



Intel[®] FPGA SDK for OpenCL[™] Pro Edition

Custom Platform Toolkit User Guide

Updated for Intel[®] Quartus[®] Prime Design Suite: **20.3**



Online Version



Send Feedback

UG-OCL007

ID: **683085**

Version: **2022.03.28**

Contents

1. Intel® FPGA SDK for OpenCL™ Pro Edition Custom Platform Toolkit User Guide.....	4
1.1. Prerequisites for the Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit.....	5
1.2. Overview of the Intel FPGA SDK for OpenCL Pro Edition Custom Platform.....	6
1.2.1. Directories and Files in an Intel FPGA SDK for OpenCL Pro Edition Custom Platform.....	6
1.2.2. Recommendations for Structuring the Custom Platform Directory.....	7
1.3. Custom Platform Automigration for Forward Compatibility.....	7
1.3.1. Customizing Automigration.....	8
1.4. Creating an Intel FPGA SDK for OpenCL Custom Platform.....	8
1.4.1. Designing the Board Hardware.....	9
1.4.2. Creating the Board XML Files.....	15
1.4.3. Creating the MMD Library.....	21
1.4.4. Setting Up the FPGA Client Driver.....	23
1.4.5. Providing Intel FPGA SDK for OpenCL Utilities Support.....	24
1.4.6. Testing the Hardware Design.....	26
1.5. Applying for the Intel FPGA SDK for OpenCL Pro Edition Preferred Board Status.....	27
1.6. Shipping Recommendations.....	28
2. Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit Reference Material....	29
2.1. The Board Platform Designer Subsystem.....	29
2.1.1. Intel FPGA SDK for OpenCL Pro Edition-Specific Platform Designer System Components.....	30
2.2. XML Elements, Attributes, and Parameters in the board_spec.xml File.....	33
2.2.1. board.....	34
2.2.2. device.....	34
2.2.3. global_mem.....	35
2.2.4. host.....	37
2.2.5. channels.....	37
2.2.6. interfaces.....	37
2.2.7. interface.....	38
2.2.8. compile.....	39
2.3. MMD API Descriptions.....	40
2.3.1. aocl_mmd_get_offline_info.....	41
2.3.2. aocl_mmd_get_info.....	43
2.3.3. aocl_mmd_open.....	45
2.3.4. aocl_mmd_close.....	45
2.3.5. aocl_mmd_read.....	45
2.3.6. aocl_mmd_write.....	46
2.3.7. aocl_mmd_copy.....	47
2.3.8. aocl_mmd_set_interrupt_handler.....	48
2.3.9. aocl_mmd_set_device_interrupt_handler.....	49
2.3.10. aocl_mmd_set_status_handler.....	50
2.3.11. aocl_mmd_yield.....	51
2.3.12. aocl_mmd_shared_mem_alloc.....	51
2.3.13. aocl_mmd_shared_mem_free.....	52
2.3.14. aocl_mmd_program.....	52
2.3.15. aocl_mmd_reprogram.....	53
2.3.16. aocl_mmd_hostchannel_create.....	53

2.3.17. aocl_mmd_hostchannel_destroy.....	54
2.3.18. aocl_mmd_hostchannel_get_buffer.....	54
2.3.19. aocl_mmd_hostchannel_ack_buffer.....	55
2.3.20. aocl_mmd_host_alloc.....	55
2.3.21. aocl_mmd_free.....	56
2.3.22. aocl_mmd_device_alloc.....	57
2.3.23. aocl_mmd_shared_alloc.....	57
2.3.24. aocl_mmd_shared_migrate.....	58
3. Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit Archives.....	60
4. Document Revision History for Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit User Guide	61

1. Intel® FPGA SDK for OpenCL™ Pro Edition Custom Platform Toolkit User Guide

The *Intel® FPGA SDK for OpenCL™ Pro Edition Custom Platform Toolkit User Guide* outlines the procedure for creating an Intel FPGA Software Development Kit (SDK) for OpenCL Pro Edition Custom Platform.

Note: If you want to work with Intel oneAPI Base Toolkit and Intel FPGA Add-on for oneAPI Base Toolkit, refer to the [Installation Guide for Intel oneAPI Toolkits](#) and [Intel oneAPI Programming Guide](#) for more information.

The Intel FPGA SDK for OpenCL⁽¹⁾⁽²⁾ Pro Edition Custom Platform Toolkit provides the necessary tools for implementing a fully functional Custom Platform. The Custom Platform Toolkit is available in the `INTELFPGAOCCLSDKROOT/board` directory, where the environment variable `INTELFPGAOCCLSDKROOT` points to the location of the SDK installation.

The goal is to enable an SDK user to target any given Custom Platform seamlessly by performing the following tasks:

1. Acquire an accelerator board and plug it into their system.
2. Acquire the Custom Platform and unpack it to a local directory.
3. Set the environment variable `QUARTUS_ROOTDIR_OVERRIDE` to point to installation directory of the Intel Quartus® Prime Pro Edition software.
4. Invoke the `aocl install <path_to_customplatform>` utility command.
5. Compile the OpenCL kernel and build the host application.
6. Set environment variables to point to the location of the memory-mapped device (MMD) library.
 - For Windows* systems, set the `PATH` environment variable.
 - For Linux* systems, set the `LD_LIBRARY_PATH` environment variable.
7. Run the host application.

(1) OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission of the Khronos Group™.

(2) The Intel FPGA SDK for OpenCL is based on a published Khronos Specification, and has passed the Khronos Conformance Testing Process. Current conformance status is available at www.khronos.org/conformance.

1.1. Prerequisites for the Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit

The *Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit User Guide* assumes that you have prior hardware design knowledge necessary for using the Custom Platform Toolkit to create an Intel FPGA SDK for OpenCL Pro Edition Custom Platform.

You must have experiences in the following hardware design areas:

- Intel Quartus Prime Pro Edition software design with Platform Designer, HDL and Tcl
- Intel FPGA intellectual property (IP) necessary to communicate with the physical interfaces of the board
- High speed design, timing analysis and Synopsys Design Constraints (SDC) constraints
- FPGA architecture, including clock and global routing, floorplanning, and I/O
- Team-based design (that is, incremental compilation)

You must install the Intel Quartus Prime Pro Edition software, the relevant device support file(s), and the SDK on your machine. Refer to the *Intel FPGA SDK for OpenCL Pro Edition Getting Started Guide* for installation instructions.

You have the following Custom Platform design options:

- Refer to the information in this document to create a Custom Platform from the templates available in the Custom Platform Toolkit.
- Refer to the information in this document and the following documents to create a Custom Platform by modifying relevant files in the Intel Arria® 10 GX FPGA Development Kit Reference Platform or the Intel Arria 10 SoC Development Kit Reference Platform, available with the SDK:
 - *Intel FPGA SDK for OpenCL Intel Arria 10 GX FPGA Development Kit Reference Platform Porting Guide*
 - *AN 807: Configuring the Intel Arria 10 GX FPGA Development Kit for the Intel FPGA SDK for OpenCL*
 - *Intel FPGA SDK for OpenCL Intel Arria 10 SoC Development Kit Reference Platform Porting Guide*

Related Information

- [Intel FPGA SDK for OpenCL Pro Edition Getting Started Guide](#)
- [Intel FPGA SDK for OpenCL Intel Arria 10 GX FPGA Development Kit Reference Platform Porting Guide](#)
- [AN 807: Configuring the Intel Arria 10 GX FPGA Development Kit for the Intel FPGA SDK for OpenCL](#)
- [Intel FPGA SDK for OpenCL Intel Arria 10 SoC Development Kit Reference Platform Porting Guide](#)
- [Intel FPGA Platforms](#)

1.2. Overview of the Intel FPGA SDK for OpenCL Pro Edition Custom Platform

An Intel FPGA SDK for OpenCL Pro Edition Custom Platform is a collection of tools and libraries necessary for the communication between the Intel FPGA SDK for OpenCL Offline Compiler and the FPGA boards.

Currently, the offline compiler targets a single Custom Platform at a time.

A given Custom Platform installation can include several board variants of the same board interface. You might have different FPGA parts, or you might want to support different subsets of board interfaces. Co-locating the board variants allows simultaneous communication with different boards in a multiple-device environment.

An Intel FPGA SDK for OpenCL Pro Edition Custom Platform contains the following components:

- **Intel Quartus Prime skeleton project**—A Intel Quartus Prime project for your board, which the SDK's offline compiler modifies to include the compiled kernel. This project must include a post-place-and-route partition for all logic not controlled by the kernel clock.
- **Board installation setup**—A description of your board and its various components.
- **Generic I/O interface**—An MMD software library that implements basic I/O between the host and the board.
- **Board utilities**—An implementation of the SDK's utilities for managing the accelerator board, including tasks such as installing and testing the board.

1.2.1. Directories and Files in an Intel FPGA SDK for OpenCL Pro Edition Custom Platform

Populate your Intel FPGA SDK for OpenCL Pro Edition Custom Platform with files, libraries and drivers that allow an OpenCL kernel to run on the target FPGA board.

Table 1. Contents within the Top-Level Custom Platform Directory

Content	Description
board_env.xml	The XML file that describes the board installation to the SDK.
<hardware>	Directory containing the Intel Quartus Prime projects for the supported boards within a given Custom Platform. Specify the name of this directory in the board_env.xml file. Within this directory, the SDK assumes that any subdirectory containing a board_spec.xml file is a board.
include	Directory containing board-specific header files.
source	Directory containing board-specific files, libraries and drivers.
linux64 (on Linux) and windows64 (on Windows)	Directory containing platform-specific (for example, x86_64 Linux) drivers and utilities.

1.2.2. Recommendations for Structuring the Custom Platform Directory

For ease of use, consider adopting the Intel-recommended directory structure and naming convention when you create an Intel FPGA SDK for OpenCL Pro Edition Custom Platform.

- Make the `INTELFPGAOCCLSDKROOT/board` directory the location of the board installation, where `INTELFPGAOCCLSDKROOT` points to the location of the SDK installation.

Attention: Do not remove any existing subdirectories from the `INTELFPGAOCCLSDKROOT/board` directory.

- Create a `<board_vendor_name>` subdirectory within the `INTELFPGAOCCLSDKROOT/board` directory to store the Custom Platform.
- Store the contents of a given Custom Platform in a `INTELFPGAOCCLSDKROOT/board/<board_vendor_name>/<board_family_name>` subdirectory.
- Assign unique names to software libraries (for example, `lib<board_vendor_name>_<board_family_name>.so`) to avoid name collisions.

For example, if you (ABC Incorporated) create a Custom Platform for a family of boards named XYZ, set up your Custom Platform such that the SDK user can access XYZ by installing the XYZ Custom Platform in `INTELFPGAOCCLSDKROOT/board/ABC/XYZ`, where `INTELFPGAOCCLSDKROOT` is the environment variable that points to the absolute path to the SDK installation package.

1.3. Custom Platform Automigration for Forward Compatibility

The automigration feature updates an existing Intel-registered Custom Platform for use with the current version of the Intel Quartus Prime Design Suite and the Intel FPGA SDK for OpenCL Pro Edition.

Important: Automigration is more likely to complete successfully if your Custom Platform resembles an Intel FPGA Reference Platform as closely as possible.

The following information applies to a Custom Platform that is version 14.0 and beyond:

1. To update a Custom Platform for use with the current version of the Intel Quartus Prime Design Suite, which includes the SDK, do not modify your Custom Platform. The automigration capability detects the version of your Custom Platform based on certain characteristics and updates it automatically.
2. If you have modified a Custom Platform and you want to update it for use with the current version of the Intel Quartus Prime Design Suite, which includes the SDK, implement all features mandatory for the current version of the Custom Platform. After you modify a Custom Platform, automigration can no longer correctly detect its characteristics. Therefore, you must upgrade your Custom Platform manually.

A successfully-migrated Custom Platform preserves its original functionality. In most cases, new features in a new Intel Quartus Prime Design Suite or SDK version does not interfere with Custom Platform functionality.

When the Intel FPGA SDK for OpenCL Offline Compiler compiles a kernel, it probes the `board_spec.xml` file for the following information:

1. The version of the Custom Platform, as specified by the `version` attribute of the `board` XML element.
2. The platform type, as specified by the `platform_type` parameter of the `auto_migrate` attribute within the `compile` XML element.

Based on the information, the SDK names a set of fixes it must apply during Custom Platform migration. It applies the fixes to the Intel Quartus Prime project that the offline compiler uses to compile the OpenCL kernel. It also generates an `automigration.rpt` report file in the SDK user's current working directory describing the applied fixes.

The automigration process does not modify the installed Custom Platform.

Note: If automigration fails, contact your local Intel FPGA product's field applications engineer for assistance.

1.3.1. Customizing Automigration

You and the Intel FPGA SDK for OpenCL user both have the ability to disable the automigration of an installed Custom Platform. In addition, you may choose which named fixes, identified by the SDK, you want to apply to your Custom Platform.

1. Disable automigration in one of the following manners:
 - If you are a board developer, within the `compile` XML element in the `board_spec.xml` file, set the `platform_type` parameter of the `auto_migrate` attribute to `none`.
 - If you are an SDK user, invoke the `aoc -no-auto-migrate` command.
2. To explicitly include or exclude fixes that the SDK identifies, in the `board_spec.xml` file, subscribe or unsubscribe to each fix by listing it in the `include_fixes` or `exclude_fixes` parameter, respectively. The `include_fixes` and `exclude_fixes` parameters are part of the `auto_migrate` attribute within the `compile` element. When listing multiple fixes, separate each fix by a comma.

Refer to the `automigration.rpt` file for the names of the fixes that you specify in the `include_fixes` and `exclude_fixes` parameters.

1.4. Creating an Intel FPGA SDK for OpenCL Custom Platform

The following topics outline the tasks you must perform to create a Custom Platform for use with the Intel FPGA SDK for OpenCL.

1. [Designing the Board Hardware](#) on page 9
2. [Creating the Board XML Files](#) on page 15
3. [Creating the MMD Library](#) on page 21
4. [Setting Up the FPGA Client Driver](#) on page 23
5. [Providing Intel FPGA SDK for OpenCL Utilities Support](#) on page 24
6. [Testing the Hardware Design](#) on page 26

1.4.1. Designing the Board Hardware

To design an accelerator board for use with the Intel FPGA SDK for OpenCL Pro Edition, you must create all the board and system components, and the files that describe your hardware design to the Intel FPGA SDK for OpenCL Offline Compiler.

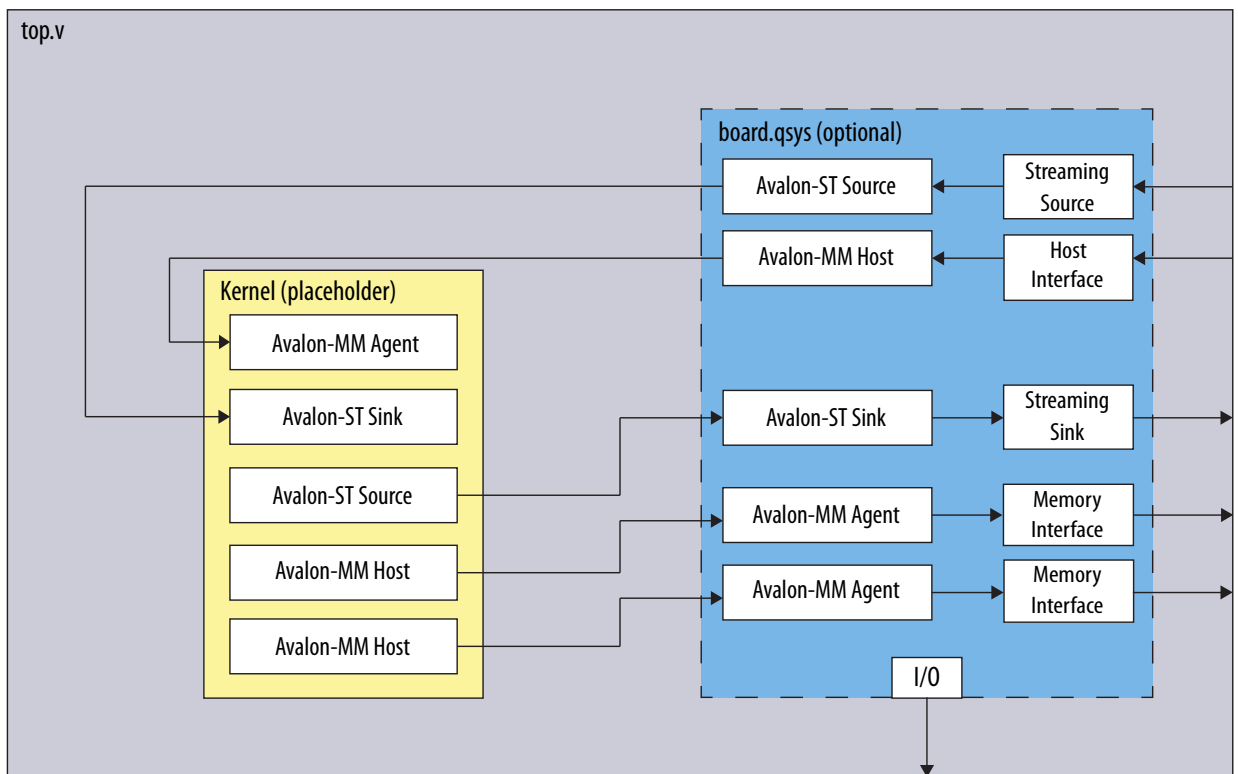
Each board variant in the Custom Platform consists of an Intel Quartus Prime project, and a `board_spec.xml` XML file that describes the system to the offline compiler. The `board_spec.xml` file describes the interfaces necessary to connect to the kernel. The offline compiler generates a custom circuit based on the data from the `board_spec.xml` file. Then it incorporates the OpenCL kernel into the Platform Designer system you create for all nonkernel logic.

