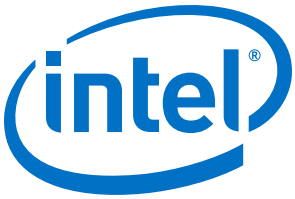

AN 881: PCI Express* Gen3 x16 Avalon®-MM DMA with On-Chip, External or HBM2 Memory Reference Designs

Subscribe
Send Feedback

AN-881 | 2018.12.27
Latest document on the web: PDF | HTML

Contents

1. Introduction

Table 1. Reference Designs Summary

Reference Design	Hardware	Throughput (GB/s)		Design Link
		Read	Write	
Avalon®-MM Intel® Stratix® 10 MX Hard IP+ DMA with Internal Memory	Intel Stratix 10 MX FPGA Development Kit	12.95	12.95	PCI Express Gen3 x16 AVMM DMA with On-Chip Memory Reference Design
Avalon-MM Intel Stratix 10 MX Hard IP + DMA with External DDR4 Memory	Intel Stratix 10 MX FPGA Development Kit	12.76	12.80	PCI Express Gen3 x16 AVMM DMA with DDR4 Memory Reference Design
Avalon-MM Intel Stratix 10 MX Hard IP + DMA with HBM2 Memory	Intel Stratix 10 MX FPGA Development Kit	13.22	13.01	PCI Express Gen3 x16 AVMM DMA with HBM2 Memory Reference Design

Note: To download these reference designs, first make sure that you have access to the Intel Design Store by logging into Design Store. You can then click on the links provided in the table above to download the designs.

Note: Although the theoretical maximum throughput for either Read or Write operations is 16 GB/s, the real throughput will be less than that (as shown in the table above) due to the overhead inherent in the PCI Express* protocol.

This document covers three reference designs using the PCI Express Avalon Memory-Mapped (Avalon-MM) Direct Memory Access (DMA) with Memory IP Interfaces. These reference designs demonstrate the performance of the Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express, a high-performance DMA controller and three types of memory solutions: on-chip, external (DDR4) and HBM2 memories.

Each of the reference designs includes a Linux* software driver to set up the DMA transfers with high-throughput data movers for DMA support. The Read Data Mover moves data from the system memory to the on-chip or external or HBM2 memory in Avalon-MM space . The Write Data Mover moves data from the on-chip or external or HBM2 memory in the application logic to the system memory in PCIe* space. These reference designs allow you to evaluate the performance of the Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express while using the Avalon-MM interface with high-performance DMA with different memory IPs.

Figure 1. Avalon-MM DMA Block Diagram

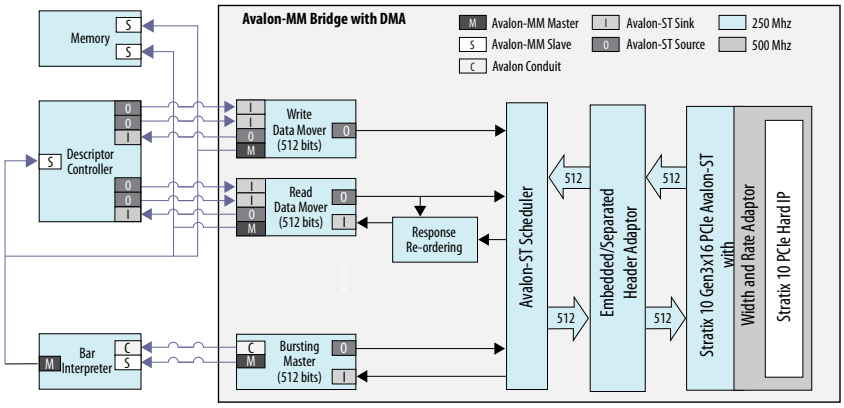

For Intel Arria® 10, Intel Cyclone® 10 GX or Intel Stratix 10 SX, GX or TX devices and Avalon-MM DMA configurations up to Gen3 x8, refer to:

- Intel Stratix 10 Avalon-MM Interface for PCI Express Solutions User Guide
- AN 829: PCI Express Avalon-MM DMA Reference Design
- Intel Arria 10 or Intel Cyclone 10 GX Avalon-MM DMA Interface for PCI Express Solutions User Guide

Related Information

- Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express Solutions User Guide
- PCI Express Base Specification Revision 3.0
- High Bandwidth Memory (HBM2) Interface Intel FPGA IP User Guide
- External Memory Interfaces Intel Stratix 10 FPGA IP User Guide
- Intel Stratix 10 MX FPGA Development Kit

1.1. PCI Express Gen3 x16 Avalon-MM DMA with On-Chip, External and HBM2 Memory Hardware and Software Requirements

Hardware Requirements

The three reference designs (for the three different types of memory - on-chip, DDR4 or HBM2) all run on the following development kits:

- Intel Stratix 10 FPGA MX development kit (using the on-chip or DDR4 or HBM2 memory). This development kit uses a 1SM21BHU2F53E2VGS1 Intel Stratix 10 device.

Each reference design requires two computers:

- A computer with a PCIe Gen3 x16 slot running Linux. This computer is computer number 1.
- A second computer with the Intel Quartus® Prime Pro Edition software version 18.1 installed. This computer downloads the FPGA SRAM Object File (.sof) to the FPGA on the development kit. This computer is computer number 2.

Software Requirements

- The reference design software is installed on computer number 1. The reference designs are available in the Intel FPGA Design Store. The Intel Quartus Prime Pro Edition Platform Archive File (.par) includes the recommended synthesis, fitter, and timing analysis settings for the parameters specified in the reference designs.
- The Intel Quartus Prime Pro Edition software is installed on computer number 2. You can download this software from the Intel Quartus Prime Pro Edition Software Features/Download web page.
- The Linux driver is configured specifically for these reference designs.

Note: The driver was developed and tested on CentOS 7.0, 64-bit with 3.10.514 kernel compiled for the x86_64 architecture.

Related Information

- Intel Stratix 10 PCI Express Gen3 x16 Avalon-MM DMA with On-Chip Memory Reference Design
- Intel Stratix 10 PCI Express Gen3 x16 Avalon-MM DMA with DDR4 Memory Reference Design
- Intel Stratix 10 PCI Express Gen3 x16 Avalon-MM DMA with HBM2 Memory Reference Design

To download the reference designs and the design software from the Design Store, go to:

- Intel Quartus Prime Pro Edition Download Center: http://fpgasoftware.intel.com/?edition=pro

1.2. Avalon-MM Bridge with DMA Module Descriptions

Refer to the *Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express Solutions User Guide*.

