PCI Express Performance Measurement

User Guide

January 2015

Copyright and Legal Notice

History

The following table shows the History of this document.

Date	Version	Revision
13.02.2015	1.0	Initial Release

Table of Contents

1 Features of Smartlogic's PCIe DMA IP Core

Smartlogic's PCI Express IP Core Suite allows the development of complex PCI-Express endpoints within a Xilinx FPGA with only basic PCI Express protocol Know-how. The IP is built around the Xilinx PCI-Express Hard IP Macrocell and greatly simplifies the development of a powerful PCI-Express Endpoint due to standardized AXI Stream Slave and Full AXI Master Interfaces.

With this IP Core it is possible to manage up to 16 independent buffers in the Host Memory. With the powerful addressing and arbitration scheme it is possible to control the priority of the DMA data and to manage nearly any geometry in the Host Memory.

IP-Features

- PCI-SIG compliant solution
- Integrates into Xilinx IP-Integrator Flow
- Supports X1, X2, X4 or X8 Link Widths at Gen1, Gen2 or Gen3 speed
- Up to 8 AXI4 Masters (Burst & Single Transfers) to allow easy access of other FPGA peripherals
- Supports 32 Bit and 64 Bit DMA Addressing
- Support for either Scatter Gather or contiguous Memory
- Flexible and easy customizable Register Interface
- MSI and Legacy Interrupt Support
- Powerful customization of the included Xilinx Hardmacrocell regarding BARs, Reference Clock Speed and many other parameters
- Supported Technologies : Virtex 6, Artix 7, Kintex 7, Virtex 7, Zynq
- Up to 16 independent AXI Stream Slaves for DMA Write transfers
- Up to 16 independent AXI Stream Masters for DMA Read transfers
- Independent Clocking of each AXI Stream channel possible
- Parametrizable Priority Control
- Memory Size up to 4 GByte per Streaming Channel
- Performance only limited by PCI-Express Bandwidth
- High Speed Simulation Mode

Block Diagram of the IP-Core:

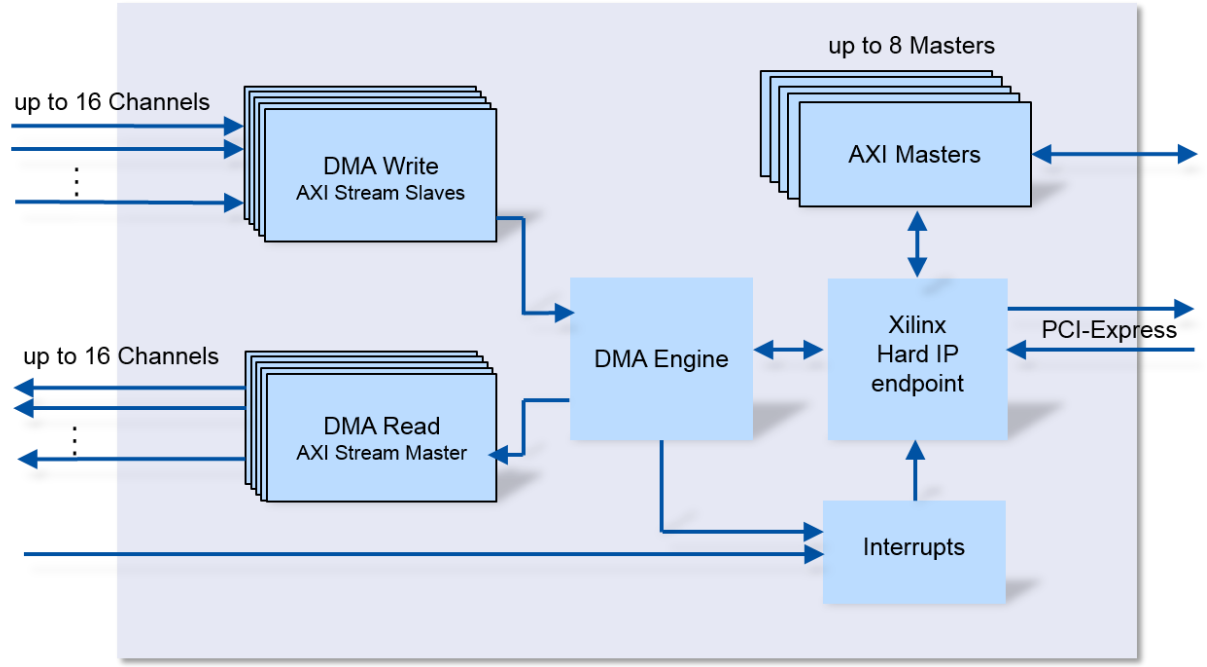

FPGA Resources:

The following table lists different example configurations and the needed FPGA Resources:

FPGA Family	Configuration Width / Speed	User Registers[1]	FPGA-Resources[2] LUTs	FFs	BRAM	PCIE	Misc
Artix 7	64 bit Architecture X4, Gen 1 4 DMA Channels, 1 Master	15 x 32	8.803	8.521	12	1	1 MMCM
	256-Bit Architecture, X4, Gen 1 4 DMA Channels, 1 Master	15 x 32	15.934	13.156	12	1	1 MMCM

(1) The amount of user registers is design dependant

(2) These numbers include the resources of the Xilinx Endpoint Block Plus Hardcore IP

Link Speeds and Link Widths:

The following table lists the supported PCI-Express Link Speeds and Link widths for the different versions of the IP Core:

Product Version	FPGA Family	Supported Link Speed / Link Width						
		Gen 1, 2.5 GBit/s			Gen 2, 5.0 GBit/s		Gen 3, 8.0 GBit/s	
		X1	X4	X8	X1	X4	X1	X4
64-Bit Core	Virtex 6, Kintex, Artix*, Zynq	✓	✓	✓	✓	✓		
256-Bit Core	Virtex 7, Kintex, Artix*	✓	✓	✓	✓	✓	✓**	✓**

(*) Artix does not support a x8 Link Width, Gen2 X4 is only supported for the 256-Bit Version of the Core

(**) Gen3 Link Speed only with Virtex-7 Devices possible

Supported Toolchain:

The following table lists the supported tools:

Tool	Version
P&R	Vivado 2013.3 or higher for IPI Flow ISE 14.7 or Vivado 2013.3 for HDL Flow
Simulator	Modelsim 10.0 and 10.1, 10.3c, Vivado Simulator**

* Note, that the Xilinx Simulation models do not work with ModelSim PE/DE 10.2 up to 10.3b ! See Xilinx Answer Record 61652 for more details. In order to workaround this issue in the not supported ModelSim PE Versions, either use ModelSim SE or use only the High Speed Simulation Mode.

** usage of the Vivado Simulator is not recommended, since not all Versions simulate correctly. Especially Vivado ISIM 2014.4 does not work

2 PCI Express Performance Demodesign and Windows GUI

The PCI Express Performance Demodesign is a free evaluation application of our PCI Express IP Core which measures various significant PCI Express parameters. The measurement results will reveal the overall performance of the Host and important parameters like Maximum Payload Size (MPS), Maximum Read Request Size (MRS), Payload throughput, CRC Errors and others.

Compiled Bitstreams are available from Smartlogic for various demoboardplatforms as Xilinx AC701, KC705, VC709 and others for evaluation purposes at no cost.

The Demodesign is bundled with a Windows 7 Driver and a Graphical User Interface (GUI).

Supplied files and folder structure:

Folder	Remark
amd64\ amd64\sldemo_64.cat amd64\sldemo_64.inf amd64\sldemo_64.sys	Compiled Windows 7 driver 64bit
perf_test\ Perftest.exe QtCore4.dll QtGui4.dll	Compiled Performancetest GUI for Windows

Note: For security reasons, the downloaded ZIP File is password protected. The password can be obtained at no cost by simply sending a request email to perf_eval[at]smartlogic.de . The password will be sent immediately to you. If you need a 32-Bit Windows 7 Device Driver, please get in contact with us and write an email to ip[at]smartlogic.de .

The following paragraphs provide guidelines how to install the Bitstream and the Driver and how to work with the GUI.

Installation:

The first installation step is to program the performance demo bitstream into the desired demoboard platform.

The following User Guides from Xilinx show in detail how to do this:

Demoboard	Xilinx User Guide
AC701	XTP227 See Sections “AC701 Setup”, “Hardware Setup” and “Program AC701 Flash with PCIe Design”
KC705	XTP197 See Sections “KC705 Setup”, “Program KC705 Flash with PCIe Design”
VC709	XTP237 See Sections “VC709 Setup”, “Program BPI Flash with PCIe Design”

In order to download these user manuals, go to the Xilinx Webpage (www.xilinx.com) and type in the Xilinx User Guide Name in the Search field (e.g. “XTP227” for the Artix Version). You may need a Xilinx Account which can be created very easiliy at no cost.

When the Bitstream is programmed into the Demoboard, connect the board with the Host and power the system on.

After Windows has bootet the new PCIe Device will be detected from Windows and you will be asked from Windows to install a device driver.

Decompress the downloaded zip folder that ontains the device driver to a local folder on your harddrive and instruct windows to install the driver from the local folder.

After successful installation re-boot Windows.

Installation of the GUI

The GUI itself does not need a specific installation procedure. Simply execute perftest.exe from the local folder. Make sure, that the two DLL Diles QTCore4.dll and QTgui4.dll reside in the same filder where perftest.exe is located.