You must preserve the design of all nonkernel logic. You can preserve your design in the Intel Quartus Prime Pro Edition software via one of the following methods:

- Create a design partition containing all nonkernel logic under a single HDL hierarchy and then export the partition. For example, you may create and export a `board.qsys` Platform Designer subsystem (see figure below). The top-level `system.qsys` Platform Designer system can then instantiate this exported board Platform Designer subsystem.
- Implement the Configuration via Protocol (CvP) configuration scheme, which preserves all logic outside a design partition. In this case, you only need to create a partition around the kernel logic. You may place all nonkernel logic into a single top-level Platform Designer system file (for example, `system.qsys`).

You must design all the components and the `board_spec.xml` file that describe the system to the SDK.

Figure 1. Example System Hierarchy with a Board Platform Designer Subsystem



For more information, refer to the [System Design of Intel Stratix® 10 Reference Platform](#) in the *Intel FPGA SDK for OpenCL: Intel Stratix 10 GX FPGA Development Kit Reference Platform Porting Guide*

1. [Creating the Board Platform Designer System](#) on page 11
2. [Establishing Guaranteed Timing Flow](#) on page 14

1.4.1.1. Creating the Board Platform Designer System

When designing your board hardware, you have the option to create a Platform Designer subsystem within `system.qsys` that contains all the board logic. In addition to organizing your design code, having this subsystem allows you to create a Intel Quartus Prime partition that you can preserve. To create your board system in a Platform Designer subsystem, you may modify the `board.qsys` template in the Custom Platform Toolkit.

An implementation of a board Platform Designer subsystem might include the following components:

- Proper reset sequencing
- Intel FPGA SDK for OpenCL Pro Edition-specific components
- Host-to-FPGA communication IP
- Memory IP used for SDK's global memory
- Streaming channels to board-specific interfaces

Refer to *The Board Platform Designer System* section for more information.

You can refer to the templates of the following hardware design files in one of the reference platform BSPs available at the [FPGA Software Download Center](#):

- `board.qsys`
- `top.v`
- `top.qpf`
- `board_spec.xml`

Refer to the template of the `post_flow.tcl` file in one of the BSPs available at [FPGA Software Download Center](#). If you need help with downloading the BSP, refer to the [Downloading the Intel FPGA SDK for OpenCL Pro Edition](#) topic in *Intel FPGA SDK for OpenCL Pro Edition: Getting Started Guide*.

Note: Starting from the 20.3 release, Intel has removed the support for Windows and Linux OpenCL BSPs. Use version 20.2 or older OpenCL BSPs available at [FPGA Software Download Center](#) as a reference. If you want to migrate your OpenCL BSP to a newer version, follow the recommended steps provided in the Reference Platform Porting Guides available under [Intel FPGA SDK for OpenCL documentation](#).

To create non-kernel logic, perform the following tasks in the `board.qsys` top-level Platform Designer system or in a board Platform Designer subsystem:

1. In Platform Designer, add your host and memory IPs to the Platform Designer system, and establish all necessary connections and exports.

Attention: You might need to acquire separate IP licenses. For a list of available licensed and unlicensed IP solutions, visit the All Intellectual Property page of the Intel FPGA website. For more information about each IP, click the link in the Product Name column to navigate to the product page.

- a. Connect your host interface clock such that it drives `por_reset_controller/clk`. Your design's global reset and clock inputs are fed to a reset counter (`por_reset_counter`). This reset counter then synchronizes to the host interface clock in the Merlin Reset Controller (`por_reset_controller`).

The `por_reset_counter` ACL SW Reset component implements the power-on reset. It resets all the device hardware by issuing a reset for a number of cycles after the FPGA completes its configuration.

- b. Modify the parameters of the `pipe_stage_host_ctrl` Avalon® Memory-Mapped (Avalon-MM) Pipeline Bridge component such that it can receive requests from your host IP. Connect your host interface's Avalon-MM host port to the `s0` port of `pipe_stage_host_ctrl`. Connect the `m0` port of `pipe_stage_host_ctrl` to all the peripherals that must communicate with your host interface, including the OpenCL Kernel Clock Generator and the OpenCL Kernel Interface components.
- c. Adjust the number of `clock_cross_kernel_mem_<N>` Avalon-MM Clock Crossing Bridge components to match the number of memory interfaces on your board. This component performs clock crossing between the kernel and memory interfaces. Modify the parameters of each component so that they are consistent with the parameters of the OpenCL Memory Bank Divider component and the `interface` attribute described in `board_spec.xml`. Connect the `m0` host, clock, and reset ports of `clock_cross_kernel_mem_<N>` (that is, `m0`, `m0_clk`, and `m0_reset`, respectively) to your memory IP.

Important: Connect `m0_reset` in such a way that assertion of `kernel_reset` from the OpenCL Memory Bank Divider component triggers this reset.

2. Customize the SDK-specific Platform Designer (Standard) system components.
 - a. In Platform Designer, click **Tools** ► **Options**. In the **Options** dialog box, add `INTELFPGAOCSDKROOT/ip/board` to the **Qsys IP Search Path** and then click **Finish**.
 - b. Instantiate the OpenCL Kernel Clock Generator component. Specify the component parameters, and connect the signals and ports as outlined in the *OpenCL Kernel Clock Generator* section.
 - c. Instantiate the OpenCL Kernel Interface component. Specify the component parameters, and connect the signals and ports as outlined in the *OpenCL Kernel Interface* section.
 - d. For each global memory type, instantiate the OpenCL Memory Bank Divider component. Specify the component parameters, and connect the signals and ports as outlined in the *OpenCL Memory Bank Divider* section.

Attention: Set the parameters such that the resulting bank hosts have the equivalent address bits and burst widths as those from the kernel, as defined in the `interface` attribute of the `global_mem` element in the `board_spec.xml` file. For each memory bank, Platform Designer generates a host that inherits the same characteristics as your specifications.

3. If you choose to create a Platform Designer subsystem for the nonkernel logic, export any necessary I/Os to the top-level `system.qsys` Platform Designer system.
4. Edit the top-level `top.v` file to instantiate `system.qsys` and connect any board-specific I/Os.
5. Set up the `top.qpf` Intel Quartus Prime project with all the necessary settings for your board design.
6. Modify the `post_flow.tcl` file to include the Tcl code that generates the `fpga.bin` file during Intel Quartus Prime compilation.
The `fpga.bin` file is necessary for programming the board.
7. Edit the `board_spec.xml` file to include board-specific descriptions.

Related Information

- [All Intellectual Property page on the Intel FPGA website](#)
- [OpenCL Kernel Clock Generator](#) on page 30
- [OpenCL Kernel Interface](#) on page 31
- [OpenCL Memory Bank Divider](#) on page 32
- [The Board Platform Designer Subsystem](#) on page 29

1.4.1.1.1. General Quality of Results Considerations for the Exported Board Partition

When generating a post-place-and-route partition, take into account several design considerations for the exported board partition that might have unexpected consequences on the Intel FPGA SDK for OpenCL Pro Edition compilation results. The best approach to optimizing the board partition is to experiment with a range of different OpenCL kernels.

The list below captures some of the parameters that might impact the quality of SDK compilation results:

- **Resources Used**
Minimize the number of resources the partition uses to maximize the resources available for the OpenCL kernels.
- **Kernel Clock Frequency**
Intel recommends that the kernel clock has a high clock constraint (for example, greater than 350 MHz for a Stratix V device). The amount of logic in the partition clocked by the kernel clock should be relatively small. This logic should not limit the kernel clock speed for even the simplest OpenCL kernels. Therefore, at least within the partition, the kernel clock should have a high clock constraint.
- **Host-to-Memory Bandwidth**
The host-to-memory bandwidth is the transfer speed between the host processor to the physical memories on the accelerator card. To measure this memory bandwidth, compile and run the host application included with the Custom Platform Toolkit.
- **Kernel-to-Memory Bandwidth**
The kernel-to-memory bandwidth is the maximum transfer speed possible between the OpenCL kernels and global memory.
To measure this memory bandwidth, compile and run the host program included in the `/tests/boardtest/host` directory of the Custom Platform Toolkit.
- **Fitter Quality of Results (QoR)**
To ensure that OpenCL designs consuming much of the device's resources can still achieve high clock frequencies, region-constrain the partition to the edges of the FPGA. The constraint allows OpenCL kernel logic to occupy the center of the device, which has the most connectivity with all other nodes.
Test compile large designs to ensure that other Fitter-induced artifacts in the partition do not interfere with the QoR of the kernel compilations.
- **Routability**
The routing resources that the partition consumes can affect the routability of a compiled OpenCL design. A kernel might use every digital signal processing (DSP) block or memory block on the FPGA; however, routing resources that the partition uses might render one of these blocks unroutable. This routing issue causes compilation of the Intel Quartus Prime project to fail at the fitting step. Therefore, it is imperative that you test a partition with designs that use all DSP and memory blocks.

1.4.1.2. Establishing Guaranteed Timing Flow

Deliver a design partition for nonkernel logic that has a clean timing closure flow as part of your Custom Platform.

1. Create a placed and routed design partition using the Intel Quartus Prime Pro Edition software. This is the design partition for nonkernel logic.
2. Import the post-fit partition from 1 on page 15 into the top-level design as part of the compilation flow.
3. Run the `INTELFPGAOCCLSDKROOT/ip/board/bsp/adjust_plls.tcl` script as a post-flow process, where `INTELFPGAOCCLSDKROOT` points to the path of the Intel FPGA SDK for OpenCL Pro Edition installation.

The `adjust_plls.tcl` script determines the maximum kernel clock frequency and stores it in the `pll_rom` on-chip memory of the OpenCL Kernel Clock Generator component.

1.4.2. Creating the Board XML Files

Your Custom Platform must include the XML files that describe your Custom Platform and each of your hardware system to the Intel FPGA SDK for OpenCL Pro Edition. You may create these XML files in simple text editors (for example, WordPad for Windows, and vi for Linux).

[Creating the board_env.xml File](#) on page 15

[Creating the board_spec.xml File](#) on page 17

1.4.2.1. Creating the board_env.xml File

The `board_env.xml` file describes your Custom Platform to the SDK. Together with the other contents of the Custom Platform, the `board_env.xml` file sets up the board installation that enables the offline compiler to target a specific accelerator board.

A `board_env.xml` template is available in the `/board_package` directory of the Custom Platform Toolkit.

1. Create a `board_env` top-level XML element. Within `board_env`, include the following XML elements:
 - `hardware`
 - `platform`
 - Include a `platform` element for each operating system that your Custom Platform supports.
2. Within each `platform` element, include the following XML elements:
 - `mmdlib`
 - `linkflags`
 - `linklibs`
 - `utilbindir`
3. Parameterize each element and corresponding attribute(s) with information specific to your Custom Platform, as outline in the table below:

Table 2. Specifications of XML Elements and Attributes in the board_env.xml File

Element	Attribute Description
board_env	<p>version: The SDK's Custom Platform Toolkit release you use to create your Custom Platform.</p> <p>Attention: The Custom Platform version must match the SDK version you use to develop the Custom Platform.</p> <p>name: Name of the board installation directory containing your Custom Platform.</p>
hardware	<p>dir: Name of the subdirectory, within the board installation directory, that contains the board variants.</p> <p>default: The default board variant that the offline compiler targets when the SDK user does not specify an explicit argument for the --board <board_name> offline compiler option.</p>
platform	<p>name: Name of the operating system.</p> <p>Refer to the <i>Intel FPGA SDK for OpenCL Pro Edition Getting Started Guide</i> and the <i>Intel FPGA RTE for OpenCL Pro Edition Getting Started Guide</i> for more information.</p>
mmdlib	<p>A string that specifies the path to the MMD library of your Custom Platform. To load multiple libraries, specify them in an ordered, comma-separated list. The host application loads the libraries in the order that they appear in the list.</p> <p><i>Tip:</i> You can use %b to reference your board installation directory.</p>
linkflags	<p>A string that specifies the linker flags necessary for linking with the MMD layer available with the board.</p> <p><i>Tip:</i> You can use %a to reference the SDK installation directory and %b to reference your board installation directory.</p>
linklibs	<p>A string that specifies the libraries the SDK must link against to use the MMD layer available with the board.</p>
utilbindir	<p>Directory in which the SDK expects to locate its utility executables (that is, install, uninstall, program, diagnose, and flash).</p> <p><i>Tip:</i> You can use %a to reference the SDK installation directory and %b to reference your board installation directory.</p>

Your board_env.xml file should resemble the following example:

```
<?xml version="1.0"?>
<board_env version="<SDK_version>" name="<Custom_Platform_name>">
  <hardware dir="hardware" default="<board_name>"></hardware>
  <platform name="linux64">
    <mmdlib>%b/linux64/lib/libaltera_<Custom_Platform_name>_mmd.so</mmdlib>
    <linkflags>-L%b/linux64/lib</linkflags>
    <linklibs>-laltera_<Custom_Platform_name>_mmd</linklibs>
    <utilbindir>%b/linux64/libexec</utilbindir>
  </platform>

  <platform name="windows64">
    <mmdlib>%b/windows64/bin/altera_<Custom_Platform_name>_mmd.dll</mmdlib>
    <linkflags>/libpath:%b/windows64/lib</linkflags>
    <linklibs>altera_<Custom_Platform_name>_mmd.lib</linklibs>
    <utilbindir>%b/windows64/libexec</utilbindir>
  </platform>
</board_env>
```

Where:

- <SDK_version> is the Intel FPGA SDK for OpenCL Pro Edition version.
- <Custom_Platform_name> is the name of your Custom Platform (for example, a10_ref).
- <board_name> is the name of the FPGA board.

Related Information

- [Prerequisites for the Intel FPGA SDK for OpenCL Pro Edition](#)
- [Prerequisites for the Intel FPGA RTE for OpenCL Pro Edition](#)

1.4.2.1.1. Testing the board_env.xml File

After you generate the `board_env.xml` file, test the file within your board installation directory to ensure that the Intel FPGA SDK for OpenCL Offline Compiler recognizes the board installation.

1. At the command prompt, invoke the `aocl board-xml-test` command to verify that the Intel FPGA SDK for OpenCL Pro Edition can locate the correct field values.