Related Information

Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express Solutions User Guide

1.3. DMA Procedure Steps

Software running on the host completes the following steps to initiate the DMA and verify the results.

Refer to *AN 829: PCI Express Avalon-MM DMA Reference Design* for the DMA procedure.

Related Information

https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/an/an829.pdf

1.4. Setting Up the Hardware

1. Power down computer number 1.
2. Plug the FPGA MX development kit card into a PCIe slot that supports Gen3 x16.
3. For the Intel Stratix 10 FPGA MX development kit, connector ATX AUX PWR (shown as J11 in the board schematic) powers the card. After inserting the card into an available PCIe slot, connect either a 2x4-pin or a 2x3-pin PCIe power cable from the power supply of computer number 1 to the ATX AUX PWR (J11) connector of the PCIe card.
4. Connect a USB cable from computer number 2 to the FPGA MX development kit. The development kit includes an on-board Intel FPGA Download Cable for FPGA programming.
5. To power up the FPGA MX development kit via the PCIe slot, power on computer number 1. You need to connect ATX AUX PWR to power up the board. External power is required when the development kit is connected to the PCIe slot. Use the external power adapter that ships with the kit.
6. On computer number 2, bring up the Intel Quartus Prime programmer and configure the FPGA through an Intel FPGA Download Cable.

   Note: You must reconfigure the FPGA whenever the FPGA development kit loses power.
7. To force system enumeration to discover the PCIe device, restart computer number 1.

1.5. Installing the DMA Test Driver and Running the Linux DMA Software

For instructions on how to install the DMA test driver and run the Linux DMA application, refer to sections 2.6 and 2.7 in the User Guide for the Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express (Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express Solutions User Guide).

Note: After downloading the reference design from the Design Store and decompressing the archived project (from the .qar file) using Intel Quartus Prime 18.1, you will find a few folders generated in your project directory. You will also have a `software.zip` folder with the software driver and tool. Make sure to unzip this folder in Windows. You may run into driver installation issues if you unzip this folder in Linux.

