS/math.h` in your function by adding the following line:

```
#include "HLS/math.h"
```

To use the nonstandard math functions that are optimized for Intel FPGAs, include `HLS/extendedmath.h` in your function by adding the following line:

```
#include "HLS/extendedmath.h"
```

The `extendedmath.h` header is compatible only with Intel HLS Compiler. It is not compatible with GCC or Microsoft Visual Studio.

If your component uses arbitrary precision fixed-point datatypes provided in the `ac_fixed.h` header, you use some of the datatypes with some math functions by including the following line:

```
#include "HLS/ac_fixed_math.h"
```

To see examples of how to use the math functions provided by these header files, review the following tutorial: `<quartus_installdir>/hls/examples/tutorials/best_practices/single_vs_double_precision_math`.

B.1. Math Functions Provided by the `math.h` Header File

The Intel HLS Compiler supports a subset of functions that are present in your native compiler `HLS/math.h` header file.

For each `math.h` function listed below, "•" indicates that the HLS compiler supports the function; "X" indicates that the function is not supported.

The math functions supported on Linux operating systems might differ from the math functions supported on Windows operating systems. Review the comments in the `HLS/math.h` header file to see which math functions are supported on the different operating systems.

Table 36. Trigonometric Functions

	Trigonometric Function	Supported by the HLS Compiler?
Double-precision floating point functions	cos	•
	sin	•
	tan	•
	acos	•
	asin	•
	atan	•
	atan2	•
Single-precision floating point functions	cosf	•
	sinf	•
	tanf	•
	acosf	•
	asinf	•
	atanf	•
	atan2f	•

Table 37. Hyperbolic Functions

Hyperbolic Function	Supported by the HLS Compiler?
cosh	•
sinh	•
tanh	•
acosh	X
asinh	X
atanh	X

Table 38. Exponential and Logarithmic Functions

Exponential or Logarithmic Function	Supported by the HLS Compiler?
exp	•
frexp	•
ldexp	•
log	•
log10	•
modf	•
exp2	•
expm1	X
ilogb	•
log1p	X

continued...



Exponential or Logarithmic Function	Supported by the HLS Compiler?
log2	•
logb	•
scalbn	X
scalbln	X

Table 39. Power Functions

Power Function	Supported by the HLS Compiler?
pow	•
sqrt	•
cbrt	X
hypot	X

Table 40. Error and Gamma Functions

Error or Gamma Function	Supported by the HLS Compiler?
erf	X
erfc	X
tgamma	X
lgamma	X

Table 41. Rounding and Remainder Functions

Rounding or Remainder Function	Supported by the HLS Compiler?
ceil	•
floor	•
fmod	•
trunc	•
round	•
lround	X
llround	X
rint	•
lrint	X
llrint	X
nearbyint	X
remainder	X
remquo	X

Table 42. Floating-Point Manipulation Functions

Floating-Point Manipulation Function	Supported by the HLS Compiler?
copysign	X
nan	X
nextafter	X
nexttoward	X

Table 43. Minimum, Maximum, and Difference Functions

Minimum, Maximum, or Difference Function	Supported by the HLS Compiler?
fdim	•
fmax	•
fmin	•

Table 44. Other Functions

Function	Supported by the HLS Compiler?
fabs	•
abs	X
fma	X

Table 45. Classification Macros

Classification Macro	Supported by the HLS Compiler?
fpclassify	X
isfinite	•
isinf	•
isnan	•
isnormal	X
signbit	X

Table 46. Comparison Macros

Comparison Macro	Supported by the HLS Compiler?
isgreater	X
isgreaterequal	X
isless	X
islessequal	X
islessgreater	X
isunordered	X



B.2. Math Functions Provided by the `extendedmath.h` Header File

Adding the `HLS/extendedmath.h` header file adds support for the following functions:

Table 47. Extended math functions

Data type	Function
Double-precision floating point	<ul style="list-style-type: none"> • <code>sincos</code> • <code>acospi</code> • <code>asinpi</code> • <code>atanpi</code> • <code>cospi</code> • <code>sinpi</code> • <code>tanpi</code> • <code>pown</code> • <code>powr</code> • <code>rsqrt</code>
Single-precision floating point	<ul style="list-style-type: none"> • <code>sincosf</code> • <code>acospif</code> • <code>asinpif</code> • <code>atanpif</code> • <code>cospif</code> • <code>sinpif</code> • <code>tanpif</code> • <code>pownf</code> • <code>powrf</code> • <code>rsqrtf</code>

In addition, the `HLS/extendedmath.h` header file supports the following versions of the `popcount` function:

Table 48. Popcount function

Data type	Function
Unsigned char	<code>popcountc</code>
Unsigned short	<code>popcounts</code>
Unsigned int	<code>popcount</code>
Unsigned long	<code>popcountl</code>
Unsigned long long	<code>popcountll</code>

To see an example of how to use the math functions provided by the `extendedmath.h` header file and how to override a math function in the header file so that you can compile your design with GCC or Microsoft Visual Studio, review the following example design: `<quartus_installdir>/hls/examples/QRD`.

B.3. Math Functions Provided by the `ac_fixed_math.h` Header File

Adding the `ac_fixed_math.h` header file adds support for the following arbitrary precision fixed-point (`ac_fixed`) datatype functions:



- `sqrt_fixed`
- `reciprocal_fixed`
- `reciprocal_sqrt_fixed`
- `sin_fixed`
- `cos_fixed`
- `sincos_fixed`
- `sinpi_fixed`
- `cospi_fixed`
- `sincospi_fixed`
- `log_fixed`
- `exp_fixed`

For details about inputs type restrictions, input value limits, and output type propagation rules, review the comments in the `ac_fixed_math.h` header file.