The SDK generates an output similar to the one below:

```
board-path      = <path_to_customplatform>
board-version   = <SDK_version>
board-name      = <customplatform_name>
board-default   = <board_name>
board-hw-path   = <path_to_customplatform>/hardware/<customplatform_name>
board-link-flags = /libpath:<path_to_customplatform>/windows64/lib
board-libs      = altera_<customplatform_name>_mmd.lib
board-util-bin  = <path_to_customplatform>/windows64/libexec
board-mmdlib    = <path_to_customplatform>/windows64/bin/
altera_<customplatform_name>_mmd.dll
```

2. Invoke the `aoc -list-boards` command to verify that the offline compiler can identify and report the board variants in the Custom Platform. For example, if your Custom Platform includes two FPGA boards, the SDK generates an output similar to the one below:

```
Board list:
  <board_name_1>
  <board_name_2>
```

The last board installation test takes place when you use the offline compiler to generate a design for your board.

Related Information

- [Creating the board_env.xml File](#) on page 15
- [Testing the Hardware Design](#) on page 26

1.4.2.2. Creating the board_spec.xml File

The `board_spec.xml` XML file contains metadata necessary to describe your hardware system to the Intel FPGA SDK for OpenCL Pro Edition.

For detailed descriptions on the type of information you must include in the `board_spec.xml` file, refer to the *XML Elements, Attributes, and Parameters in the board_spec.xml File* section.

1. Structure the `board_spec.xml` file to include the following XML elements and attributes:

Table 3. XML Elements and Attributes Specified in the board_spec.xml File

Element	Attribute
board	version, name
device	device_model, used_resources
global_mem	name, max_bandwidth, interleaved_bytes, config_addr, [default], interface
host	kernel_config
[channels]	interface
interfaces	interface, kernel_clk_reset
compile	project, revision, qsys_file, generic_kernel, generate_cmd, synthesize_cmd, auto_migrate

- For the board element, specify the board version and the name of the accelerator board. The name of the board must match the name of the directory in which the board_spec.xml file resides.

Important: The board version must match the SDK version you use to develop the Custom Platform.

Attention: The board name must contain a combination of only letters, numbers, underscores (_), hyphens (-), or periods (.) (for example: a10_ref).

- For the device element, perform the following steps to specify the name of the device model file.
 - Navigate to the `INTELFPGAOCCLSDKROOT/share/models/dm` directory, where `INTELFPGAOCCLSDKROOT` points to the path to the SDK installation. The directory contains a list of device models files that describe available FPGA resources on accelerator boards.
 - If your device is listed in the `dm` directory, specify the `device_model` attribute with the name of the device model file. Proceed to Step 4.
For example, `device_model="10ax115s2f45i2sges_dm.xml"`
 - If your device is not listed in the `dm` directory, or if your board uses an FPGA that does not have a device model, create a new device model by performing the tasks described in Steps d to g:
 - Copy a device model from the `INTELFPGAOCCLSDKROOT/share/models/dm` directory (for example, `10ax115s2f45e2lg_dm.xml`).
 - Place your copy of the device model in the Custom Platform subdirectory in which your `board_spec.xml` file resides.
 - Rename the file, and modify the values to describe the part your board uses.
Important: The file name must conform to the format `<part_number>_dm.xml`, where `<part_number>` is the official Intel FPGA device part number (for example, `10ax115s2f45i1sg`).
You can search for the Intel FPGA device part number on the [Buy Intel FPGA](#) page.
 - In the `board_spec.xml` file, update the `device_model` attribute of the device element with the name of your file.

4. For the `device` element, specify the parameters in the `used_resources` attribute to describe the FPGA resources that the board design consumes in the absence of any OpenCL kernel.

If your design includes a defined partition around all the board logic, you can extract the data from the Partition Statistics section of the Fitter report.

5. For each global memory type, specify the following information:
 - a. Name of the memory type.
 - b. The combined maximum global memory bandwidth.
You can calculate this bandwidth value from datasheets of your memories.
 - c. The size of the data that the Intel FPGA SDK for OpenCL Offline Compiler interleaves across memory banks.
 - d. If you have a homogeneous memory system, proceed to Step e. If you have a heterogeneous memory system, for each global memory type, specify the `config_addr` attribute with the base address of the ACL Mem Organization Control Platform Designer component (`mem_org_mode`).
 - e. If you choose to set a global memory type as default, assign a value of 1 to the optional `default` attribute. The default memory must start at address 0x0

If you do not include this attribute, the first memory defined in the `board_spec.xml` file becomes the default memory.

- f. Specify the parameters in the `interface` attribute to describe the characteristics of each memory interface.
6. For the `host` element, specify the parameters in the `kernel_config` attribute to describe the offset at which the kernel resides. Determine the start of the offset from the perspective of the `kernel_cra` host in the OpenCL Kernel Interface Platform Designer component.
7. If your board provides channels for direct OpenCL kernel-to-I/O accesses, include the `channels` element for all channel interfaces. Specify the parameters in the `interface` attribute to describe the characteristics of each channel interface.
8. Include the `interfaces` element to describe the kernel interfaces connecting to and controlling OpenCL kernels. Include one of each interface types (that is `host`, `irq`, and `streamsource`).

- a. Specify the parameters in the `interface` attribute to describe the characteristics of each kernel interface.

For the `streamsource` interface type, also specify the `clock` attribute with the name of the clock the snoop stream uses. Usually, this clock is the kernel clock.

Important: Update the width of the snoop interface (`acl_internal_snoop`) specified with the `streamsource` kernel interface. Updating the width ensures that the `global_mem` interface entries in `board_spec.xml` match the characteristics of the bank<N> Avalon Memory-Mapped (Avalon-MM) hosts from corresponding OpenCL Memory Bank Divider component for the default memory.

- b. Specify the parameters in the `kernel_clk_reset` attribute to include the exported kernel clock and reset interfaces as kernel interfaces.
9. Include the `compile` element and specify its attributes to control the Intel Quartus Prime compilation, registration, and automigration.

Below is the XML code of an example `board_spec.xml` file:

```
<?xml version="1.0"?>
<board version="22.3" name="a10gx">

  <compile name="top" project="top" revision="top" qsys_file="none"
generic_kernel="1">
    <generate cmd="quartus_sh -t scripts/pre_flow_pr.tcl"/>
    <synthesize cmd="quartus_cdb -t import_compile.tcl"/>
    <auto_migrate platform_type="a10_ref" >
      <include fixes=""/>
      <exclude fixes=""/>
    </auto_migrate>
  </compile>

  <compile name="base" project="top" revision="base" qsys_file="none"
generic_kernel="1">
    <generate cmd="quartus_sh -t scripts/pre_flow_pr.tcl base"/>
    <synthesize cmd="quartus_sh --flow compile top -c base"/>
    <auto_migrate platform_type="a10_ref" >
      <include fixes=""/>
      <exclude fixes="pre_skipbak,post_skipbak"/>
    </auto_migrate>
  </compile>

  <compile name="flat" project="top" revision="flat" qsys_file="none"
generic_kernel="1">
    <generate cmd="quartus_sh -t scripts/pre_flow_pr.tcl flat"/>
    <synthesize cmd="quartus_sh --flow compile top -c flat"/>
    <auto_migrate platform_type="a10_ref" >
      <include fixes=""/>
      <exclude fixes="pre_skipbak,post_skipbak"/>
    </auto_migrate>
  </compile>

  <compile name="regenerate_cache" project="top" revision="top" qsys_file="none"
generic_kernel="1">
    <generate cmd="quartus_sh -t scripts/pre_flow_pr.tcl"/>
    <synthesize cmd="quartus_sh -t scripts/regenerate_cache.tcl a10_ref a10gx"/>
    <auto_migrate platform_type="a10_ref" >
      <include fixes=""/>
      <exclude fixes="pre_skipbak"/>
    </auto_migrate>
  </compile>

  <device device_model="10ax115s2f45i2sges_dm.xml">
    <used_resources>
      <alms num="33400"/> <!-- Total ALMs - ALMs available to kernel_system_inst
-->
      <ffs num="133600"/>
      <dsps num="0"/>
      <rams num="182"/>
    </used_resources>
  </device>

  <!-- DDR4-2400 -->
  <global_mem name="DDR" max_bandwidth="19200" interleaved_bytes="1024"
config_addr="0x018">
    <interface name="board" port="kernel_mem0" type="agent" width="512"
maxburst="16" address="0x00000000" size="0x80000000" latency="240" addpipe="1"/>
  </global_mem>

</host>
```

```
<kernel_config start="0x00000000" size="0x0100000" />
</host>

<interfaces>
  <interface name="board" port="kernel_cra" type="host" width="64" misc="0"/>
  <interface name="board" port="kernel_irq" type="irq" width="1"/>
  <interface name="board" port="acl_internal_snoop" type="streamsource"
enable="SNOOPENABLE" width="31" clock="board.kernel_clk"/>
  <kernel_clk_reset clk="board.kernel_clk" clk2x="board.kernel_clk2x"
reset="board.kernel_reset"/>
</interfaces>

</board>
```

Related Information

XML Elements, Attributes, and Parameters in the board_spec.xml File on page 33

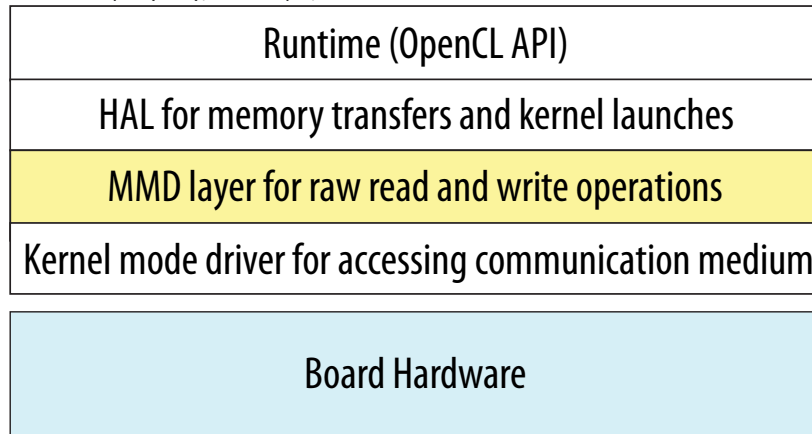
1.4.3. Creating the MMD Library

Your Custom Platform requires an MMD layer necessary for communication with the accelerator board.

You must implement a file I/O-like software interface such as open, read, write, and close to communicate with the accelerator board over any medium. The result of your implementation is a set of linker arguments that allows an OpenCL host application to link against the MMD layer of the target board. A dynamic link library (DLL) that fully implements the MMD layer is also necessary for the communication.

Figure 2. Intel FPGA SDK for OpenCL Software Architecture

This figure depicts the four layers of the Intel FPGA SDK for OpenCL Pro Edition software architecture: runtime, hardware abstraction layer (HAL), MMD layer, and kernel mode driver.



The following tasks outline the procedure for creating an MMD library for use with PCI Express® (PCIe®).

1. Name a new library file that implements the MMD layer in the following manner:

```
<board_vendor_name>_<board_family_name>[_<unique_string>]_mmd.
<a|so|lib|dll>
```

Where:

- <board_vendor_name> is the entity responsible for the accelerator board.
- <board_family_name> is the board family name that the library supports.

- `<unique_string>` is a designation that you create. Intel recommends that you include information such as revision and interface type.
- `<a|so|lib|dll>` is the file extension. It can be an archive file (.a), a shared object file (.so), a library file (.lib), or a dynamic link library file (.dll).

Example library file name: `altera_svdevkit_pcierev1_mmd.so`

2. Include the `aocl_mmd.h` header file (refer to `source/include` directory in one of the reference platform BSPs available at the [FPGA Software Download Center](#)) in the operating system-specific implementation of the MMD layer.

The `aocl_mmd.h` file and the *MMD API Descriptions* reference section contain full details on the MMD application programming interface (API) descriptions, their arguments, and their return values.

3. Implement the MMD layer for your Custom Platform, and compile it into a C/C++ library.

Example source codes of a functional MMD library are available in the `<path_to_a10_ref>/source/host/mmd` directory of the Intel Arria 10 GX FPGA Development Kit Reference Platform. In particular, the `acl_pcie.cpp` file implements the API functions defined in the `aocl_mmd.h` file.

If the SDK users need to load a particular library at runtime, deliver the library in a directory that the operating system can locate. Instruct the SDK users to add the library path to the `LD_LIBRARY_PATH` (for Linux) or `PATH` (for Windows) environment variable at runtime.

4. Modify the `mmdlib` and `linkflags` elements in the `board_env.xml` file by specifying the library flags necessary for linking with the MMD layer.

Related Information

[MMD API Descriptions](#) on page 40

1.4.3.1. Kernel Power-up State

The OpenCL kernel is an unknown state after you power-up your system or reprogram your FPGA. As a result, the MMD layer does not enable or respond to any interrupts from the kernel during these periods. The kernel is in a known state only after `aocl_mmd_set_interrupt_handler` is called. Therefore, enable interrupts from the kernel only after the handler becomes available to the MMD layer.

The general sequence of calls for a single host application is as follows:

1. `get_offline_info`
2. `open`
3. `get_info`
4. `set_status_handler`
5. `set_interrupt_handler`
6. `get_info /read/write/copy/yield`
7. `close`

1.4.4. Setting Up the FPGA Client Driver

The Intel FPGA SDK for OpenCL Pro Edition supports the FPGA Client Driver custom extension. The FPGA Client Driver (FCD) allows the SDK to automatically find and load the Custom Platform libraries at host run time.

Attention: To allow SDK users to use FCD, you must remove the MMD library from the `linklibs` element in the `board_env.xml` file.

Enumerating the Custom Platform FCD on Windows

Specify the Custom Platform libraries in the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\Intel\OpenCL\Boards`. Specify the value name to be the path to the library, and specify the data to be a `DWORD` that is set to 0.