If everything was installed correctly, the GUI should come up with the following screen:

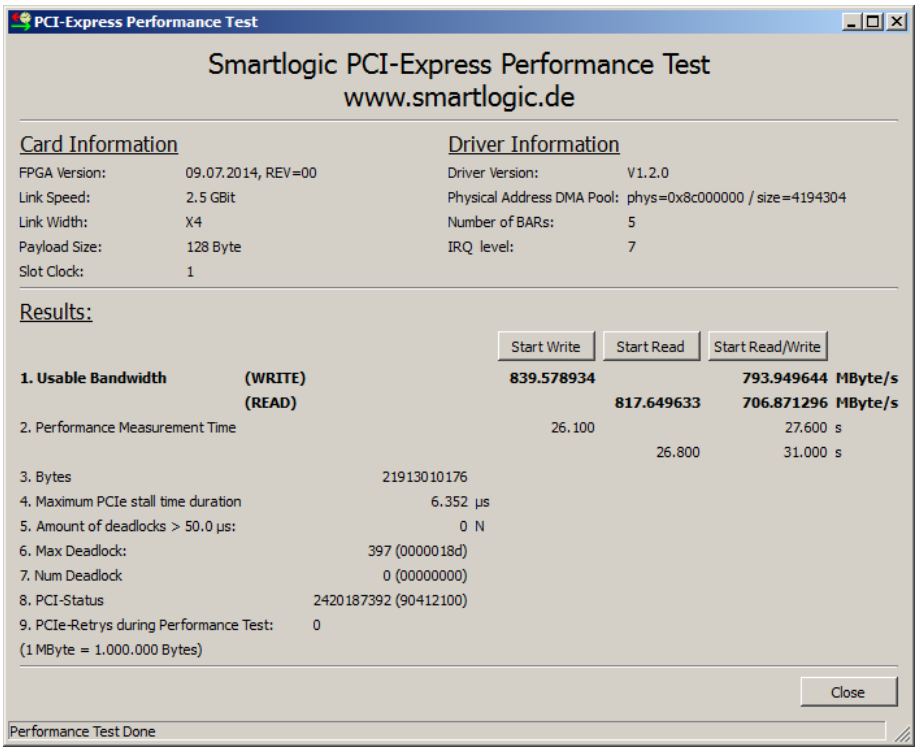

De-Installation

In order to remove the GUI from you PC, simply delete the EXE and the DLL Files from the local folder.

The Driver can be removed easily from the “Device Manager”. Simply click on the SL_DEMO Device and de-install the driver.

3 Working with the GUI – Measurement Results

The GUI displays several informational details about the connected endpoint in the “Card Information” Section.

In order to measure the real payload throughput (Payload = User Data without protocol Overhead) simply press one of the three buttons “Start Write”, “Start Read” or “Start Read/Write”. This will generate a well known amount of data traffic in the desired direction. The FPGA measures the time to transfer the data. After the traffic generators have sent their data, the time and the MB/s value is displayed on the GUI.

“Write” refers to the direction from the FPGA to the Host and “Read” refers to the direction from Host to FPGA.

Note: These measurements provide the real data throughputs and contains all impacts like protocol overhead, possibly poor flow control update and power management setting. The results can differ significantly on different host platforms.

Further Results after a performance measurement:

The following table lists all further informations that are available after a measurement run:

Result	Explanation
3. Bytes	This is the amount of bytes that was transferred over PCI Express for measuring the performance. The amount is calculated out of Link Width and Link Speed in order to have a test duration of approximately 30 Seconds
4. Maximum PCIe Stall time duration	This is the longest time duration during the measurement, where no data could be transferred over PCIe. Typical values are 1us to 12 us. Greater values indicate, that there are severe system impacts like power management issues. This value is important do know in order to design FIFO depths big enough without data overflow.
5. Amount of deadlocks > 50 us	This value shows how often a PCIe Stall happened with a duration of greater 50 us while the performance test was running.
6. Max Deadlock	This is the same information as 4 (raw information)
7. Num Deadlock	This is the same information as 5 (raw information)
8. PCI-Status	This value contains side information on the PCI Express Link. This is a 32-Bit value (hex in brackets) and has the following meaning: Bit 0 : Correctable Error Reporting Enable Bit 1: non fatal Error Reporting Enable Bit 2: Non_fatal reporting enable Bit 3: Unsupported Request Reporting enable Bit 4 : Enable relaxed ordering Bits 7:5 : Maximum Payload Size Bit 8: Extended Tag supported Bit 9 : Phantom Functions enable Bit 10 : Auxiliary Power PM Enable Bit 11 : Enable no snoop Bits 14:12 : Maximum Read Request size Bits 19:16: current Link Speed Bits 25:20: current Link Width Bit 31: MSI Enabled
9. PCIe Retries during Performance Test	This field reports the number of CRC Errors encountered during the measurements.