Figure 2. Link Test GUI

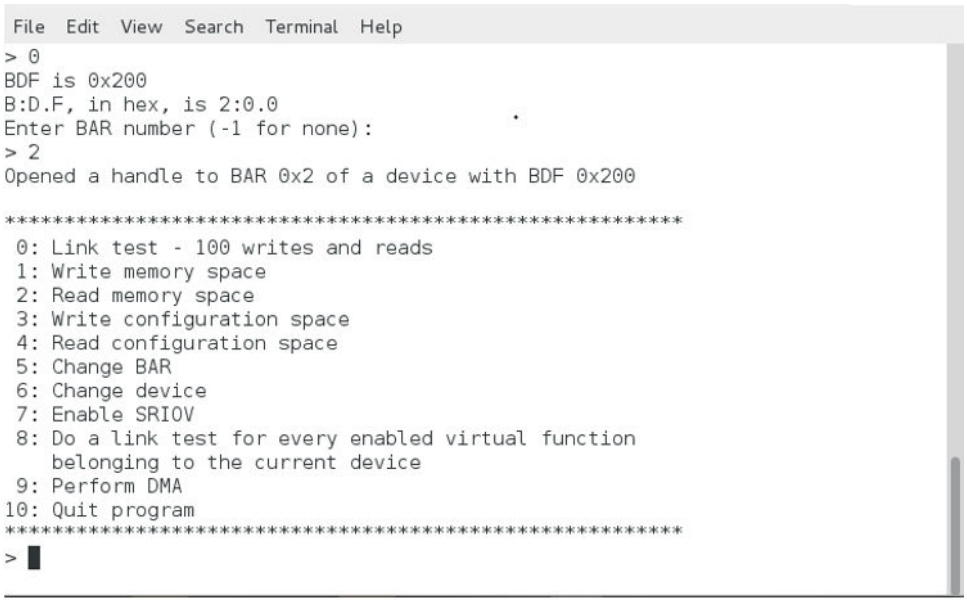

Figure 3. Link Test Pass Result

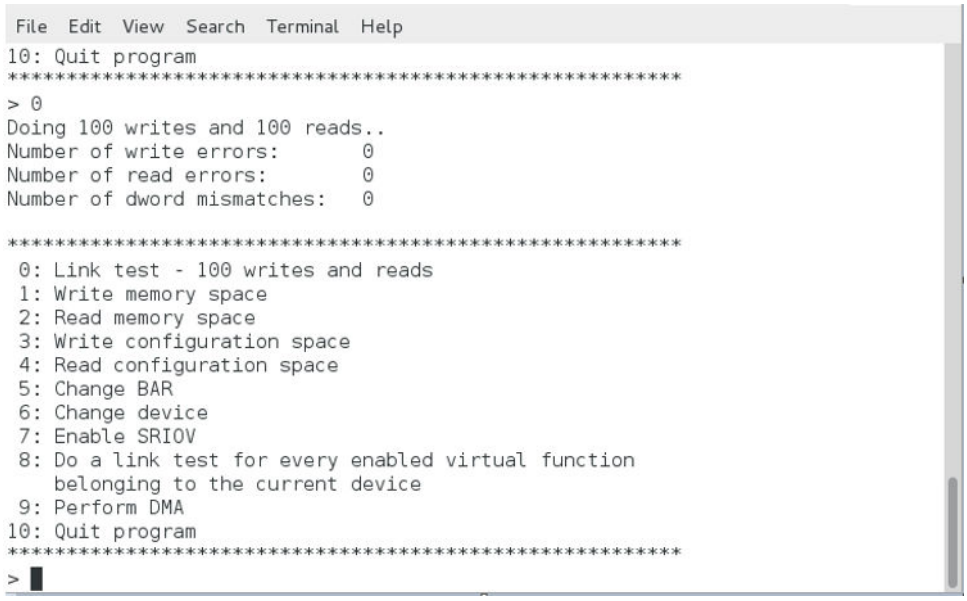

Figure 4. DMA GUI

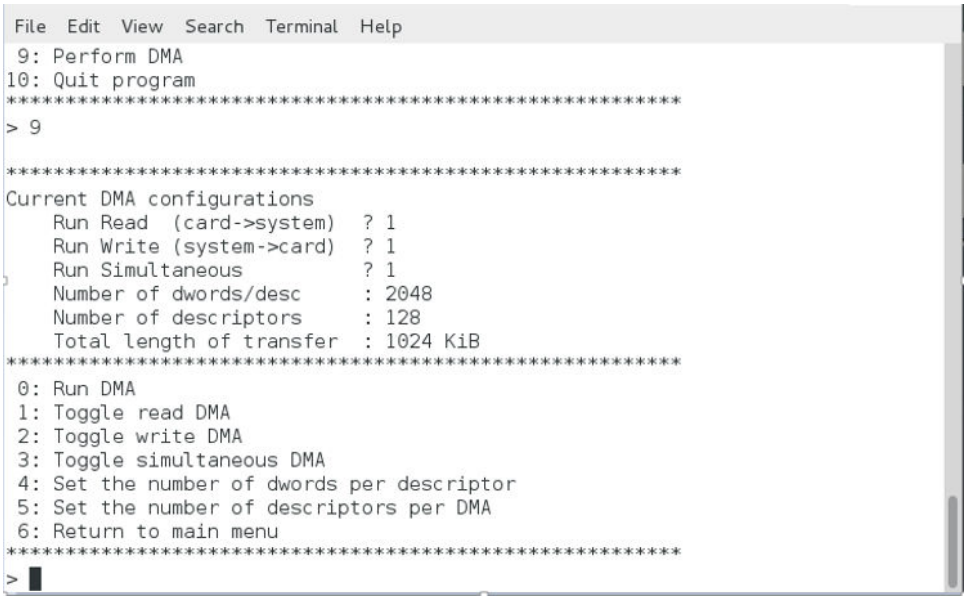

2. Reference Design Description

This application note covers three reference designs:

1. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with Internal Memory
2. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with External DDR4 Memory
3. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with HBM2 Memory

2.1. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with Internal Memory

Figure 5. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with Internal Memory Reference Design

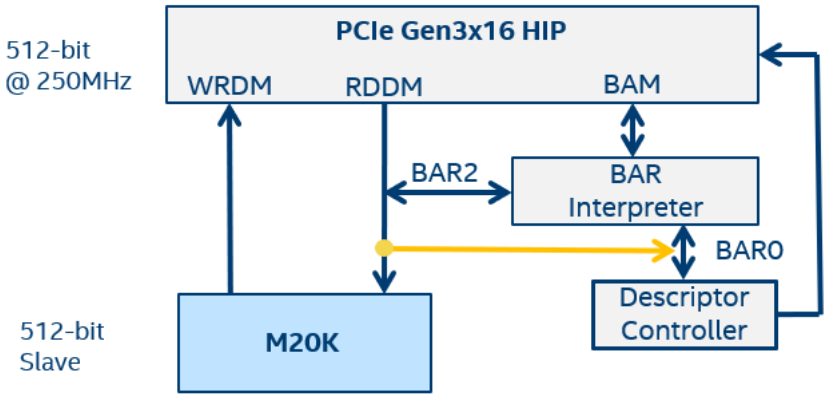

2.1.1. Project Hierarchy

The reference design uses the following directory structure:

Figure 6. Directory Structure

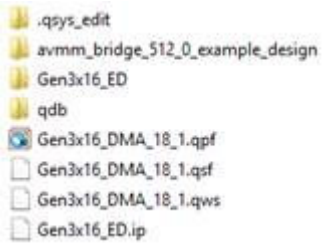

- pcie_example_design.v: The top-level module.

2.1.2. Parameter Settings for PCI Express Hard IP+ Variations

This reference design supports a 512-byte maximum payload size. The following tables list the values for all the parameters.

Table 2. System Settings

Parameter	Value
Number of lanes	Intel Stratix 10 MX: 16
Lane rate	Intel Stratix 10 Gen3: 16 Gbps
RX buffer credit allocation - performance for received request	Intel Stratix 10: Not available
Hard IP Mode	By default, the Hard IP mode is set to Gen3 x16, with a 512-bit interface to the Application Layer running at 250 MHz.

Table 3. Base Address Register (BAR) Settings

Parameter	Value	BAR Size
BAR0	64-bit prefetchable memory	16 bits
BAR1	Disabled	
BAR2	64-bit prefetchable memory	30 bits
BAR3	Disabled	
BAR4	Disabled	
BAR5	Disabled	

Table 4. Device Identification Register Settings

Parameter	Value
Vendor ID	0x00001172
Device ID	0x0000E003
Revision ID	0x00000001
Class Code	0x00000000
Subsystem Vendor ID	0x00000000
Subsystem Device ID	0x00000000

Table 5. PCI Express/PCI* Capabilities

Parameter	Value
Maximum payload size	512 bytes
Completion timeout range	None
Implement completion timeout	Disabled

Table 6. Error Reporting Settings

Parameter	Value
Advanced Error Reporting (AER)	Enabled
ECRC checking	Disabled
ECRC generation	Disabled

Table 7. Link Settings

Parameter	Value
Link port number	1
Slot clock configuration	Enabled

Table 8. Message Signaled Interrupts (MSI) and MSI-X Settings

Parameter	Value
Number of MSI messages requested	4
Implement MSI-X	Disabled
Table size	0
Table offset	0x0000000000000000
Table BAR indicator	0
Pending bit array (PBA) offset	0x0000000000000000
PBA BAR indicator	0

Table 9. Power Management

Parameter	Value
Endpoint L0s acceptable latency	Maximum of 64 ns
Endpoint L1 acceptable latency	Maximum of 1 us

Table 10. PCIe Address Space Setting

Parameter	Value
Address width of accessible PCIe memory space	40

2.1.3. PCIe Avalon-MM DMA Reference Design Platform Designer System

The following image show the modules in the Platform Designer system for this reference design.