For example:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Intel\OpenCL\Boards]
"c:\board_vendor a\my_board_mmd.dll"=dword:00000000
```

To enumerate Custom Platform FCDs on Windows, the FCD Loader scans the value in the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\Intel\OpenCL\Boards`. For each value in this key, the name specifies the path to a DLL and the data is a `dword`. If the `dword` data is 0, the Installable Client Driver (ICD) Loader attempts to open the corresponding DLL. If the DLL is an MMD library, then the SDK attempts to open any board that is associated with that library.

In this case, the FCD opens the library `c:\board_vendor a\my_board_mmd.dll`.

If the registry key specifies multiple libraries, the Loader loads the libraries in the order that they appear in the key. If there is an order dependency between the libraries available with your Custom Platform, ensure that you list the libraries accordingly in the registry key.

Enumerating the Custom Platform FCD on Linux

Enter the absolute paths of Custom Platform libraries in a `.fcd` file (formerly, a `.acd` file). Store the `.fcd` file in the `/opt/Intel/OpenCL/Boards/` directory.

To enumerate Custom Platform FCDs on Linux, the FCD Loader scans the files with the extension `.fcd` in the path `/opt/Intel/OpenCL/Boards/`. The FCD Loader opens each `.fcd` file in this path as a text file. Each `.fcd` file should contain the absolute path to every library in the Custom Platform, one library per line. The ICD Loader attempts to open each library. If the library is an MMD library, then the SDK attempts to open any board that is associated with that library.

For example, consider the file `/opt/Intel/OpenCL/Boards/PlatformA.fcd`. If it contains the line `/opt/PlatformA/libPlatformA_mmd.so`, the FCD Loader loads the library `/opt/PlatformA/libPlatformA_mmd.so`.

If the `.fcd` file specifies multiple libraries, the Loader loads the libraries in the order that they appear in the file. If there is an order dependency between the libraries available with your Custom Platform, ensure that you list the libraries accordingly in the `.fcd` file.

Starting in the Intel FPGA SDK for OpenCL version 16.1, FCD allows the SDK users to load multiple heterogeneous boards from different Custom Platforms and use them together in a single host application. However, this feature has limited support for legacy Custom Platforms that are released prior to 16.1.

For more information about how SDK users link their host applications to the ICD and FCD, refer to the *Linking Your Host Application to the Khronos ICD Loader Library* section in the *Intel FPGA SDK for OpenCL Pro Edition Programming Guide*.

Related Information

[Linking Your Host Application to the Khronos ICD Loader Library](#)

1.4.5. Providing Intel FPGA SDK for OpenCL Utilities Support

Each Custom Platform you develop for use with the Intel FPGA SDK for OpenCL Pro Edition must support a set of SDK utilities. These utilities enable users to manage the accelerator board through the SDK.

If you create a new Custom Platform, perform the following tasks to create executables of the SDK utilities and then store them in the `utilbindir` directory of your Custom Platform:

Tip:

Within the `<path_to_a10_ref>/source/util` directory of the Intel Arria 10 10 GX FPGA Development Kit Reference Platform, you can find source code for the `program` and `flash` utilities in the `reprogram` and `flash` subdirectories, respectively. Scripts for the `install` and `uninstall` utilities are available in the `<path_to_a10_ref>/<OS_platform>/libexec` directory.

You can find the source code for the `diagnose` utility in the `<path_to_a10_ref>/source/util/diagnostic` directory within the Intel Arria 10 GX FPGA Development Kit Reference Platform.

1. Create an `install` utility executable that sets up the current host computer with the necessary drivers to communicate with the board via the MMD layer. The `install` utility takes no argument.

For example, the PCIe-based MMD might need to install PCIe drivers into the host operating system.

Executable call: `aocl install <path_to_customplatform>`

2. Create an `uninstall` utility executable that removes the current host computer drivers (for example, PCIe drivers) used for communicating with the board. The `uninstall` utility takes no argument.

Executable call: `aocl uninstall <path_to_customplatform>`

3. Create a `diagnose` utility executable that confirms the board's integrity and the functionality of the MMD layer.

The diagnose utility must support the following internal calling modes:

Calling Mode	Description
-probe	Prints all available devices in a Custom Platform. For a given hardware configuration, the utility lists the devices in the same order, and each device is associated with the same identification string each time.
-probe <device_name>	Queries the specified device and prints summary statistics about the device.
<device_name> where <device name> is the string that corresponds to the FPGA device.	Performs a full diagnostic test for the specified device. The utility generates the message DIAGNOSTIC_PASSED as the output. Otherwise, the utility generates the message DIAGNOSTIC_FALIED.

When users invoke the diagnose utility command without an argument, it queries the devices in the Custom Platform and supplies a list of valid <device_name> strings assigned to the list of devices.

Executable call without argument: `aocl diagnose`

When users invoke the diagnose utility command with a <device_name> argument, the utility runs your diagnostic test for the board. A user may give a board a different logical device name than the physical device name associated with the Custom Platform. The aocl utility simply converts the user-side logical device name to the Custom Platform-side physical device name. If the diagnostic test runs successfully, the utility generates the message DIAGNOSTIC_PASSED as the output. Otherwise, the utility generates the message DIAGNOSTIC_FALIED.

Executable call with argument: `aocl diagnose <device_name>`.

4. Create a program utility executable that receives the fpga.bin file and the Intel FPGA SDK for OpenCL Offline Compiler executable file (.aocx) and configures that design onto the FPGA. Although the main method for FPGA programming is via the host and the MMD, make this utility available to users who do not have a host system or who perform offline reprogramming.

The program utility command takes <device_name>, fpga.bin, and <kernel_filename>.aocx as arguments. When users invoke the command, the SDK extracts the fpga.bin file and passes it to the program utility.

Important: Intel highly recommends that the program utility links with and calls the aocl_mmd_program function implemented in the MMD layer. Refer to the aocl_mmd_program and Program Support reference sections for more information.

Executable call: `aocl program <device_name> <kernel_filename>.aocx`, where <device name> is the acl number that corresponds to the FPGA device.

5. Create a flash utility executable that receives the fpga.bin file and programs that design into the flash memory on the board. The flash utility command takes <device_name> and a .aocx file name as arguments. When users invoke the command, the SDK extracts the fpga.bin file and passes it to the flash utility.

Executable call: `aocl flash <device_name> <kernel_filename>.aocx`, where <device name> is the acl number that corresponds to the FPGA device.

When users invoke a utility command, the utility probes the current Custom Platform's board_env.xml file and executes the <utility_executable> file within the utilbindir directory.

Related Information

- [Creating the board_env.xml File](#) on page 15
- [Program Support](#) on page 53
- [aocl_mmd_program](#) on page 52

1.4.6. Testing the Hardware Design

After you create the software utilities and the MMD layer, and your hardware design achieves timing closure, test the design.

To test the hardware design, perform the following tasks:

1. Navigate to the `boardtest.cl` OpenCL kernel within the `INTELFPGAOCCLSDKROOT/board/custom_platform_toolkit/tests/boardtest` directory.
`INTELFPGAOCCLSDKROOT` points to the location of the Intel FPGA SDK for OpenCL Pro Edition installation.
2. Compile your kernel to generate an Intel FPGA SDK for OpenCL Offline Compiler executable file (`.aocx`) by invoking the `aoc -no-interleaving=default boardtest.cl` command.
3. Program the accelerator board by invoking the `aocl program aocl0 boardtest.aocx` command.
4. Invoke the commands `aocl compile-config` and `aocl link-config`. Confirm they include flags necessary for your MMD layer to compile and link successfully.
5. Build the `boardtest` host application.
 - For Windows systems, you may invoke the `make -f Makefile.windows` command or use Microsoft Visual Studio.
If you invoke the `make` command, use the `Makefile.windows` file located in the `INTELFPGAOCCLSDKROOT/board/custom_platform_toolkit/tests/boardtest` directory.
If you build your host application in Microsoft Visual Studio, the `boardtest.sln` and `main.cpp` files are located in the `INTELFPGAOCCLSDKROOT/board/custom_platform_toolkit/tests/boardtest/host` directory.
 - For Linux systems, invoke the `make -f Makefile.linux` command.
The `Makefile.linux` file is located in the `INTELFPGAOCCLSDKROOT/board/custom_platform_toolkit/tests/boardtest` directory.
6. Run the `boardtest` executable.

Attention: To ensure that your hardware design produces consistent performance, you might have to test it using multiple OpenCL kernels in addition to `boardtest.cl`.

To qualify as an Intel FPGA preferred board, rigorous testing across multiple boards is necessary. Specifically, you should perform overnight testing of all Custom Platform tests and executes the SDK example designs on multiple boards. All board variants within a Custom Platform must go through the testing process.

1.5. Applying for the Intel FPGA SDK for OpenCL Pro Edition Preferred Board Status

Registering your Custom Platform and the supported FPGA boards in the Intel FPGA SDK for OpenCL Pro Edition Preferred Board Partner Program allows them to benefit from ongoing internal testing across versions of the Intel Quartus Prime Design Suite. Intel-tested Custom Platforms and boards are more likely to be forward compatible with future Intel Quartus Prime Design Suite versions.

For your Custom Platform and the supported FPGA boards to achieve the Intel FPGA SDK for OpenCL Preferred Board status, you must generate the following data and submit it to Intel:

1. The output from invoking the `aocl board-xml-test` command.
2. The output from invoking the `aoc -list-boards` command.
3. The outputs from the host compilation, host execution, and all Intel Quartus Prime report files (`.rpt`). Also, for each board in your Custom Platform, the `acl_quartus_report.txt` file from the following tests:
 - a. All tests included in the `INTELFPGAOCSDKROOT/board/custom_platform_toolkit/tests` directory, where `INTELFPGAOCSDKROOT` points to location of the Intel FPGA SDK for OpenCL Pro Edition installation.
 - b. Compilations of the following examples on the OpenCL Design Examples page of the Intel FPGA website:
 - i. Vector Addition
 - ii. Matrix Multiplication
 - iii. FFT (1D)
 - iv. FFT (2D)
 - v. Sobel Filter
 - vi. Finite Difference Computation (3D)
4. For each board in the Custom Platform, a summary of the following:
 - a. `HOST-TO-MEMORY BANDWIDTH` as reported by the `boardtest` test in the Custom Platform Toolkit (`/tests/boardtest`).
 - b. `KERNEL-TO-MEMORY BANDWIDTH` as reported by the `boardtest` test.
 - c. Throughput in `swap-and-execute(s)` reported by the `swapper` test in the Custom Platform Toolkit (`/tests/swapper`).
 - d. Actual `clock freq` as reported in the `acl_quartus_report.txt` file from the blank test in the Custom Platform Toolkit (`INTELFPGAOCSDKROOT/board/custom_platform_toolkit/tests/blank`).

Important: Use global routing to reduce consumption of local routing resources. Using global routing is necessary because it helps meet timing and improve kernel performance (Fmax). Use global or regional routing for any net with fan-out greater than 5000, and for kernel clock, 2x clock and reset. Check the Non-Global High Fan-Out Signals report in the Resource subsection, which is under the Fitter section of the Compilation Report.

5. Submit the necessary number of boards to Intel for in-house regression testing. Regression testing tests the out-of-the-box experience for each board variant on each operating system that your Custom Platform supports. Ensure that you test the procedure outlined below before you submit your boards:
 - a. Install the board into the physical machine.
 - b. Boot the machine and invoke the `aocl install <path_to_customplatform>` utility command.
 - c. Invoke the `aocl diagnose` command.
 - d. Run the SDK's test programs. The tester can also invoke the `aocl program <device_name> <kernel_filename>.aocx` command to verify the functionality of the `program` utility.

Related Information

[OpenCL Design Examples on the Intel FPGA website](#)

1.6. Shipping Recommendations

Before shipping your Intel-verified board to Intel FPGA SDK for OpenCL users, program the flash memory of the board with the `hello_world` OpenCL design example. Programming the flash memory of the board with the `hello_world.aocx` hardware configuration file allows the SDK user to power on the board and observe a working kernel.

Use the `hello_world` OpenCL design example available in the `$INTELFPGAOCSDKROOT/examples_aoc` directory.

For more information, refer to the `README` file available with the `hello_world` OpenCL design example and the *Programming the Flash Memory of an FPGA* sections in the *Intel FPGA SDK for OpenCL Pro Edition Getting Started Guide*.

Related Information

- [OpenCL Design Examples on the Intel FPGA website](#)
- [Programming the Flash Memory of an FPGA on Windows](#)
- [Programming the Flash Memory of an FPGA on Linux](#)


2. Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit Reference Material

The *Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit Reference Material* chapter provides supplementary information that can assist you in the implementation of your Custom Platform.

2.1. The Board Platform Designer Subsystem

When designing your board hardware, you have the option to create a Platform Designer subsystem within the top-level Platform Designer system (`board.qsys`) that contains all nonkernel logic. The board Platform Designer subsystem is the main design entry point for a new accelerator board. It is the location where the instantiations of the OpenCL host and global memory interfaces occur. Your board design must have a minimum of 128 kilobytes (KB) of external memory. Any Avalon Memory-Mapped (Avalon-MM) agent interface (for example, a block RAM) can potentially be a memory interface.

The diagram below represents a board system implementation in more details:



Note: Blocks denoted with an asterisk (*) are blocks that you have to add to the board Platform Designer subsystem.

The OpenCL host communication interface and global memory interface are the main components of the board system. The memory-mapped device (MMD) layer communicates, over some medium, with the intellectual property (IP) core instantiated in this Platform Designer system.

For example, an MMD layer executes on a PCI Express (PCIe)-based host interface, and the host interface generates Avalon interface requests from an Intel PCIe endpoint on the FPGA.

Within the board Platform Designer subsystem, you can also define the global memory system available to the OpenCL kernel. The global memory system may consist of different types of memory interfaces. Each memory type may consist of one, two, four, or eight banks of physical memory. All the banks of a given memory type must be the same size in bytes and have equivalent interfaces. If you have streaming I/O, you must also include the corresponding IP in the board Platform Designer system. In addition, you must update the `board_spec.xml` file to describe the channel interfaces.

2.1.1. Intel FPGA SDK for OpenCL Pro Edition-Specific Platform Designer System Components

The Platform Designer system for your board logic includes components specific to the Intel FPGA SDK for OpenCL Pro Edition that are necessary for implementing features that instantiate host communication and global memory interfaces.

The board Platform Designer system must export an Avalon-MM host for controlling OpenCL kernels. It must also export one or more Avalon-MM agent ports that the kernels use as global memory interfaces. The `INTELFPGAOCCLSDKROOT/ip/board` directory of the SDK includes a library that contains SDK-specific Platform Designer system components, where `INTELFPGAOCCLSDKROOT` points to the location of the SDK installation. These components are necessary for implementing features such as Avalon-MM interfaces, organizing programmable banks, cache snooping, and supporting Intel's guaranteed timing closures.

1. [OpenCL Kernel Clock Generator](#) on page 30
2. [OpenCL Kernel Interface](#) on page 31
3. [OpenCL Memory Bank Divider](#) on page 32

2.1.1.1. OpenCL Kernel Clock Generator

The OpenCL Kernel Clock Generator is a Platform Designer component that generates a clock output and a clock 2x output for use by the OpenCL kernels. An Avalon-MM agent interface allows reprogramming of the phase-locked loops (PLLs) and kernel clock status information.

Table 4. Parameter Settings for the OpenCL Kernel Clock Generator Component

Parameter	Description
REF_CLK_RATE	Frequency of the reference clock that drives the kernel PLL (that is, <code>pll_refclk</code>).
KERNEL_TARGET_CLOCK_RATE	Frequency that the Intel Quartus Prime Pro Edition software attempts to achieve during compilation. Keep this parameter at its default setting.

Table 5. Signals and Ports for the OpenCL Kernel Clock Generator Component

Signal or Port	Description
kernel_pll_refclk	The reference clock for the kernel PLL. The frequency of this clock must match the frequency you specify for the REF_CLK_RATE component parameter.
clk	The clock used for the host control interface. The clock rate of <code>clk</code> can be slow.
reset	The reset signal that resets the PLL and the control logic. Resetting the PLL disables the kernel clocks temporarily. Connect this reset signal to the power-on reset signal in your system.
ctrl	The agent port used to connect to the OpenCL host interface and to adjust the frequency based on the OpenCL kernel.
kernel_clk kernel_clk2x	The kernel clock and its 2x variant that runs on twice the speed. The <code>kernel_clk2x</code> signal is directly exported from this interface. Because <code>kernel_clk</code> has internal Platform Designer connections, export it using a clock source component. You can also use the clock source to export the kernel reset. In addition, clock all logic at the board Platform Designer system interface with <code>kernel_clk</code> , except for any I/O that you add.
kernel_pll_locked	(Optional) If the PLL is locked onto the reference clock, the value of this signal is 1. The host interface manages this signal normally; however, this signal is made available in the board Platform Designer system.

2.1.1.2. OpenCL Kernel Interface

The OpenCL Kernel Interface is a Platform Designer component that allows the host interface to access and control the OpenCL kernel.

Table 6. Parameter Settings for the OpenCL Kernel Interface Component

Parameter	Description
Number of global memory systems	Number of global memory types in your board design.

Table 7. Signals and Ports for the OpenCL Kernel Interface Component

Signal or Port	Description
clk	The clock input used for the host control interface. The clock rate of <code>clk</code> can be slow.
reset	This reset input resets the control interface. It also triggers the <code>kernel_reset</code> signal, which resets all kernel logic.
ctrl	Use this agent port to connect to the OpenCL host interface. This interface is a low-speed interface with which you set kernel arguments and start the kernel's execution.
kernel_clk	The <code>kernel_clk</code> output from the OpenCL Kernel Clock Generator drives this clock input.
kernel_cra	This Avalon-MM host interface communicates directly with the kernels generated by the Intel FPGA SDK for OpenCL Offline Compiler. Export the Avalon-MM interface to the OpenCL Kernel Interface and name it in the <code>board_spec.xml</code> file.
sw_reset_in	When necessary, the OpenCL host interface resets the kernel via the <code>ctrl</code> interface. If the board design requires a kernel reset, it can do so via this reset input. Otherwise, connect the interface to a global power-on reset.
kernel_reset	Use this reset output to reset the kernel and any other hardware that communicates with the kernel.

continued...

Signal or Port	Description
	Warning: This reset occurs between the MMD open and close calls. Therefore, it must not reset anything necessary for the operation of your MMD.
sw_reset_export	This reset output is the same as kernel_reset, but it is synchronized to the clk interface. Use this output to reset logic that is not in the kernel_clk clock domain but should be reset whenever the kernel resets.
acl_bsp_memorg_host	The memory interfaces use these signals. Based on the number of global memory systems you specify in the OpenCL Kernel Interface component parameter editor, the Intel Quartus Prime Pro Edition software creates the corresponding number of copies of this signal, each with a different hexadecimal suffix. Connect each signal to the OpenCL Memory Bank Divider component associated with each global memory system (for example, DDR). Then, list the hexadecimal suffix in the config_addr attribute of the global_mem element in the board_spec.xml file.
kernel_irq_from_kernel	An interrupt input from the kernel. This signal is exported and named in the board_spec.xml file.
kernel_irq_to_host	An interrupt output from the kernel. This signal connects to the host interface.

2.1.1.3. OpenCL Memory Bank Divider

The OpenCL Memory Bank Divider is a Platform Designer component that takes an incoming request from the host interface on the Avalon-MM agent port and routes it to the appropriate bank host port. This component must reside on the path between the host and the global memory interfaces. In addition, it must reside outside of the path between the kernel and the global memory interfaces.

Table 8. Parameter Settings for the OpenCL Memory Bank Divider Component

Parameter	Description
Number of banks	Number of memory banks for each of the global memory types included in your board system.
Separate read/write ports	Enable this parameter so that each bank has one port for read operation and one for write operation.
Add pipeline stage to output	Enable this parameter to allow for potential timing improvements.
Data Width	Width of the data bus to the memory in bits.
Address Width (total addressable)	Total number of address bits necessary to address all global memory.
Burst size (maximum)	Set to a value equal to $\text{interleaved_bytes}/(\text{width}/8)$, where interleaved_bytes and width are defined in the interface attribute of the global_mem element in the board_spec.xml file.
Maximum Pending Reads	Maximum number of pending read transfers the component can process without asserting a waitrequest signal. Intel recommended value is 64 if BSP has two global memory banks or fewer and 128 if BSP has four or more global memory banks. Caution: A high Maximum Pending Reads value causes Platform Designer to insert a deep response FIFO buffer, between the component's host and agent, that consumes a lot of device resources. It also increases the achievable bandwidth between host and memory interfaces.
Split read/write bursts on burst word boundary	Enable splitting of read and write bursts on burst word boundary. Enable this parameter if the Number of banks parameter value is greater than 1, and the burst reads and writes that the host controller sends to the agent port crosses burst word boundary.

Table 9. Signals and Ports for the OpenCL Memory Bank Divider Component

Signal or Port	Description
clk	The bank divider logic uses this clock input. If the IP of your host and memory interfaces have different clocks, ensure that clk clock rate is not slower than the slowest of the two IP clocks.
reset	The reset input that connects to the board power-on reset.
s	The agent port that connects to the host interface controller.
kernel_clk	The kernel_clk output from the OpenCL Kernel Clock Generator drives this clock input.
kernel_reset	The kernel_reset output from the OpenCL Kernel Interface drives this reset input.
acl_bsp_snoop	Export this Avalon Streaming (Avalon-ST) source. In the board_spec.xml file, under interfaces, describe only the snoop interface for the default memory (acl_internal_snoop). If you have a heterogeneous memory design, perform these tasks only for the OpenCL Memory Bank Divider component associated with the default memory. <i>Important:</i> The memory system you build in Platform Designer alters the width of acl_bsp_snoop. You must update the width of the streamsource interface within the channels element in the board_spec.xml file to match the width of acl_bsp_snoop. <i>Important:</i> In the board_spec.xml file, update the width of the snoop interface (acl_internal_snoop) specified with the streamsource kernel interface within the interfaces element. Updating the width ensures that the global_mem interface entries in board_spec.xml match the characteristics of the bank<N> Avalon-MM hosts from corresponding OpenCL Memory Bank Divider component for the default memory.
acl_bsp_memorg_host	This conduit connects to the acl_bsp_memorg_host interface of the OpenCL Kernel Interface. <i>Note:</i> Signal present if Number of banks > 1. .
bank1, bank2, ..., bank8	The number of memory hosts available in the OpenCL Memory Bank Divider depends on the number of memory banks that were included when the unit was instantiated. Connect each bank with each memory interface in the same order as the starting address for the corresponding kernel memory interface specified in the board_spec.xml file. For example, global_mem interface that begins at address 0 must correspond to the memory host in bank1 from the OpenCL Memory Bank Divider.

Related Information

- [channels](#) on page 37
- [interfaces](#) on page 37
- [global_mem](#) on page 35
- [OpenCL Kernel Interface](#) on page 31

2.2. XML Elements, Attributes, and Parameters in the board_spec.xml File

This section describes the metadata you must include in the board_spec.xml file.

[board](#) on page 34

- [device](#) on page 34
- [global_mem](#) on page 35
- [host](#) on page 37
- [channels](#) on page 37
- [interfaces](#) on page 37
- [interface](#) on page 38
- [compile](#) on page 39

2.2.1. board

The `board` element of the `board_spec.xml` file provides the version and the name of the accelerator board.

Example eXtensible Markup Language (XML) code:

```
<board version="<version>" name="<Custom_Platform_name>">
...
</board>
```

Table 10. Attributes for the board Element

Attribute	Description
version	The version of the board. The board version must match the version of the Intel Quartus Prime Pro Edition software you use to develop the Custom Platform.
name	The name of the accelerator board, which must match the name of the directory in which the <code>board_spec.xml</code> file resides. The name must contain a combination of only letters, numbers, underscores (<code>_</code>), hyphens (<code>-</code>), or periods (<code>.</code>) (for example, <code>a10_ref</code>).

2.2.2. device

The `device` element of the `board_spec.xml` file provides the device model and the resources that the board design uses.

Example XML code:

```
<device device_model="5sgsed8k2f40c2_dm.xml">
  <used_resources>
    <alms num="45000"/>
    <!-- ALMs used in final placement - ALMs used for registers -->
    <ffs num="117500"/>
    <dsps num="0"/>
    <rams num="511"/>
  </used_resources>
</device>
```

Table 11. Attributes for the device Element

Attribute	Description
device_model	The file name of the device model file that describes the available FPGA resources on the accelerator board.
used_resources	<p>Reports the number of adaptive logic modules (ALMs), flip-flops, digital signal processor (DSP) blocks and RAM blocks that the board design consumes, in the absence of any kernel, to the Intel FPGA SDK for OpenCL Pro Edition. If you create a defined partition around all the board logic, you can obtain the used resources data from the Partition Statistics section of the Fitter report.</p> <p>Extract the information from the following parameters:</p> <ul style="list-style-type: none"> alms num—The number of logic ALMs used, excluding the number of ALMs with only their registers used. The value should correspond to [a]+[b]+[d] from part [A] of the Fitter Partition Statistics. ffs num—The number of flip flops. dsps num—The number of DSP blocks. rams num—The number of RAM blocks.

2.2.3. global_mem

The `global_mem` and `interface` elements of the `board_spec.xml` file provides information on the memory interfaces that connect to the kernel.

Example XML code:

```

<!-- DDR3-1600 -->
<global_mem name="DDR" max_bandwidth="25600" interleaved_bytes="1024"
config_addr="0x018">
  <interface name="board" port="kernel_mem0" type="agent" width="512"
maxburst="16"
  address="0x00000000" size="0x100000000" latency="240"/>
  <interface name="board" port="kernel_mem1" type="agent" width="512"
maxburst="16"
  address="0x100000000" size="0x100000000" latency="240"/>
</global_mem>

<!-- QDR II -->
<global_mem name="QDR" max_bandwidth="17600" interleaved_bytes="8"
config_addr="0x100">
  <interface name="board" type="agent" width="64" maxburst="1"
  address="0x200000000" size="0x1000000" latency="150" addpipe="1">
    <port name="kernel_qdr0_r" direction="r"/>
    <port name="kernel_qdr0_w" direction="w"/>
  </interface>
  <interface name="board" type="agent" width="64" maxburst="1"
  address="0x201000000" size="0x1000000" latency="150" addpipe="1">
    <port name="kernel_qdr1_r" direction="r"/>
    <port name="kernel_qdr1_w" direction="w"/>
  </interface>
  <interface name="board" type="agent" width="64" maxburst="1"
  address="0x202000000" size="0x1000000" latency="150" addpipe="1">
    <port name="kernel_qdr2_r" direction="r"/>
    <port name="kernel_qdr2_w" direction="w"/>
  </interface>
  <interface name="board" type="agent" width="64" maxburst="1"
  address="0x203000000" size="0x1000000" latency="150" addpipe="1">
    <port name="kernel_qdr3_r" direction="r"/>
    <port name="kernel_qdr3_w" direction="w"/>
  </interface>
</global_mem>

```

Note: For each global memory that the kernel accesses, you must include one interface element that describes its characteristics.

Table 12. Attributes for the global_mem Element

Attribute	Description
name	The name the Intel FPGA SDK for OpenCL Pro Edition user uses to identify the memory type. Each name must be unique and must comprise of less than 32 characters.
max_bandwidth	<p>The maximum bandwidth, in megabytes per second (MB/s), of all global memory interfaces combined in their current configuration. The Intel FPGA SDK for OpenCL Offline Compiler uses <code>max_bandwidth</code> to choose an architecture suitable for the application and the board.</p> <p>Compute this bandwidth value from datasheets of your memories.</p> <p>Example <code>max_bandwidth</code> calculation for a 64-bit DDR3 interface running at 800 MHz: $max_bandwidth = 800 \text{ MHz} \times 2 \times 64 \text{ bits} \div 8\text{-bits} = 12800 \text{ MB/s}$</p> <p>You have the option to use block RAM instead of or in conjunction with external memory as global memory. The formula for calculating <code>max_bandwidth</code> for block RAM is $max_bandwidth = \text{block RAM speed} \times (\text{block RAM interface size} \div 8 \text{ bits})$</p> <p>Example <code>max_bandwidth</code> calculation for a 512-bit block RAM running at 100 MHz: $max_bandwidth = 100 \text{ MHz} \times 512 \text{ bits} \div 8 \text{ bits} = 6400 \text{ MB/s}$</p>
interleaved_bytes	<p>Include the <code>interleaved_bytes</code> attribute in the <code>board_spec.xml</code> file when you instantiate multiple interfaces for a given global memory system. This attribute controls the size of data that the offline compiler distributes across the interfaces.</p> <p>The offline compiler currently can interleave data across banks no finer than the size of one full burst. This attribute specifies this size in bytes and following are the recommended values:</p> <ul style="list-style-type: none"> • For BSP with two or fewer global memory banks: $maxburst \times width_bytes$ • For BSP with four or more global memory banks: $maxburst \times width_bytes \times 4$ <p>The <code>interleaved_bytes</code> value must be the same for the host interface and the kernels. Therefore, the configuration of the OpenCL Memory Bank Divider must match the exported kernel agent interfaces in this respect.</p> <p>For block RAM, <code>interleaved_bytes</code> equals the width of the interface in bytes.</p>
config_addr	<p>The address of the ACL Mem Organization Control Platform Designer component (<code>mem_org_mode</code>) that the host software uses to configure memory. You may omit this attribute if your board has homogeneous memory; the software uses the default address (0x18) for this component. If your board has heterogeneous memory, there is a <code>mem_org_mode</code> component in the board system for each memory type.</p> <p>Enter the <code>config_addr</code> attribute and set it to the value of the base address of the <code>mem_org_mode</code> component(s).</p>
default	<p>Include this optional attribute and assign a value of 1 to set the global memory as the default memory interface. The default memory must start at address 0x0.</p> <p>If you do not implement this attribute, the first memory type defined in the <code>board_spec.xml</code> file becomes the default memory interface.</p>
interface	See the <i>interface</i> section for the parameters you must specify for each interface.
allocation_type	<p>A list that specifies which USM allocator is used to allocate from the <code>global_mem</code> element. Values allowed in this list are <code>host</code>, <code>shared</code>, and <code>device</code>. The following conditions apply:</p> <ul style="list-style-type: none"> • If there are multiple <code>global_mem</code> elements with the same <code>allocation_type</code> attribute, the first <code>allocation_type</code> attribute in the <code>board_spec.xml</code> is assumed to be the one used by the specified allocator. • If there is a single <code>global_mem</code> element with multiple <code>allocation_type</code> attributes, this indicates that allocations of the specified types use this <code>global_mem</code> interface. • [Legacy support] If you have not specified the <code>allocation_type</code> attribute, it is assumed that all global memory interfaces have the <code>device</code> <code>allocation_type</code>.

Related Information

[interface](#) on page 38

2.2.4. host

The `host` element of the `board_spec.xml` file provides information on the interface from the host to the kernel.

Example XML code:

```
<host>
  <kernel_config start="0x00000000" size="0x0100000"/>
</host>
```

Table 13. Attributes for the host Element

Attribute	Description
kernel_config	This attribute informs the Intel FPGA SDK for OpenCL Offline Compiler at what offset the kernel resides, from the perspective of the <code>kernel_cra</code> host on the <code>kernel_interface</code> module. start: the starting address of the kernel. Normally, this attribute has a value of 0 because the <code>kernel_cra</code> host should not host anything except kernels. size: keep this parameter at the default value of 0x0100000.

2.2.5. channels

The Intel FPGA SDK for OpenCL Pro Edition supports data streaming directly between kernels and I/O via explicitly named *channels*. Include the `channels` element in the `board_spec.xml` file if your accelerator board provides channels for direct kernel-to-I/O accesses. For the `channels` element, you must identify all the channel interfaces, which are implemented using the Avalon-ST specification. Specify each channel interface via the `interface` attribute. Refer to the *interface* section for the parameters you must specify for each interface. The channel interface only supports data, and valid and ready Avalon-ST signals. The I/O channel defaults to 8-bit symbols and big-endian ordering at the interface level.

Example XML code:

```
<channels>
  <interface name="udp_0" port="udp0_out" type="streamsource" width="256"
    chan_id="eth0_in"/>
  <interface name="udp_0" port="udp0_in" type="streamsink" width="256"
    chan_id="eth0_out"/>
  <interface name="udp_0" port="udp1_out" type="streamsource" width="256"
    chan_id="eth1_in"/>
  <interface name="udp_0" port="udp1_in" type="streamsink" width="256"
    chan_id="eth1_out"/>
</channels>
```

Related Information

[interface](#) on page 38

2.2.6. interfaces

The `interfaces` element of the `board_spec.xml` file describes the kernel interfaces that connect to OpenCL kernels and control their behaviors. For this element, include one of each interface of types `host`, `irq` and `streamsource`. Refer to the *interface* section for the parameters you must specify for each interface.

Example XML code:

```
<interfaces>
  <interface name="board" port="kernel_cra" type="host" width="64" misc="0"/>
  <interface name="board" port="kernel_irq" type="irq" width="1"/>
  <interface name="board" port="acl_internal_snoop" type="streamsourced"
    enable="SNOOPENABLE" width="31" clock="board.kernel_clk"/>
  <kernel_clk_reset clk="board.kernel_clk" clk2x="board.kernel_clk2x"
    reset="board.kernel_reset"/>
</interfaces>
```

In addition to the `host`, `irq`, and `streamsourced` interfaces, if your design includes a separate Platform Designer subsystem containing the board logic, the kernel clock and reset interfaces exported from it are also part of the `interfaces` element. Specify these interfaces with the `kernel_clk_reset` attribute and its corresponding parameters.

Table 14. Parameters for the `kernel_clk_reset` Attribute

Important: Name the kernel clock and reset interfaces in the Platform Designer connection format (that is, `<instance_name>.<interface_name>`).

For example: `board.kernel_clk`

Attribute	Description
<code>clk</code>	The Platform Designer name for the kernel clock interface. The <code>kernel_clk</code> output from the OpenCL Kernel Clock Generator component drives this interface.
<code>clk2x</code>	The Platform Designer name for the kernel clock interface. The <code>kernel_clk2x</code> output from the OpenCL Kernel Clock Generator component drives this interface.
<code>reset</code>	The Platform Designer connection for the kernel reset. The <code>kernel_reset</code> output from the OpenCL Kernel Interface component drives this interface.

Related Information

[interface](#) on page 38

2.2.7. interface

In the `board_spec.xml` file, each global memory, channel or kernel interface is comprised of individual interfaces. For the `global_mem`, `channels`, and `interfaces` XML elements, include an `interface` attribute for each interface and specify the corresponding parameters.

Table 15. Parameters for the `interface` XML Attribute

Parameter	Applicable Interface	Description
<code>name</code>	All	For <code>global_mem</code> : instance name of the Platform Designer component. For <code>channels</code> : instance name of the Platform Designer component that has the channel interface. For <code>interfaces</code> : name of the entity in which the kernel interface resides (for example, <code>board</code>).
<code>port</code>		For <code>global_mem</code> : name of the Avalon-MM interface in the Platform Designer component that corresponds to the <code>interface</code> attribute. For <code>channels</code> : name of the streaming interface in the Platform Designer component.

continued...

Parameter	Applicable Interface	Description
		For interfaces: name of the interface to the OpenCL Kernel Interface Platform Designer component. For example, <code>kernel_cra</code> is the Avalon-MM interface, and <code>kernel_irq</code> is an interrupt.
<code>type</code>		For <code>global_mem</code> : set to <code>agent</code> . For channels: <ul style="list-style-type: none"> Set to <code>streamsource</code> for a stream source that provides data to the kernel. Set to <code>streamsink</code> for a stream sink interface that consumes data from the kernel. For interfaces: set to either <code>host</code> , <code>irq</code> , or <code>streamsource</code> .
<code>width</code>		For <code>global_mem</code> : width of the memory interface in bits. For channels: number of bits in the channel interface. For interfaces: width of the kernel interface in bits.
<code>waitrequest_allowance</code>		For <code>global_mem</code> : [Optional] Amount of Avalon-MM <code>waitrequest</code> allowance supported on the agent interface (that is, kernel-facing interface) of the clock-crossing bridge that spans between the memory and the kernel clock domains. For <code>kernel_cra</code> : [Optional] Amount of Avalon-MM <code>waitrequest</code> allowance that the <code>kernel_cra</code> agent interface must support. This parameter defaults to 0 if you do not specify it in the <code>board_spec.xml</code> file. A value of 0 indicates that this <code>waitrequest</code> allowance feature is disabled.
<code>maxburst</code>	<code>global_mem</code>	Maximum burst size for the agent interface. Attention: The value of <code>width ÷ 8 × maxburst</code> must be less than the value of <code>interleaved_bytes</code> .
<code>address</code>		Starting address of the memory interface that corresponds to the host interface-side address. For example, address 0 should correspond to the <code>bank1</code> memory host from the OpenCL Memory Bank Divider. In addition, any non-zero starting address must abut the end address of the previous memory.
<code>size</code>		Size of the memory interface in bytes. The sizes of all memory interfaces should be equal.
<code>latency_type</code>		If the memory interface has variable latency, set this parameter to <code>average</code> to signify that the specified latency is considered the average case. If the complete kernel-to-memory path has a guaranteed fixed latency, set this parameter to <code>fixed</code> .
<code>chan_id</code>	<code>channels</code>	A string used to identify the channel interface. The string may have up to 128 characters.
<code>clock</code>	<code>interfaces</code>	For the <code>streamsource</code> kernel interface type, the parameter specifies the name of the clock that the snoop stream uses. Usually, this clock is the kernel clock.

2.2.8. compile

The `compile` element of the `board_spec.xml` file and its associated attributes and parameters describe the general control of Intel Quartus Prime compilation, registration, and automigration.

Example XML code:

```
<compile project="top" revision="top" qsys_file="none" generic_kernel="1">
  <generate cmd="echo"/>
  <synthesize cmd="quartus_sh -t import_compile.tcl"/>
  <auto_migrate platform_type="a10_ref" >
```

```
<include fixes="" />
<exclude fixes="" />
</auto_migrate>
</compile>
```

Attribute	Description
project	Name of the Intel Quartus Prime project file (.qpf) that the Intel Quartus Prime Pro Edition software intends to compile.
revision	Name of the revision within the Intel Quartus Prime project that the Intel Quartus Prime Pro Edition software compiles to generate the Intel FPGA SDK for OpenCL Offline Compiler executable file (.aocx).
qsys_file	Name of the Platform Designer file into which the OpenCL kernel is embedded. You have the option to assign a value of "none" to qsys_file if you do not require the Intel Quartus Prime Pro Edition software to create a top-level .qsys file for your design. In this scenario, the Intel FPGA SDK for OpenCL Offline Compiler adds a .qip file into the Intel Quartus Prime project. The Custom Platform must manually instantiate the generated HDL entity in the kernel_system.v file.
generic_kernel	Set this value to 1 if you want the offline compiler to generate a common Verilog interface for all OpenCL compilations. This setting is necessary in situations where you must set up Intel Quartus Prime design partitions around the kernel, such as in the Configuration via Protocol (CvP) flow.
generate_cmd	Command required to prepare for full compilation, such as to generate the Verilog files for the Platform Designer system into which the OpenCL kernel is embedded.
synthesize_cmd	Command required to generate the fpga.bin file from the Custom Platform. Usually, this command instructs the Intel Quartus Prime Pro Edition software to perform a full compilation.
auto_migrate	<ul style="list-style-type: none"> platform_type—Choose this value based on the value referenced in the Intel FPGA Reference Platform from which you derive your Custom Platform. Valid values are a10_ref, s10_ref, and none. include_fixes—Comma-separated list of named fixes that you want to apply to the Custom Platform. exclude_fixes—Comma-separated list of named fixes that you do not want to apply to the Custom Platform.

2.3. MMD API Descriptions

The MMD interface is an accumulation of all the MMD application programming interface (API) functions.

Important: Full details about these functions, their arguments, and their return values are available in the aocl_mmd.h file. Refer to the aocl_mmd.h file in one of the reference platform BSPs available at the [FPGA Software Download Center](#)). Include the file in the operating system-specific implementations of the MMD layer.

- [aocl_mmd_get_offline_info](#) on page 41
- [aocl_mmd_get_info](#) on page 43
- [aocl_mmd_open](#) on page 45
- [aocl_mmd_close](#) on page 45
- [aocl_mmd_read](#) on page 45
- [aocl_mmd_write](#) on page 46
- [aocl_mmd_copy](#) on page 47

- [aocl_mmd_set_interrupt_handler](#) on page 48
- [aocl_mmd_set_device_interrupt_handler](#) on page 49
- [aocl_mmd_set_status_handler](#) on page 50
- [aocl_mmd_yield](#) on page 51
- [aocl_mmd_shared_mem_alloc](#) on page 51
- [aocl_mmd_shared_mem_free](#) on page 52
- [aocl_mmd_program](#) on page 52
- [aocl_mmd_reprogram](#) on page 53
- [aocl_mmd_hostchannel_create](#) on page 53
- [aocl_mmd_hostchannel_destroy](#) on page 54
- [aocl_mmd_hostchannel_get_buffer](#) on page 54
- [aocl_mmd_hostchannel_ack_buffer](#) on page 55
- [aocl_mmd_host_alloc](#) on page 55
- [aocl_mmd_free](#) on page 56
- [aocl_mmd_device_alloc](#) on page 57
- [aocl_mmd_shared_alloc](#) on page 57
- [aocl_mmd_shared_migrate](#) on page 58

2.3.1. aocl_mmd_get_offline_info

The `aocl_mmd_get_offline_info` function obtains offline information about the board specified in the `requested_info_id` argument. This function is offline because it is device-independent and does not require a handle from the `aocl_mmd_open()` call.

Syntax

```
int aocl_mmd_get_offline_info( aocl_mmd_offline_info_t requested_info_id,
                             size_t param_value_size,
                             void* param_value,
                             size_t* param_size_ret )
```

Function Arguments

- `requested_info_id`—An enum value of type `aocl_mmd_offline_info_t` that indicates the offline device information returning to the caller.

Table 16. Possible Enum Values for the `requested_info_id` Argument

Name	Description	Type
AOCL_MMD_VERSION	Version of MMD layer	char*
AOCL_MMD_NUM_BOARDS	Number of candidate boards	int
AOCL_MMD_BOARD_NAMES	Names of available boards Attention: Separate each board name by a semicolon (;) delimiter.	char*
AOCL_MMD_VENDOR_NAME	Name of board vendor	char*
<i>continued...</i>		

Name	Description	Type
AOCL_MMD_VENDOR_ID	An integer board vendor ID	int
AOCL_MMD_USES_YIELD	<p>A value of 0 instructs the Intel FPGA SDK for OpenCL Pro Edition host runtime to suspend user's processes. The host runtime resumes these processes after it receives an event update (for example, an interrupt) from the MMD layer .</p> <p>A value of 1 instructs the SDK's host runtime to continuously call the <code>aocl_mmd_yield</code> function while it waits for events to complete.</p> <p>Caution: Setting <code>AOCL_MMD_USES_YIELD</code> to 1 might cause high CPU utilization if the <code>aocl_mmd_yield</code> function does not suspend the current thread.</p>	int
AOCL_MMD_MEM_TYPES_SUPPORTED	<p>A bit field listing all memory types that the Custom Platform supports. You may combine the following enum values in this bit field:</p> <p><code>AOCL_MMD_PHYSICAL_MEMORY</code>—Custom Platform includes IP to communicate directly with physical memory (for example, DDR and QDR).</p> <p><code>AOCL_MMD_SVM_COARSE_GRAIN_BUFFER</code>—Custom Platform supports both the caching of shared virtual memory (SVM) pointer data for OpenCL <code>cl_mem</code> objects and the requirement of explicit user function calls to synchronize the cache between the host processor and the FPGA.</p> <p><i>Note:</i> Currently, Intel does not support this level of SVM except for a subset of <code>AOCL_MMD_SVM_FINE_GRAIN_SYSTEM</code> support.</p> <p><code>AOCL_MMD_SVM_FINE_GRAIN_BUFFER</code>—Custom Platform supports caching of SVM pointer data for individual bytes. To synchronize the cache between the host processor and the FPGA, the Custom Platform requires information from the host runtime collected during pointer allocation. After an SVM pointer receives this additional data, the board interface synchronizes the cache between the host processor and the FPGA automatically.</p> <p><i>Note:</i> Currently, Intel does not support this level of SVM except for a subset of <code>AOCL_MMD_SVM_FINE_GRAIN_SYSTEM</code> support.</p> <p><code>AOCL_MMD_SVM_FINE_GRAIN_SYSTEM</code>—Custom Platform supports caching of SVM pointer data for individual bytes and it does not require additional information from the host runtime to synchronize the cache between the host processor and the FPGA. The board interface synchronizes the cache between the host processor and the FPGA automatically for all SVM pointers.</p> <p>Attention: Intel's support for this level of SVM is preliminary. Some features might not be fully supported.</p>	int (bit field)

- `param_value_size`—Size of the `param_value` field in bytes. This `size_t` value should match the size of the expected return type that the enum definition indicates.
For example, if `AOCL_MMD_NUM_BOARDS` returns a value of type `int`, set the `param_value_size` to `sizeof (int)`. You should see the same number of bytes returned in the `param_size_ret` argument.
- `param_value`—A `void*` pointer to the variable that receives the returned information.
- `param_size_ret`—A pointer argument of type `size_t*` that receives the number of bytes of returned data.

Return Value

A negative return value indicates an error.

2.3.2. `aocl_mmd_get_info`

The `aocl_mmd_get_info` function obtains information about the board specified in the `requested_info_id` argument.

Syntax

```
int aocl_mmd_get_info( int handle,
                    aocl_mmd_info_t requested_info_id,
                    size_t param_value_size,
                    void* param_value,
                    size_t* param_size_ret );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `requested_info_id`—An enum value of type `aocl_mmd_offline_info_t` that indicates the device information returning to the caller.

Table 17. Possible Enum Values for the `requested_info_id` Argument

Name	Description	Type
<code>AOCL_MMD_NUM_KERNEL_INTERFACES</code>	Number of kernel interfaces	<code>int</code>
<code>AOCL_MMD_KERNEL_INTERFACES</code>	Kernel interfaces	<code>int*</code>
<code>AOCL_MMD_PLL_INTERFACES</code>	Kernel clock handles	<code>int*</code>
<code>AOCL_MMD_MEMORY_INTERFACE</code>	Global memory handle	<code>int</code>
<code>AOCL_MMD_TEMPERATURE</code>	Temperature measurement	<code>float</code>
<code>AOCL_MMD_PCIE_INFO</code>	PCIe information	<code>char*</code>
<code>AOCL_MMD_BOARD_NAME</code>	Board name	<code>char*</code>
<code>AOCL_MMD_BOARD_UNIQUE_ID</code>	Unique board ID	<code>char*</code>
<code>AOCL_MMD_CONCURRENT_READS</code>	Number of parallel reads A value of 1 indicates serial reads.	<code>int</code>

continued...

Name	Description	Type
AOCL_MMD_CONCURRENT_WRITES	Number of parallel writes A value of 1 indicates serial writes.	int
AOCL_MMD_CONCURRENT_READS_OR_WRITES	Total number of concurrent read and write operations.	int
AOCL_MMD_MIN_HOST_MEMORY_ALIGNMENT	Minimum alignment that the BSP supports for host allocations.	size_t
AOCL_MMD_HOST_MEM_CAPABILITIES	Capabilities of <code>aocl_mmd_host_alloc()</code> function	unsigned int
AOCL_MMD_SHARED_MEM_CAPABILITIES	Capabilities of <code>aocl_mmd_shared_alloc()</code> function	unsigned int
AOCL_MMD_DEVICE_MEM_CAPABILITIES	Capabilities of <code>aocl_mmd_device_alloc()</code> function	unsigned int
AOCL_MMD_HOST_MEM_CONCURRENT_GRANULARITY	Granularity of concurrent host accesses	size_t
AOCL_MMD_SHARED_MEM_CONCURRENT_GRANULARITY	Granularity of concurrent shared accesses	size_t
AOCL_MMD_DEVICE_MEM_CONCURRENT_GRANULARITY	Granularity of concurrent device accesses	size_t

- param_value_size**—Size of the `param_value` field in bytes. This `size_t` value should match the size of the expected return type that the enum definition indicates.
 For example, if `AOCL_MMD_TEMPERATURE` returns a value of type `float`, set the `param_value_size` to `sizeof (float)`. You should see the same number of bytes returned in the `param_size_ret` argument.
- param_value**—A `void*` pointer to the variable that receives the returned information.
- param_size_ret**—A pointer argument of type `size_t*` that receives the number of bytes of returned data.

Capability Values

Table 18. Capability Values for `aocl_mmd_get_info` Function

Value	Description
AOCL_MMD_MEM_CAPABILITY_SUPPORTED	If you do not set this value, allocation function is not supported even if other capabilities are set.
AOCL_MMD_MEM_CAPABILITY_ATOMIC	Supports atomic access to the memory by either the host or the device.
AOCL_MMD_MEM_CAPABILITY_CONCURRENT	Supports concurrent access to the memory either by the host or the device if the accesses are not on the same block. Block granularity is defined by <code>AOCL_MMD_*_MEM_CONCURRENT_GRANULARITY</code> . Blocks are aligned to this granularity.
AOCL_MMD_MEM_CAPABILITY_P2P	Memory can be accessed by multiple devices at the same time.

Return Value

A negative return value indicates an error.

2.3.3. aocl_mmd_open

The `aocl_mmd_open` function opens and initializes the specified device.

Syntax

```
int aocl_mmd_open( const char* name );
```

Function Arguments

`name`—The function opens the board with a name that matches this `const char*` string. The name typically matches the one specified by the `AOCL_MMD_BOARD_NAMES` offline information.

The OpenCL runtime first queries the `AOCL_MMD_BOARD_NAMES` offline information to identify the boards that it might be able to open. Then it attempts to open all possible devices by calling `aocl_mmd_open` and using each of the board names as argument.

Important: The name must be a C-style NULL-terminated ASCII string.

Return Value

If `aocl_mmd_open()` executes successfully, the return value is a positive integer that acts as a handle to the board.

If `aocl_mmd_open()` fails to execute, a negative return value indicates an error. In the event of an error, the OpenCL runtime proceeds to open other known devices. Therefore, it is imperative that the MMD layer does not exit the application if an open call fails.

2.3.4. aocl_mmd_close

The `aocl_mmd_close` function closes an opened device via its handle.

Syntax

```
int aocl_mmd_close( int handle );
```

Function Arguments

`handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.

Return Value

If the `aocl_mmd_close()` executes successfully, the return value is 0.

If `aocl_mmd_close()` fails to execute, a negative return value indicates an error.

2.3.5. aocl_mmd_read

The `aocl_mmd_read` function is the read operation on a single interface.