Figure 7. Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express with On-Chip Memory Platform Designer Module

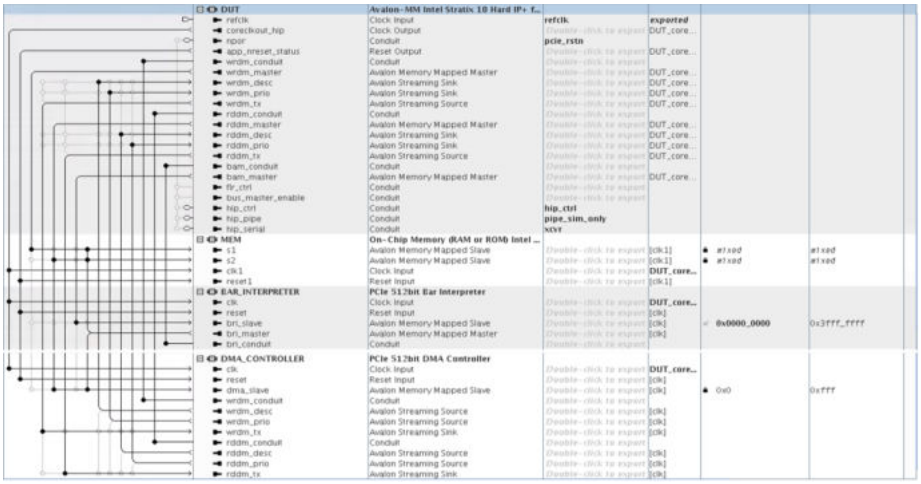

Table 11. Platform Designer Port Descriptions

Port	Function	Description
RDDM	Read Data Mover	This interface transfers DMA data from the PCIe system memory to the memory in Avalon-MM address space.
WRDM	Write Data Mover	This interface transfers DMA data from the memory in Avalon-MM address space to the PCIe system memory.
BAM	Bursting Avalon-MM Master	This interface provides host access to the registers and memory in Avalon-MM address space. The Bursting Avalon-MM Master module converts PCIe Memory Reads and Writes to Avalon-MM Reads and Writes.
On-chip Memory	64 KB Dual-Port RAM	This is a 64-KB dual-port on-chip memory. To prevent data corruption, software divides the memory into separate regions for reads and writes. The regions do not overlap.

2.1.4. Intel Stratix 10 MX DMA On-Chip Memory Throughput

Figure 8. Intel Stratix 10 MX DMA Using On-Chip Memory Throughput

```
****************************************
Current DMA configurations
    Run Read  (card->system)  ? 1
    Run Write (system->card)  ? 1
    Run Simultaneous          ? 1
    Number of dwords/desc     : 2048
    Number of descriptors     : 128
    Total length of transfer  : 1024 KiB

Current run #: 1
Current time : Tue Jul 24 09:37:04 2018

DMA throughputs, in GB/s (10^9B/s)
    Current Read Throughput   : 12.95
    Average Read Throughput   : 12.95
    Current Write Throughput  : 12.95
    Average Write Throughput  : 12.95
    Current Simul Throughput  : 22.80
    Average Simul Throughput  : 22.80
****************************************
```

2.2. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with External DDR4 Memory

Figure 9. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with External DDR4 Memory Reference Design

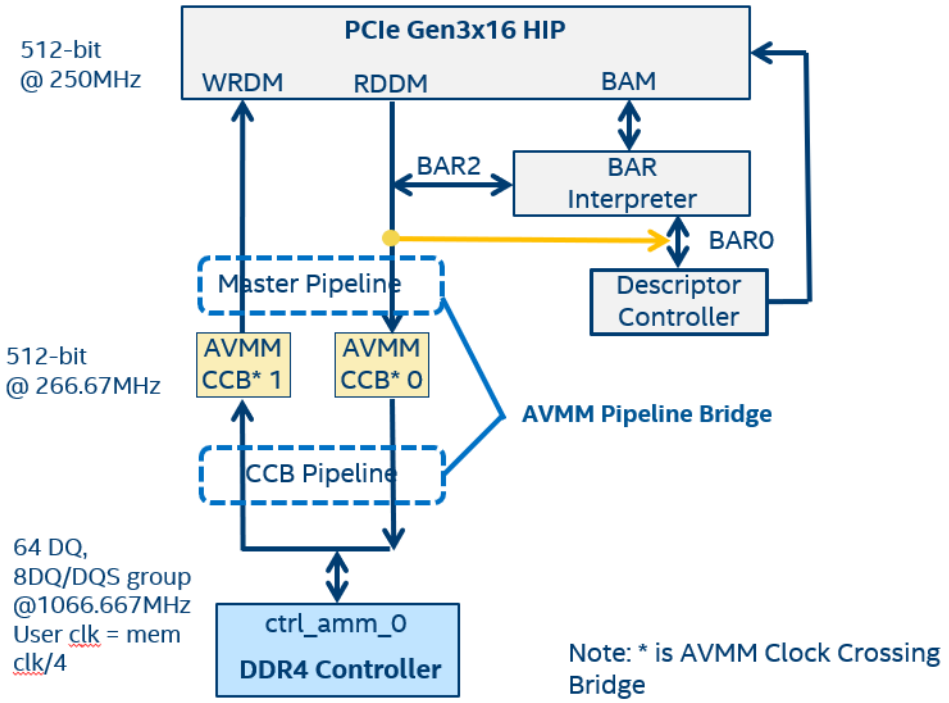

2.2.1. Project Hierarchy

The reference design uses the following directory structure:

Figure 10. Directory Structure

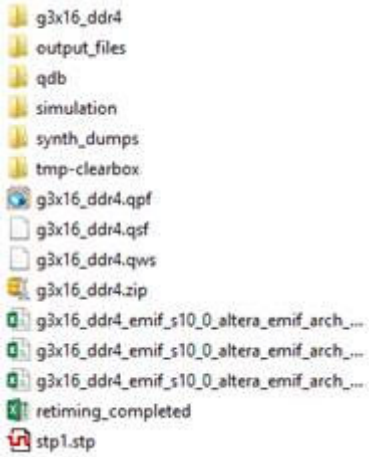

- g3x16_ddr4.v: The top-level module.

2.2.2. Parameter Settings for PCI Express Hard IP+ Variations

This reference design supports a 512-byte maximum payload size. The following tables list the values for all the parameters.

Table 12. System Settings

Parameter	Value
Number of lanes	Intel Stratix 10 MX: 16
Lane rate	Intel Stratix 10 Gen3: 16 Gbps
RX buffer credit allocation - performance for received request	Intel Stratix 10: Not available
Hard IP Mode	By default, the Hard IP mode is set to Gen3 x16, with a 512-bit interface to the Application Layer running at 250 MHz.

Table 13. Base Address Register (BAR) Settings

Parameter	Value	BAR Size
BAR0	64-bit prefetchable memory	16 bits
BAR1	Disabled	
BAR2	64-bit prefetchable memory	30 bits
BAR3	Disabled	
BAR4	Disabled	
BAR5	Disabled	

Table 14. Device Identification Register Settings

Parameter	Value
Vendor ID	0x00001172
Device ID	0x0000E003
Revision ID	0x00000001
Class Code	0x00000000
Subsystem Vendor ID	0x00000000
Subsystem Device ID	0x00000000

Table 15. PCI Express/PCI* Capabilities

Parameter	Value
Maximum payload size	512 bytes
Completion timeout range	None
Implement completion timeout	Disabled

Table 16. Error Reporting Settings

Parameter	Value
Advanced Error Reporting (AER)	Enabled
ECRC checking	Disabled
ECRC generation	Disabled

Table 17. Link Settings

Parameter	Value
Link port number	1
Slot clock configuration	Enabled

Table 18. Message Signaled Interrupts (MSI) and MSI-X Settings

Parameter	Value
Number of MSI messages requested	4
Implement MSI-X	Disabled
Table size	0
Table offset	0x0000000000000000
Table BAR indicator	0
Pending bit array (PBA) offset	0x0000000000000000
PBA BAR indicator	0

Table 19. Power Management

Parameter	Value
Endpoint L0s acceptable latency	Maximum of 64 ns
Endpoint L1 acceptable latency	Maximum of 1 us

Table 20. PCIe Address Space Setting

Parameter	Value
Address width of accessible PCIe memory space	40

2.2.3. PCIe Avalon-MM DMA Reference Design with DDR4 Memory Platform Designer System

The following image shows the modules in the Platform Designer system for this reference design.

Figure 11. Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express with DDR4 Memory Platform Designer Module

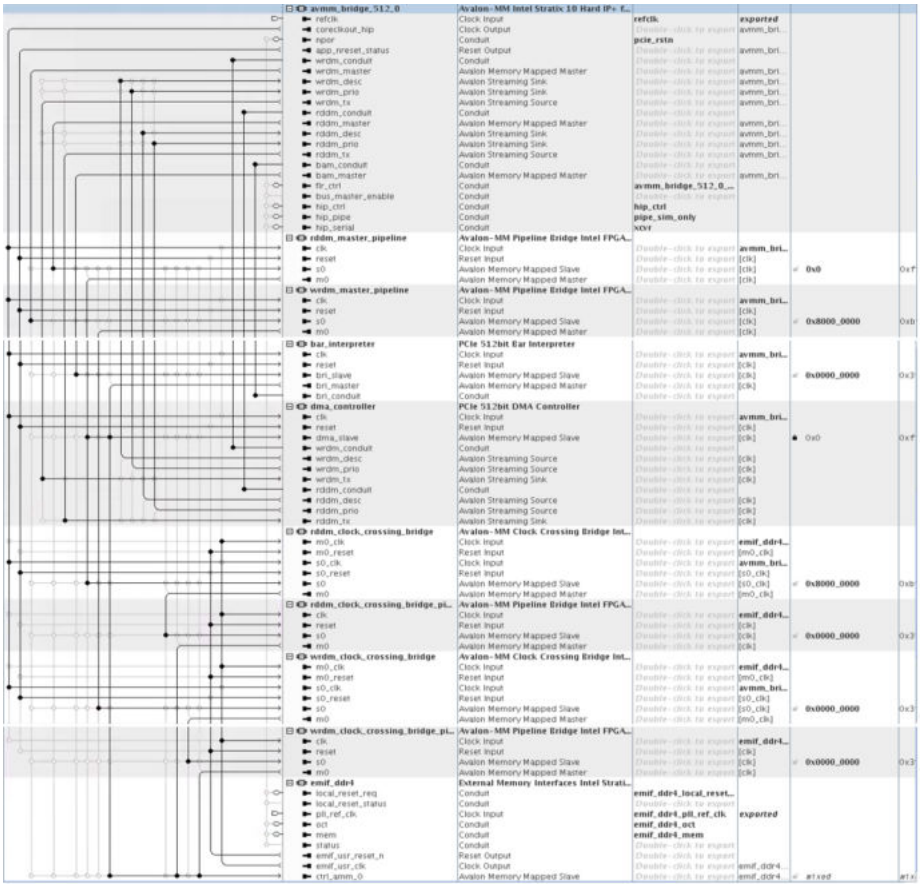

Table 21. Platform Designer Port Descriptions

Port	Function	Description
RDDM	Read Data Mover	This interface transfers DMA data from the PCIe system memory to the memory in Avalon-MM address space.
WRDM	Write Data Mover	This interface transfers DMA data from the memory in Avalon-MM address space to the PCIe system memory.
BAM	Bursting Avalon-MM Master	This interface provides host access to the registers and memory in Avalon-MM address space. The Bursting Avalon-MM Master module converts PCIe Memory Reads and Writes to Avalon-MM Reads and Writes.
Intel DDR4 Controller	DDR4 Controller	This is a single-port DDR4 controller with 64-DQ width and 8 DQ per DQS group.

2.2.4. Intel Stratix 10 MX DMA DDR4 Memory Throughput

Figure 12. Intel Stratix 10 MX DMA Using DDR4 Memory Throughput

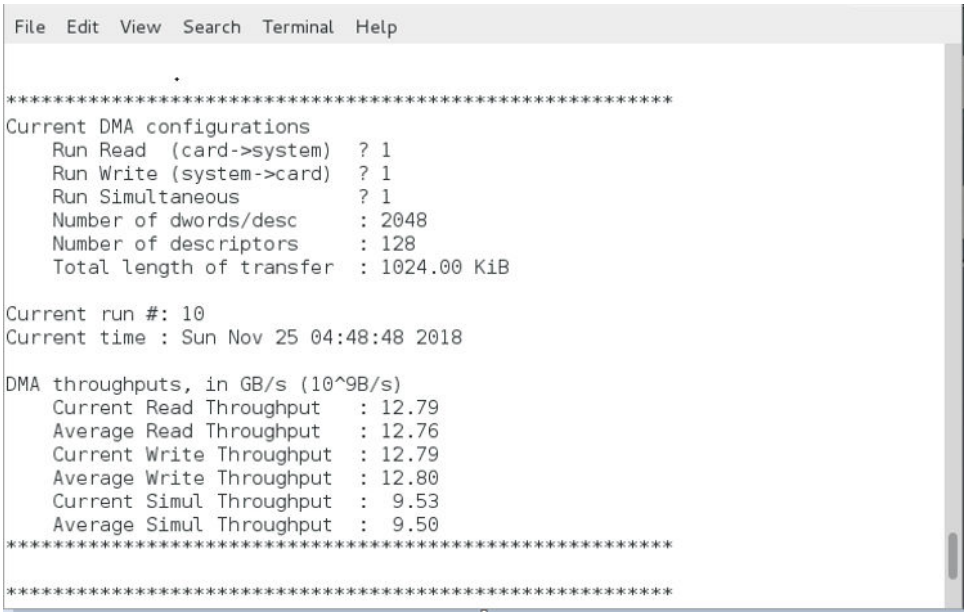

2.3. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with HBM2 Memory

Figure 13. Avalon-MM Intel Stratix 10 MX Hard IP+ DMA with HBM2 Memory Reference Design