Syntax

```
int aocl_mmd_read( int handle,
                  aocl_mmd_op_t op,
                  size_t len,
                  void* dst,
                  aocl_mmd_interface_t interface,
                  size_t offset );
```

Function Arguments

- **handle**—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- **op**—The operation object of type `aocl_mmd_op_t` used to track the progress of the operation. If `op` is `NULL`, the call must block, and return only after the operation completes.

Note: `aocl_mmd_op_t` is defined as follows:

```
typedef void* aocl_mmd_op_t;
```

- **len**—The size of the data, in bytes, that the function transfers. Declare `len` with type `size_t`.
- **dst**—The host buffer, of type `void*`, to which data is written.
- **interface**—The handle to the interface that `aocl_mmd_read` is accessing. For example, to access global memory, this handle is the enum value `aocl_mmd_get_info()` returns when its `requested_info_id` argument is `AOCL_MMD_MEMORY_INTERFACE`. The interface argument is of type `aocl_mmd_interface_t`, and can take one of the following values:

Name	Description
AOCL_MMD_KERNEL	Control interface into the kernel interface
AOCL_MMD_MEMORY	Data interface to device memory
AOCL_MMD_PLL	Interface for reconfigurable PLL

- **offset**—The `size_t` byte offset within the interface at which the data transfer begins.

Return Value

If the read operation is successful, the return value is 0.

If the read operation fails, a negative return value indicates an error.

2.3.6. aocl_mmd_write

The `aocl_mmd_write` function is the write operation on a single interface.

Syntax

```
int aocl_mmd_write( int handle,
                   aocl_mmd_op_t op,
                   size_t len,
```

```
const void* src,  
aocl_mmd_interface_t interface,  
size_t offset );
```

Function Arguments

- **handle**—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- **op**—The operation object of type `aocl_mmd_op_t` used to track the progress of the operation. If `op` is `NULL`, the call must block, and return only after the operation completes.

Note: `aocl_mmd_op_t` is defined as follows:

```
typedef void* aocl_mmd_op_t;
```

- **len**—The size of the data, in bytes, that the function transfers. Declare `len` with type `size_t`.
- **src**—The host buffer, of type `const void*`, from which data is read.
- **interface**—The handle to the interface that `aocl_mmd_write` is accessing. For example, to access global memory, this handle is the enum value `aocl_mmd_get_info()` returns when its `requested_info_id` argument is `AOCL_MMD_MEMORY_INTERFACE`. The interface argument is of type `aocl_mmd_interface_t`, and can take one of the following values:

Name	Description
<code>AOCL_MMD_KERNEL</code>	Control interface into the kernel interface
<code>AOCL_MMD_MEMORY</code>	Data interface to device memory
<code>AOCL_MMD_PLL</code>	Interface for reconfigurable PLL

- **offset**—The `size_t` byte offset within the interface at which the data transfer begins.

Return Value

If the read operation is successful, the return value is 0.

If the read operation fails, a negative return value indicates an error.

2.3.7. `aocl_mmd_copy`

The `aocl_mmd_copy` function is the copy operation on a single interface.

Syntax

```
int aocl_mmd_copy( int handle,  
aocl_mmd_op_t op,  
size_t len,  
aocl_mmd_interface_t intf,  
size_t src_offset,  
size_t dst_offset );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `op`—The operation object of type `aocl_mmd_op_t` used to track the progress of the operation. If `op` is `NULL`, the call must block, and return only after the operation completes.

Note: `aocl_mmd_op_t` is defined as follows:

```
typedef void* aocl_mmd_op_t;
```

- `len`—The size of the data, in bytes, that the function transfers. Declare `len` with type `size_t`.
- `intf`—The handle to the interface that `aocl_mmd_read` is accessing. For example, to access global memory, this handle is the enum value `aocl_mmd_get_info()` returns when its `requested_info_id` argument is `AOCL_MMD_MEMORY_INTERFACE`. The interface argument is of type `aocl_mmd_interface_t`, and can take one of the following values:

Name	Description
<code>AOCL_MMD_KERNEL</code>	Control interface into the kernel interface
<code>AOCL_MMD_MEMORY</code>	Data interface to device memory
<code>AOCL_MMD_PLL</code>	Interface for reconfigurable PLL

- `src_offset`—The `size_t` byte offset within the source interface at which the data transfer begins.
- `dst_offset`—The `size_t` byte offset within the destination interface at which the data transfer begins.

Return Value

If the copy operation is successful, the return value is 0.

If the copy operation fails, a negative return value indicates an error.

2.3.8. `aocl_mmd_set_interrupt_handler`

The `aocl_mmd_set_interrupt_handler` function sets the interrupt handler for the opened device. When the device internals identify an asynchronous kernel event (for example, a kernel completion), the interrupt handler is called to notify the OpenCL runtime of the event.

Attention: Ignore the interrupts from the kernel until this handler is set.

Syntax

```
int aocl_mmd_set_interrupt_handler( int handle,
                                  aocl_mmd_interrupt_handler_fn fn,
                                  void* user_data );
```


Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `fn`—The callback function to invoke when a kernel interrupt occurs. The `fn` argument is of type `aocl_mmd_interrupt_handler_fn`, which is defined as follows:

```
typedef void (*aocl_mmd_interrupt_handler_fn)( int handle, void* user_data );
```
- `user_data`—The `void*` type user-provided data that passes to `fn` when it is called.

Return Value

If the function executes successfully, the return value is 0.

If the function fails to execute, a negative return value indicates an error.

2.3.9. `aocl_mmd_set_device_interrupt_handler`

The `aocl_mmd_set_device_interrupt_handler` function sets the device interrupt handler for the opened device. When the device internals identify an asynchronous exception event (for example, a bit correction event), the device interrupt handler is called to notify the OpenCL runtime of the event.

Attention: Ignore the interrupts from the device until this handler is set.

Syntax

```
int aocl_mmd_set_device_interrupt_handler( int handle,  
                                          aocl_mmd_device_interrupt_handler_fn fn,  
                                          void* user_data );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `fn`—The callback function to invoke when a kernel interrupt occurs. The `fn` argument is of type `aocl_mmd_device_interrupt_handler_fn`, which is defined as follows:

```
typedef void (*aocl_mmd_device_interrupt_handler_fn)( int handle,  
aocl_mmd_interrupt_info* data_in, void* user_data );
```

`aocl_mmd_interrupt_info` is defined as:

```
typedef struct {  
    unsigned long long int exception_type;  
    void *user_private_info;  
    size_t user_cb;  
} aocl_mmd_interrupt_info;
```

Where:

- `exception_type` acts as a bitfield that contains exactly one bit, corresponding to an exception number.
 - `user_private_info` and `user_cb` represent pointers to binary data that the OpenCL implementation return. These pointers log additional information that is helpful for debugging the error.
- `user_data`—The `void*` type user-provided data that passes to `fn` when it is called.

Return Value

If the function executes successfully, the return value is 0.

If the function fails to execute, a negative return value indicates an error.

2.3.10. `aocl_mmd_set_status_handler`

The `aocl_mmd_set_status_handler` function sets the operation status handler for the opened device. The operation status handler is called under the following circumstances:

- When the operation completes successfully and status is 0.
- When the operation completes with errors and status is a negative value.

Syntax

```
int aocl_mmd_set_status_handler( int handle,  
                                aocl_mmd_status_handler_fn fn,  
                                void* user_data );
```

Function Arguments

- **handle**—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- **fn**—The callback function to invoke when a status update occurs. The `fn` argument is of type `aocl_mmd_status_handler_fn`, which is defined as follows:

```
type void (*aocl_mmd_status_handler_fn)( int handle, void* user_data,  
aocl_mmd_op_t op, int status );
```

- **user_data**—The `void*` type user-provided data that passes to `fn` when it is called.

Return Value

If the function executes successfully, the return value is 0.

If the function fails to execute, a negative return value indicates an error.

2.3.11. `aocl_mmd_yield`

The `aocl_mmd_yield` function is called when the host interface is idle. The host interface might be idle because it is waiting for the device to process certain events.

Syntax

```
int aocl_mmd_yield( int handle );
```

Function Arguments

handle—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.

Return Value

A nonzero return value indicates that the yield function performed work necessary for proper device functioning such as processing direct memory access (DMA) transactions.

A return value of 0 indicates that the yield function did not perform work necessary for proper device functioning.

Note: The yield function might be called continuously if it reports that it has necessary work to perform.

2.3.12. `aocl_mmd_shared_mem_alloc`

Starting with the Intel FPGA SDK for OpenCL version 20.3, the `aocl_mmd_shared_mem_alloc` function is deprecated and provided only for backward compatibility. Use the `aocl_mmd_host_alloc` function instead.

Related Information

[aocl_mmd_host_alloc](#) on page 55

2.3.13. aocl_mmd_shared_mem_free

Starting with the Intel FPGA SDK for OpenCL version 20.3, the `aocl_mmd_shared_mem_free` function is deprecated and provided only for backward compatibility. Use the `aocl_mmd_free` function instead.

Related Information

[aocl_mmd_free](#) on page 56

2.3.14. aocl_mmd_program

The `aocl_mmd_program` function is the program operation for the specified device. The host must guarantee that no other OpenCL operations are executing on the device during the program operation. During `aocl_mmd_program` execution, the kernels are idle and no read, write, or copy operation can occur.

Disable interrupts and program the FPGA with the data from `user_data`, which has a size specified by the `size` argument. The host then calls `aocl_mmd_set_status_handler` and `aocl_mmd_set_interrupt_handler` again, which enable the interrupts. If events such as interrupts occur during `aocl_mmd_program` execution, race conditions or data corruption might occur.

This function replaces `aocl_mmd_program`, which is deprecated starting with Intel FPGA SDK for OpenCL Version 18.1.

Syntax

```
int aocl_mmd_program( int handle,
                    void* user_data,
                    size_t size,
                    aocl_mmd_program_mode_t program_mode );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `user_data`—The `void*` type binary contents of the `fpga.bin` file that are created during compilation.
- `size`—The size of `user_data` in bytes. The size argument is of `size_t`.
- `program_mode`—The bit-field that specifies the mode of device programming.

Table 19. Possible Values for the `program_mode` Argument

<code>program_mode</code> Argument Value	Description
<code>AOCL_MMD_PROGRAM_PRESERVE_GLOBAL_MEMORY</code>	This flag specifies that during programming the global memory on the devices are preserved.

Return Value

If `aocl_mmd_program` executes successfully, the return value is the pointer value that the host uses to access shared memory.

2.3.14.1. Program Support

For Intel FPGA SDK for OpenCL Pro Edition users who program their FPGAs with the `clCreateProgramWithBinary` flow (that is, program-on-the-fly), the `aocl_mmd_program` subroutine is used to configure the FPGA from within the host applications. The host ensures that this call executes only when the FPGA is idle, meaning that no kernels are running and no transfers are outstanding. The MMD layer must then reconfigure the device with the data in the `user_data` argument of `aocl_mmd_program`.

The data in the `user_data` argument is the same `fpga.bin` data created during Intel Quartus Prime compilation. The Intel FPGA SDK for OpenCL Offline Compiler packages the exact contents of `fpga.bin` into the `.aocx` file during compilation. The contents of the `fpga.bin` is irrelevant to the offline compiler. It passes the file contents through the host and to the `aocl_mmd_program` call via the `user_data` argument.

For more information about the `clCreateProgramWithBinary` function, refer to the *OpenCL Specification version 1.0* and the *Programming an FPGA via the Host* section of the *Intel FPGA SDK for OpenCL Pro Edition Programming Guide*.

Related Information

- [OpenCL Specification version 1.0](#)
- [Programming an FPGA via the Host](#)

2.3.15. aocl_mmd_reprogram

Starting with Intel FPGA SDK for OpenCL Version 18. 1, the `aocl_mmd_reprogram` function is deprecated and provided only for backwards compatibility. Use `aocl_mmd_program` instead.

2.3.16. aocl_mmd_hostchannel_create

The `aocl_mmd_hostchannel_create` function creates a channel interface.

Syntax

```
int aocl_mmd_hostchannel_create( int handle,
                                char *channel_name,
                                size_t queue_depth,
                                int direction );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `channel_name`—Name of the channel to be initialized. The channel name is same as that used in the `board_spec.xml` file.
- `queue_depth`—The size of pinned internal buffer in bytes. Pointer to the internal buffer is provided when the user calls the `aocl_mmd_hostchannel_get_buffer()` function.
- `direction`—The direction of the channel.

Return Value

If the function executes successfully, the return value is positive and is handle to the channel.

If the function fails to execute, a negative return value indicates an error.

2.3.17. aocl_mmd_hostchannel_destroy

The `aocl_mmd_hostchannel_destroy` function destroys the channel interface.

Syntax

```
int aocl_mmd_hostchannel_destroy( int handle,  
                                int channel );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `channel`—A positive `int` value representing handle to the channel to close obtained from the `aocl_mmd_hostchannel_create()` call.

Return Value

If the function executes successfully, the return value is 0.

If the function fails to execute, a negative return value indicates an error.

2.3.18. aocl_mmd_hostchannel_get_buffer

The `aocl_mmd_hostchannel_get_buffer` function provides a host with a pointer to the buffer they can access to write or read from the channel interface, along with the space or data available in the buffer, in bytes.

Syntax

```
void *aocl_mmd_hostchannel_get_buffer( int handle,  
                                      int channel,  
                                      size_t *buffer_size,  
                                      int *status );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `channel`—A positive `int` value representing handle to the channel to close obtained from the `aocl_mmd_hostchannel_create()` call.
- `buffer_size`—A pointer to `size_t` that the function writes available buffer space or size to.
- `status`—A pointer to `int` that the function writes result of the call to.

Return Value

If the function executes successfully, `int` pointed to by the `status` pointer is 0. Returned `void*` may still be NULL, in which case `size_t` pointed by the `buffer_size` is 0.

If the function fails to execute, `int` pointed by the `status` pointer is a negative value.

2.3.19. `aocl_mmd_hostchannel_ack_buffer`

You can acknowledge write or read from the channel by calling `aocl_mmd_hostchannel_ack_buffer`.

Syntax

```
size_t aocl_mmd_hostchannel_ack_buffer( int handle,
                                       int channel,
                                       size_t send_size,
                                       int *status );
```

Function Arguments

- `handle`—A positive `int` value representing the handle to the board obtained from the `aocl_mmd_open()` call.
- `channel`—A positive `int` value representing handle to the channel to close obtained from the `aocl_mmd_hostchannel_create()` call.
- `send_size`—The size in bytes that the user is acknowledging.
- `status`—A pointer to `int` that the function writes result of the call to.

Return Value

If the function executes successfully, `int` pointed to by `status` pointer is 0. Also, there is no guarantee that the user's `send_size` is the actual size that gets acknowledged. The returned `size_t` is the amount of bytes that was actually acknowledged.

If the function fails to execute, `int` pointed by `status` pointer is a negative value.

2.3.20. `aocl_mmd_host_alloc`

Host allocations provide memory that is allocated on the host. This memory must be deallocated with the `aocl_mmd_free` function. Host allocations are accessible by the host and one or more devices. The same pointer to a host allocation may be used on the host and all supported devices. They have address equivalence.

Once the device has signaled completion through the `aocl_mmd_interrupt_handler_fn` function, the host can assume it has access to the latest contents of the memory, allocated by the `aocl_mmd_host_alloc` function call.

Syntax

```
void* aocl_mmd_host_alloc (int* handles,
                          size_t num_devices,
                          size_t size,
                          size_t alignment,
                          aocl_mmd_mem_properties_t *properties,
                          int* error);
```

Function Arguments

- `handles`—Handles for devices that needs access to this memory.
- `num_devices`—Number of devices in the handles.
- `size`—The size of the memory region.
- `alignment`—The alignment (in bytes) of the allocation.
- `properties`—Specifies additional information about the allocated memory, described by a property type name and its corresponding value. Each property type name is immediately followed by the corresponding desired value. The list is terminated with a zero. Supported values are described above. For example, [`<property1>`, `<value1>`, `<property2>`, `<value2>`, 0]
- `error`—The error code defined by `AOCL_MMD_ERROR*`:
 - `AOCL_MMD_ERROR_SUCCESS`—No error occurred.
 - `AOCL_MMD_ERROR_INVALID_HANDLE`—The device handle provided is invalid.
 - `AOCL_MMD_ERROR_OUT_OF_MEMORY`—Ran out of memory.
 - `AOCL_MMD_ERROR_UNSUPPORTED_ALIGNMENT`—The device does not support the provided alignment.
 - `AOCL_MMD_ERROR_UNSUPPORTED_PROPERTY`—The device does not support the provided property.