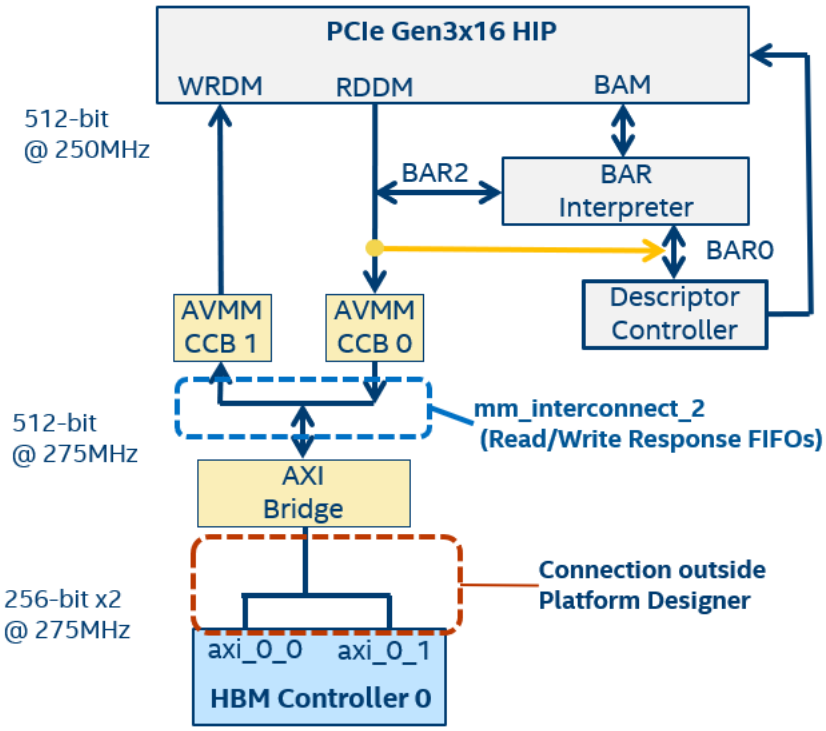

This Gen3 x16 DMA with HBM2 Platform Designer system (captured in the file `g3x16_hbm2.qsys`) instantiates two Avalon-MM clock-crossing bridges and AXI bridge IP cores between PCIe Avalon-MM Masters and HBMC AXI Slave. The purpose of those IPs is as follows:

- PCIe Hard IP and HBMC clock domain crossing
- Burst length adaptation
- Exporting the AXI Master interface
- Controlling the Read/Write Response FIFO depth

PCIe Hard IP and HBM clock domain crossing

The Gen3 x16 IP user interface is 512-bit @ 250 MHz. The 250 MHz is the frequency of the `coreclkout_hip` generated by the PCIe Hard IP. The HBM Controller AXI interface in the design is 256-bit @ 275 MHz. The HBM Controller core clock is generated by an IOPLL. Two Avalon-MM clock-crossing bridges are used to handle the clock crossing.

Burst length adaptation

The Gen3 x16 IP Write Data Mover (WRDM) and Read Data Mover (RDDM) Avalon-MM interfaces are bursting masters that issue Read/Write transactions in burst mode (the maximum burst count supported is 8). However, the HBM Controller AXI4 slave only supports single-burst transfers (burst length of 1). To resolve this, the maximum burst size in the Avalon-MM clock crossing bridges is set to 1.

Exporting the AXI Master interface

The design uses an AXI bridge to export the AXI Master interface from the Platform Designer system. The exported AXI Master interface is connected externally to the HBMC AXI Slave interfaces.

The AXI Bridge Read/Write address drives both HBMC AXI slaves.

The AXI Bridge Read/Write 512-bit data bus is split into two 256-bit data busses.

Controlling the Read/Write Response FIFO depth

The AXI Bridge Read/Write Acceptance Capability parameter setting dictates the Interconnect Read/Write Response FIFO depth generated by Platform Designer in the `altera_merlin_axi_slave_ni` module. The Response FIFO depth affects the Avalon-MM transaction performance.

If the Read/Write Response FIFO depth is not deep enough and the FIFO becomes full, it creates backpressure, impacting the throughput.

The default Read/Write Acceptance Capability parameter value is set to 16. Intel Quartus Prime 18.1 allows up to 32. In this design, the Read/Write Response FIFO depth generated by Platform Designer is manually changed to 64 in the `altera_merlin_axi_slave_ni` module in order to support Gen3 x16 throughput. In future Intel Quartus Prime releases, the AXI Bridge will support a higher maximum value.

2.3.1. Project Hierarchy

The reference design uses the following directory structure:

Figure 14. Directory Structure

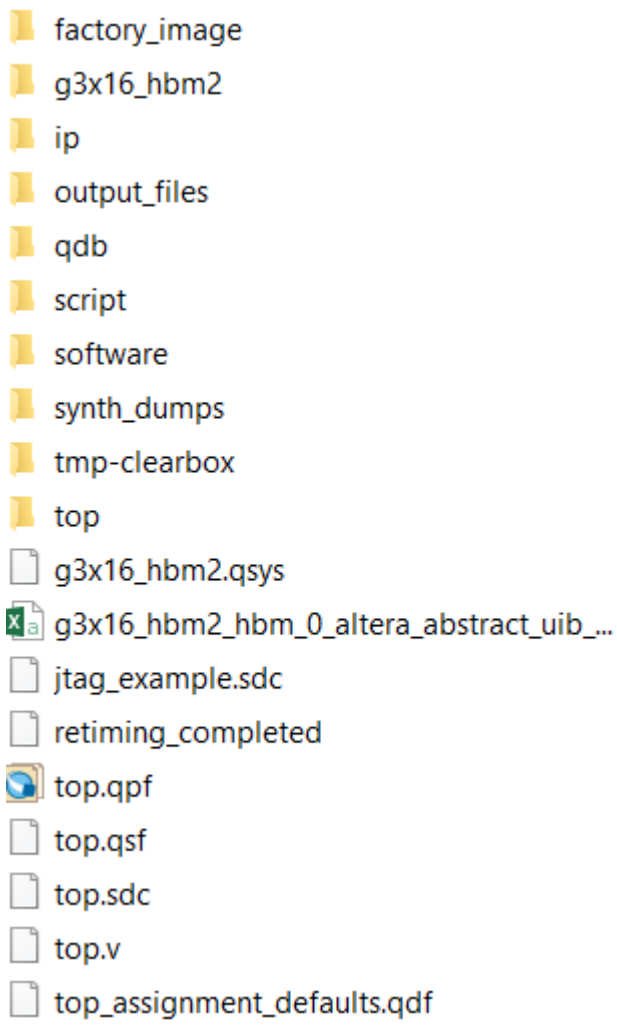

- top.v: The top-level module.

2.3.2. Parameter Settings for PCI Express Hard IP Variations

This reference design supports a 512-byte maximum payload size. The following tables list the values for all the parameters.

Table 22. System Settings

Parameter	Value
Number of lanes	Intel Stratix 10 MX: 16
Lane rate	Intel Stratix 10 Gen3: 16 Gbps
RX buffer credit allocation - performance for received request	Intel Stratix 10: Not available
Hard IP Mode	By default, the Hard IP mode is set to Gen3 x16, with a 512-bit interface to the Application Layer running at 250 MHz.

Table 23. Base Address Register (BAR) Settings

Parameter	Value	BAR Size
BAR0	64-bit prefetchable memory	16 bits
BAR1	Disabled	
BAR2	64-bit prefetchable memory	30 bits
BAR3	Disabled	
BAR4	Disabled	
BAR5	Disabled	

Table 24. Device Identification Register Settings

Parameter	Value
Vendor ID	0x00001172
Device ID	0x0000E003
Revision ID	0x00000001
Class Code	0x00000000
Subsystem Vendor ID	0x00000000
Subsystem Device ID	0x00000000

Table 25. PCI Express/PCI* Capabilities

Parameter	Value
Maximum payload size	512 bytes
Completion timeout range	None
Implement completion timeout	Disabled

Table 26. Error Reporting Settings

Parameter	Value
Advanced Error Reporting (AER)	Enabled
ECRC checking	Disabled
ECRC generation	Disabled

Table 27. Link Settings

Parameter	Value
Link port number	1
Slot clock configuration	Enabled

Table 28. Message Signaled Interrupts (MSI) and MSI-X Settings

Parameter	Value
Number of MSI messages requested	4
Implement MSI-X	Disabled

continued...

Parameter	Value
Table size	0
Table offset	0x0000000000000000
Table BAR indicator	0
Pending bit array (PBA) offset	0x0000000000000000
PBA BAR indicator	0

Table 29. Power Management

Parameter	Value
Endpoint L0s acceptable latency	Maximum of 64 ns
Endpoint L1 acceptable latency	Maximum of 1 us

Table 30. PCIe Address Space Setting

Parameter	Value
Address width of accessible PCIe memory space	40

2.3.3. PCIe Avalon-MM DMA Reference Design with HBM2 Memory Platform Designer System

The following image shows the modules in the Platform Designer system for this reference design.

Figure 15. Avalon-MM Intel Stratix 10 Hard IP+ for PCI Express with HBM2 Memory Platform Designer Module

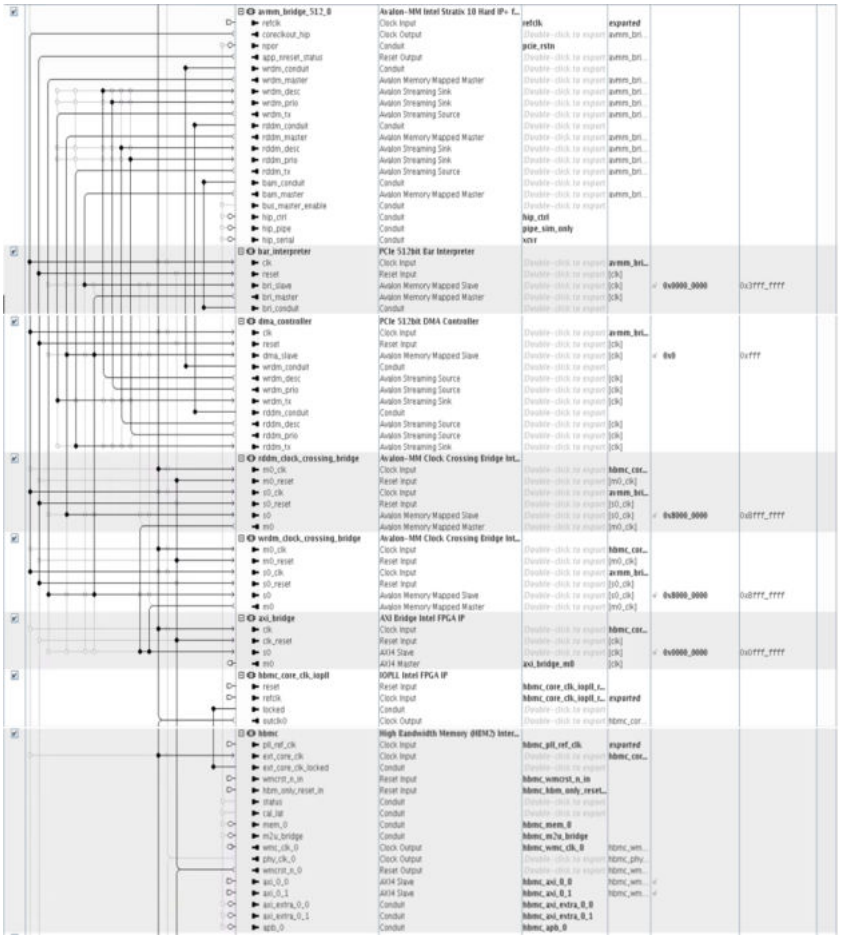

Table 31. Platform Designer Port Descriptions

Port	Function	Description
RDDM	Read Data Mover	This interface transfers DMA data from the PCIe system memory to the memory in Avalon-MM address space.
WRDM	Write Data Mover	This interface transfers DMA data from the memory in Avalon-MM address space to the PCIe system memory.
BAM	Bursting Avalon-MM Master	This interface provides host access to the registers and memory in Avalon-MM address space. The Bursting Avalon-MM Master module converts PCIe Memory Reads and Writes to Avalon-MM Reads and Writes.
HBM2 Memory Controller	HBM Controller	This is a single channel with 2 pseudo channel HBM Controllers. The user interface to the HBM2 Controller uses the AXI4 protocol. Each Controller has one AXI4 interface per pseudo channel or 2 AXI4 interfaces per channel.