Return Value

If the `aocl_mmd_host_alloc` function executes successfully, the return value is a valid pointer value. Otherwise, the return value is `NULL`.

2.3.21. `aocl_mmd_free`

Releases memory that was allocated by MMD.

Syntax

```
int aocl_mmd_free (void* mem);
```

Function Arguments

`mem`—The pointer to the memory region. Must be a pointer that is allocated by the MMD.

Return Values

Returns one of the following error code:

- `AOCL_MMD_ERROR_SUCCESS`—No error occurred
- `AOCL_MMD_ERROR_INVALID_POINTER`—Invalid pointer provided

2.3.22. `aocl_mmd_device_alloc`

Allocate memory that is owned by the device. This pointer can only be accessed by the kernel. It cannot be accessed by the host. The host is able to manipulate the pointer (for example, increment it) and not just access the underlying data. This memory must be deallocated by the `aocl_mmd_free()` function.

Syntax

```
void * aocl_mmd_device_alloc (int handle,  
                             size_t size,  
                             size_t alignment,  
                             aocl_mmd_mem_properties_t *properties,  
                             int* error);
```

Function Arguments

- `handle`—Device that has access to this memory.
- `size`—The size of the memory region.
- `alignment`—The alignment (in bytes) of the memory region.
- `properties`—Specifies additional information about the allocated memory, described by a property type name and its corresponding value. Each property type name is immediately followed by the corresponding desired value. The list is terminated with a zero. Supported values are described above. For example, [`<property1>`, `<value1>`, `<property2>`, `<value2>`, 0]

Return Values

Returns one of the following error code:

- `AOCL_MMD_ERROR_SUCCESS`—No error occurred
- `AOCL_MMD_ERROR_INVALID_HANDLE`—The device handle provided is invalid.
- `AOCL_MMD_ERROR_OUT_OF_MEMORY`—Ran out of memory.
- `AOCL_MMD_ERROR_UNSUPPORTED_ALIGNMENT`—The device does not support the provided alignment.
- `AOCL_MMD_ERROR_UNSUPPORTED_PROPERTY`—The device does not support the provided property.

2.3.23. `aocl_mmd_shared_alloc`

Shared allocations can migrate between the host and one or more associated device. The same pointer to a shared allocation can be used on the host and the supported device. They have address equivalence.

Syntax

```
void * aocl_mmd_shared_alloc (int handle,
                             size_t size,
                             size_t alignment,
                             aocl_mmd_mem_properties_t* properties,
                             int* error);
```

Function Arguments

- `handle`—Device that needs access to this memory.
- `size`—The size of the memory region.
- `alignment`—The alignment (in bytes) of the allocation.
- `properties`—Specifies additional information about the allocated memory described by a property type name and its corresponding value. Each property type name is immediately followed by the corresponding desired value. The list is terminated with a zero. Supported values are described above. For example, [`<property1>`, `<value1>`, `<property2>`, `<value2>`, 0]
- `error`—The error code defined by `AOCL_MMD_ERROR*`.
 - `AOCL_MMD_ERROR_SUCCESS`—No error occurred.
 - `AOCL_MMD_ERROR_INVALID_HANDLE`—The device handle provided is invalid.
 - `AOCL_MMD_ERROR_OUT_OF_MEMORY`—Ran out of memory.
 - `AOCL_MMD_ERROR_UNSUPPORTED_ALIGNMENT`—The device does not support the provided alignment.
 - `AOCL_MMD_ERROR_UNSUPPORTED_PROPERTY`—The device does not support the provided property.

Return Values

If the `aocl_mmd_shared_alloc` function executes successfully, the return value is a valid pointer value. Otherwise, the return value is `NULL`.

2.3.24. aocl_mmd_shared_migrate

A call to the `aocl_mmd_shared_migrate()` function must be made for non-concurrent shared allocations any time the accessor of the allocation changes. For example, the `aocl_mmd_shared_migrate()` function should be called indicating that the allocation should be migrated to the device before a kernel accessing the allocation is launched on the device. Similarly, the `aocl_mmd_shared_migrate()` function should be called indicating that the allocation is migrated to the host before the host accesses the memory after kernel completion. For concurrent allocations, this call may be used as a performance hint, but it is not strictly required for functionality.

Syntax

```
int aocl_mmd_shared_migrate (int handle,
                             void* shared_ptr,
                             size_t size,
                             aocl_mmd_migrate_t destination);
```

Function Arguments

- `handle`—Device that has access to this memory.
- `hared_ptr`—Pointer allocated by the `aocl_mmd_shared_alloc()` function.
- `size`—Size (in bytes) of the migration. Must be a multiple of a page boundary that the BSP supports.
- `destination`—The destination of the migration.

Return Values

Returns one of the following error code:

- `AOCL_MMD_ERROR_SUCCESS`—No error occurred.
- `AOCL_MMD_ERROR_INVALID_HANDLE`—The device handle provided is invalid.
- `AOCL_MMD_ERROR_INVALID_MIGRATION_SIZE`—The migration size is not supported by the device.

3. Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit Archives

If the table does not list a software version, the user guide for the previous software version applies.

Intel Quartus Prime Version	User Guide
20.1	Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit User Guide
18.1	Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit User Guide
18.0	Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit User Guide
17.1	Intel FPGA SDK for OpenCL Custom Platform Toolkit User Guide
17.0	Intel FPGA SDK for OpenCL Custom Platform Toolkit User Guide
16.1	Intel FPGA SDK for OpenCL Custom Platform Toolkit User Guide
16.0	Altera SDK for OpenCL Custom Platform Toolkit User Guide
15.1	Altera SDK for OpenCL Custom Platform Toolkit User Guide
15.1	Altera SDK for OpenCL Custom Platform Toolkit User Guide
15.1	Altera SDK for OpenCL Custom Platform Toolkit User Guide
15.0	Altera SDK for OpenCL Custom Platform Toolkit User Guide
14.1	Altera SDK for OpenCL Custom Platform Toolkit User Guide

4. Document Revision History for Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit User Guide

Document Version	Intel Quartus Prime Version	Changes
2022.03.28	20.3	<ul style="list-style-type: none"> Updated the terms master and slave to host and agent, respectively.
2020.09.28	20.3	<ul style="list-style-type: none"> Added the <code>allocation_type</code> attribute in <i>global_mem</i> topic. Added more enum values and capability values in <i>aocl_mmd_get_info</i> topic Added the following APIs: <ul style="list-style-type: none"> <code>aocl_mmd_host_alloc</code> <code>aocl_mmd_free</code> <code>aocl_mmd_device_alloc</code> <code>aocl_mmd_shared_alloc</code> <code>aocl_mmd_shared_migrate</code> Deprecated <code>aocl_mmd_shared_mem_alloc</code> and <code>aocl_mmd_shared_mem_free</code> APIs. Updated the directories list in <i>Directories and Files in an Intel FPGA SDK for OpenCL Pro Edition Custom Platform</i> Updated the Example System Hierarchy diagram to remove the reference to <code>system.qsys</code> in <i>Designing the Board Hardware</i>. Removed all occurrences of <code>system.qsys</code> and replaced it with <code>board.qsys</code>. Removed all references to <code>INTELFPGAOCLESDKROOT/board/custom_platform_toolkit/</code> path and replaced some of them with a reference to the BSP download location. Removed a note about <code>alterahalmmd</code> library for <code>linklibs</code> element in <i>Creating the board_env.xml File</i>. Removed <code>alterahalmmd.lib</code> from the SDK output in <i>Testing the board_env.xml File</i>. Changed <code>pll_refclk</code> to <code>kernel_pll_refclk</code> in <i>OpenCL Kernel Clock Generator</i>. Changed <code>kernel_ctrl</code> to <code>ctrl</code> in <i>OpenCL Kernel Interface</i>. Added a brief note for <code>acl_bsp_memorg_host</code> in <i>OpenCL Memory Bank Divider</i>. In <i>Creating the MMD Library and MMD API Descriptions</i>, changed the reference details for <code>aocl_mmd.h</code> file. Removed the step to set <code>AOCL_BOARD_TOOLKIT_ROOT</code> environment variable in <i>Testing the board_env.xml File</i>.
2020.04.13	20.1	<ul style="list-style-type: none"> Removed a note about <code>interleaved_bytes</code> in <i>Creating the board_spec.xml File</i> topic. Updated the description of Burst size (maximum) parameter in <i>OpenCL Memory Bank Divider</i> topic. Added recommended values for the Maximum Pending Reads parameter in <i>OpenCL Memory Bank Divider</i> topic. Updated the description of <code>interleaved_bytes</code> in <i>global_mem</i> topic. Updated the description of <code>maxburst</code> in <i>interface</i> topic.
<i>continued...</i>		

Document Version	Intel Quartus Prime Version	Changes
2018.10.26	18.1	<ul style="list-style-type: none"> Updated Creating the board_spec.xml File on page 17 and global_mem on page 35 to include requirement that default global memory must start at address 0x0. Merged <i>Intel FPGA SDK for OpenCL Pro Edition Custom Platform Design Revision History</i> and <i>Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit Reference Manual Revision History</i>.
2018.07.16	18.0	<ul style="list-style-type: none"> Corrected the Windows registry key in Setting Up the FPGA Client Driver on page 23. The correct key is HKEY_LOCAL_MACHINE\SOFTWARE\Intel\OpenCL\Boards.
2018.05.04	18.0	<ul style="list-style-type: none"> Removed information pertaining to the Intel FPGA SDK for OpenCL Standard Edition and the Intel Quartus Prime Standard Edition software. In Prerequisites for the Intel FPGA SDK for OpenCL Pro Edition Custom Platform Toolkit on page 5, updated the list of related documents to refer to the Intel Arria 10 FPGA and SoC Reference Platforms. In Creating the board_spec.xml File on page 17, added a note to specify that when creating a device model (i.e. <code><part_number>_dm.xml</code>), <code><part_number></code> must be the official Intel FPGA device part number. In interface on page 38, added information on the <code>waitrequest_allowance</code> parameter. In compile on page 39: <ul style="list-style-type: none"> Added more information on compiler behavior when <code>qsys_file</code> is set to none. Added <code>s10_ref</code> and none to the list of valid values for <code>auto-migrate</code>. Under MMD API Descriptions on page 40, added a topic for the new <code>aocl_mmd_set_device_interrupt_handler</code> API.

Date	Version	Changes
November 2017	2017.11.03	<ul style="list-style-type: none"> Rebranded the references to the following: <ul style="list-style-type: none"> <code>ALTERAOCLSDKROOT</code> to <code>INTELFPGAOCCLSDKROOT</code> Arria 10 to Intel Arria 10 Qsys to Platform Designer (Standard) In Setting Up the FPGA Client Driver on page 23, updated the <code>/opt/Altera/OpenCL/Boards/</code> directory to <code>/opt/Intel/OpenCL/Boards/</code> Updated instances of <code>aocl install</code> to <code>aocl install <path_to_customplatform></code>. Updated instances of <code>aocl uninstall</code> to <code>aocl uninstall <path_to_customplatform></code>. Removed references to <code>AOCL_BOARD_PACKAGE_ROOT</code> throughout the guide since it is deprecated. In Creating the board_env.xml File on page 15, updated the <code>board_env</code> version to 17.1 in the <code>board_env.xml</code> file. In Testing the board_env.xml File on page 17, updated <code>board_default</code> and <code>board-version</code> values.
		<i>continued...</i>

Date	Version	Changes
		<ul style="list-style-type: none"> In Creating the board_spec.xml File on page 17, updated the board_spec.xml file code example. Implemented the conventions single dash and <code>-option=<value></code> in the following topics: <ul style="list-style-type: none"> Customizing Automigration on page 8 Testing the board_env.xml File on page 17 Testing the Hardware Design on page 26 Applying for the Intel FPGA SDK for OpenCL Pro Edition Preferred Board Status on page 27 Added the following new pages: <ul style="list-style-type: none"> aocl_mmd_hostchannel_create on page 53 aocl_mmd_hostchannel_destroy on page 54 aocl_mmd_hostchannel_get_buffer on page 54 aocl_mmd_hostchannel_ack_buffer on page 55
May 2017	2017.05.08	<ul style="list-style-type: none"> Rebranded Altera Client Driver (ACD) to FPGA Client Driver (FCD). Updated example code of the board_env.xml and board_spec.xml files to the 17.0 version of corresponding files in the Arria 10 GX FPGA Development Kit Reference Platform.
October 2016	2016.10.31	<ul style="list-style-type: none"> Rebranded Altera SDK for OpenCL to Intel FPGA SDK for OpenCL. Rebranded Altera Offline Compiler to Intel FPGA SDK for OpenCL Offline Compiler. Updated example code of the board_env.xml and board_spec.xml files to the 16.1 version of corresponding files in the Arria 10 GX FPGA Development Kit Reference Platform. In Setting Up the Altera Client Driver, updated the text for clarity. In Testing the Hardware Design, modified the Windows make command in Step 5 to make <code>-f Makefile.windows</code>. In OpenCL Memory Bank Divider, added information about the Split read/write bursts on burst word boundary parameter.
May 2016	2016.05.02	<ul style="list-style-type: none"> In Creating the board_spec.xml File, updated the example XML code for board_spec.xml to the current version, and updated the examples embedded in the procedure to match the example .xml file. Updated implementation requirement for the program utility in the Providing AOCL Utilities Support section. In Setting Up the Altera Client Driver, modified the Linux directory for the .acd file from <code>/opt/Altera/OpenCL_boards/</code> to <code>/opt/Altera/OpenCL/Boards/</code>. In the global_mem section under XML Elements, Attributes, and Parameters in the board_spec.xml File, added example calculations for determining the max_bandwidth value. In the interface section under XML Elements, Attributes, and Parameters in the board_spec.xml File, modified the global_mem-specific definitions for the name and port attributes. Added the option to assign a value of none to the qsys_file attribute within the compile element. Fixed a documentation error in the aocl_mmd_copy section.
November 2015	2015.11.02	<ul style="list-style-type: none"> Changed instances of <i>Quartus II</i> to <i>Quartus Prime</i>. Changed instances of <i>Altera® Complete Design Suite</i> to <i>Intel Quartus Prime Design Suite</i>. Updated the support requirement for the diagnose utility in the Providing AOCL Utilities Support section. In the Creating the board_env.xml File section, added the mmdlib XML element to the list of elements included in the board_env.xml file.
May 2015	15.0.0	<ul style="list-style-type: none"> Added the Setting Up the Altera Client Driver section.
continued...		

Date	Version	Changes
December 2014	14.1.0	<ul style="list-style-type: none"> • Specified that the Custom Platform Toolkit is available in the <code>ALTERAOCLSDKROOT/board</code> directory. • Added the <code>uninstall</code> utility executable in the <i>Providing AOCL Utilities Support</i> section. • Indicated that the version attributes in the <code>board_env.xml</code> and <code>board_spec.xml</code> files have to match the Altera Complete Design Suite and Altera SDK for OpenCL version you use to develop the Custom Platform. • Added instruction for including the <code>compile</code> eXtensible Markup Language element and its associated attributes in the <code>board_spec.xml</code> file in the section <i>Creating the board_spec.xml File</i>. • Added information on the automigration of Custom Platform in sections <i>Custom Platform Automigration</i> and <i>Customizing Automigration</i>. • Removed the <i>Generating the Rapid Prototyping Library</i> section. • Under <i>XML Elements, Attributes, and Parameters in the board_spec.xml File</i>, added information on the <code>compile</code> eXtensible Markup Language element and its associated attributes and parameters. • Under <i>MMD API Descriptions</i>, added information on the <code>AOCL_MMD_USES_YIELD</code> and the <code>AOCL_MMD_MEM_TYPES_SUPPORTED</code> enum values for the <code>requested_info_id</code> variable in the <code>aocl_mmd_get_offline_info</code> function.
October 2014	14.0.1	<ul style="list-style-type: none"> • Reorganized existing document into two chapters.
June 2014	14.0.0	<ul style="list-style-type: none"> • Initial release.