2.3.4. Intel Stratix 10 MX DMA HBM2 Memory Throughput

Figure 16. Intel Stratix 10 MX DMA Using HBM Memory Throughput

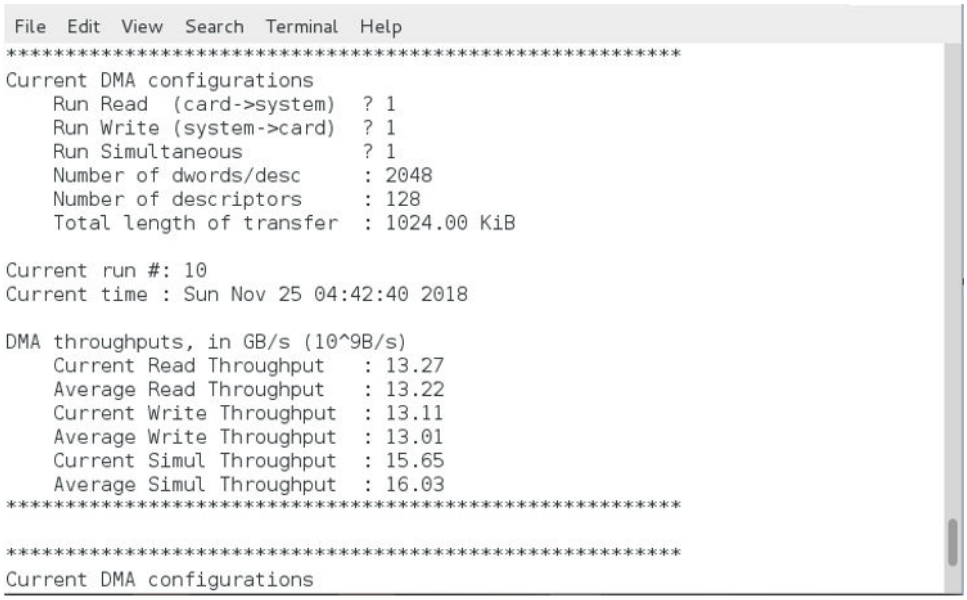

3. Understanding PCI Express throughput

3.1. Throughput for Posted Writes

The theoretical maximum throughput calculation uses the following formula:

```
Throughput  = payload size / (payload size + overhead) * link data rate
```

3.1.1. Specifying the Maximum Payload Size

The `Device Control` register, bits [7:5], specifies the maximum TLP payload size of the current system. The `Maximum Payload Size` field of the `Device Capabilities` register, bits [2:0], specifies the maximum permissible value for the payload. You specify this parameter, called **Maximum Payload Size**, using the parameter editor. The default value for this parameter is 512B. After determining the maximum TLP payload for the current system, software records that value in the `Device Control` register. This value must be less than the maximum payload specified in the `Maximum Payload Size` field of the `Device Capabilities` register.

Understanding Flow Control for PCI Express

Flow control guarantees that a TLP is not transmitted unless the receiver has enough buffer space to accept the TLP. There are separate credits for headers and payload data. A device needs sufficient header and payload credits before sending a TLP. When the Application Layer in the completer accepts the TLP, it frees up the RX buffer space in the completer's Transaction Layer. The completer sends a flow control update packet (FC Update DLLP) to replenish the consumed credits to the initiator. When a device consumes all its credits, the rate of FC Update DLLPs to replenish header and payload credits limits throughput. The flow control updates depend on the maximum payload size and the latencies of two connected devices.

3.2. Throughput for Reads

PCI Express uses a split transaction model for reads. The read transaction includes the following steps:

1. The requester sends a Memory Read Request.
2. The completer sends out the ACK DLLP to acknowledge the Memory Read Request.
3. The completer returns a Completion with Data. The completer can split the Completion into multiple completion packets.

Read throughput is typically lower than write throughput because reads require two transactions instead of a single write for the same amount of data. The read throughput also depends on the round trip delay between the time when the

Application Layer issues a Memory Read Request and the time when the requested data returns. To maximize the throughput, the application must issue enough outstanding read requests to cover this delay.

To maintain maximum throughput for the completion data packets, the requester must optimize the following settings:

- The number of completions in the RX buffer
- The rate at which the Application Layer issues read requests and processes the completion data

Read Request Size

Another factor that affects throughput is the read request size. Issuing less read requests with a larger read request size results in higher throughput than issuing more read requests with a smaller read request size. The `Maximum Read Request Size` value in `Device Control` register, bits [14:12], specifies the read request size. These reference designs support up to a 512B read request size. If the `Maximum Read Request Size` is set to a value ≤ 512B, the read request size will be that value. However, if the `Maximum Read Request Size` is set to a value > 512B, the read request size will be limited to 512B.

Outstanding Read Requests

A final factor that can affect the throughput is the number of outstanding read requests. If the requester sends multiple read requests to improve throughput, the number of available header tags limits the number of outstanding read requests. The Intel Stratix 10 read DMA can use up to 64 header tags.

3.2.1. Understanding Throughput Measurement

To measure throughput, the software driver takes two timestamps. Software takes the first timestamp shortly after the DMA is started. Software then takes the second timestamp after the DMA completes and returns the required completion status. If read DMA, write DMA and simultaneous read and write DMAs are all enabled, the driver takes six timestamps to make the three measurements.

3.3. Throughput Differences for On-Chip and External Memory

This reference design provides a choice between on-chip memory implemented in the FPGA fabric, external memory available on the PCB and HBM2 memory available as a system-in-package. The on-chip memory supports separate read and write ports. Consequently, this memory supports simultaneous read and the write DMAs.

The external memory supports a single port. Consequently, the external memory does not support simultaneous read DMA and write DMA accesses. In addition, the latency of external memory is higher than the latency of on-chip memory. These two differences between the on-chip and external memory result in lower throughput for the external memory implementation.

4. Document Revision History for AN 881: PCI Express Gen3 x16 Avalon-MM DMA with On-Chip, External or HBM2 Memory Reference Design

Document Version	Intel Quartus Prime Version	Changes
2018.12.27	18.1	Fixed typo in Figure 1.
2018.12.10	18.1	Initial release.