

Optimal Solution for Recording Application!


2-Nov-17

Design Gateway

Page 1

Agenda

- **NVMe SSD Overview**
 - SSD Trends
 - Merit of NVMe SSD for embedded system
- **NVMe-IP Introduction**
 - Summary
 - Function
 - User Interface
 - Performance and Size
 - Development Environment/Reference Design
- **Application**



2-Nov-17

Design Gateway

Page 2


SSD Trends 1

- SATA interface is now performance bottle neck
 - SSD Read/Write speed is limited to 600MB/sec SATA bandwidth
- Move to PCI Express for faster speed
 - PCIe GEN3 x4lane can provide 4GB/sec transfer speed
- M.2 and BGA package suitable for compact application
 - M.2: Wid=22mm, Len=20/42/80/120mm DIMM-like small outline
 - BGA: 20mm x 16mm x 1.5mm, 1gram package



Current 2.5" SATA SSD

2-Nov-17



Latest M.2 type PCIe SSD

Design Gateway



BGA type SSD

Page 3

SSD Trends 2

- Host Controller Standard moves from AHCI to NVMe
 - Latest standard to extract maximum performance of SSD
 - Extended Queue size, 65536 concurrent command process
 - Most OS provides NVMe driver

FlashMemory NVMe™ Driver Ecosystem

Windows 8.1, Windows Server 2012 R2, Windows 8, Windows 10, redhat, suse, ubuntu, solaris, FreeBSD, vmware ESXi 6.0, Windows Server 2012, CERTIFIED FOR Windows Server 2008 R2, Windows 7, vmware ESXi 5.5

Native / in-box
Install NVMe driver

FMS2015 "Annual Update on Interfaces" presentation

2-Nov-17

Design Gateway

Page 4

Merit of NVMe SSD for Embedded System 1

- High Bandwidth: 3.4GB/s for Read, 2.0GB/s for Write
- Cost effective: Cost difference from SATA SSD is small



\$299.99 for 512GB

Samsung
Samsung 960 PRO Series - 512GB PCIe NVMe - M.2 Internal SSD (MZ-V6P512BW)
★★★★★ 138 customer reviews | 119 answered questions

Was: \$324.00
Price: **\$299.99** & FREE Shipping. Details
You Save: \$24.01 (7%)

In Stock.
Want it Wednesday, May 24? Order within **11 hrs 26 mins** and choose **One-Day Shipping** at checkout. Details
Ships from and sold by Amazon.com. Gift-wrap available.

Capacity: 512 GB

1TB	2TB	512 GB
\$579.99	\$1,191.38	\$299.99

CrystalDiskMark 5.0.2 x64

All	Read [MB/s]	Write [MB/s]
Seq Q32T1	3469	2000
4K Q32T1	636.2	522.2
Seq	2833	1755
4K	27.84	193.5

<http://benchmarkreviews.com/41954/samsung-960-pro-nvme-ssd-review/6/>

(Amazon.com 23-May-2017)

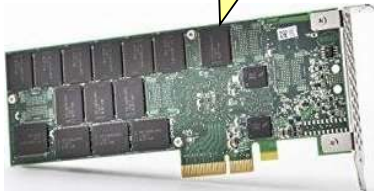
Cost and Performance of M.2 NVMe SSD (Samsung 960 PRO 512GB)

Merit of PCIe SSD for Embedded System 2

- Various form factor
 - HDDL(Half-Height,Half-Length) general PCIe board
 - M.2 cost saving module
 - SFF-8639 of 2.5" drive compatible size
 - BGA package for direct mount on PCB


Merit →

Big Capacity




HHHL PCIe board

Small, Extractable




M.2 module (length=42/60/80mm)

Hot Swap



SFF-8639 package

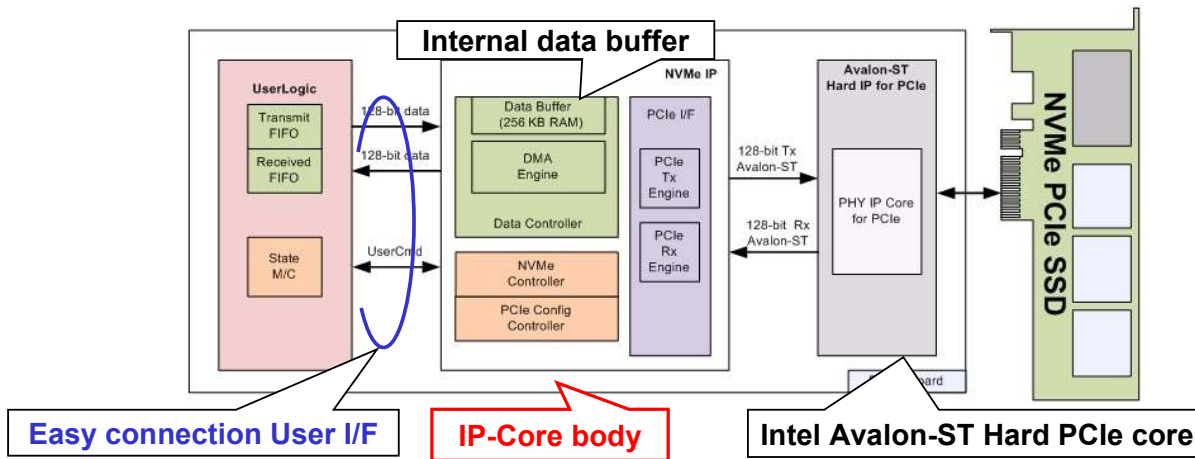
Mount on PCB



BGA package

What's NVMe-IP

- What's NVMe-IP? -> Directly connect NVMe SSD with FPGA
- How to use? -> Just connect with user logic, no need CPU, its F/W, or External DDR memory
- Application -> Best for ultra high speed data recording system
- User Merit? -> Can develop Storage Application in short period




2-Nov-17

Design Gateway

Page 7

NVMe-IP Merit

1. **Function: Full automatic access to PCIe SSD**
 - No CPU and firmware necessary, just wired logic is enough
2. **Interface: Simple and easy connection**
 - Direct connection to Intel Avalon-ST PCIe hard IP-core
 - User I/F control is parameter with pulse, data is simple FIFO
 - Use Block Mem. for data buffer (external DDR memory not required)
3. **High Performance and Compact size**
 - **Write=2145MB/s, Read=3347MB/s**
 - Support PCIe GEN3 (Operation confirmed on Arria10SoC board)
 - Core size: 1144ALM, 2120DFF (for Arria10SX case)
4. **Environment: Full reference design project**
 - Full QuartusII project with real board operation in the package




2-Nov-17

Design Gateway

Page 8

NVMe-IP Merit 1: Function

- Special PCIe Root port for NVMe SSD control
 - PCIe Initialization (Avalon-ST PCIe hard IP-core and NVMe register)
 - PCIe and SSD Status Monitor and error detection
- NVMe Control function
 - Control NVMe register by user R/W request
 - Process command sequence in NVMe standard
 - Data transfer and flow control between SSD, BlockRAM, and User FIFO



2-Nov-17


Design Gateway

Page 9

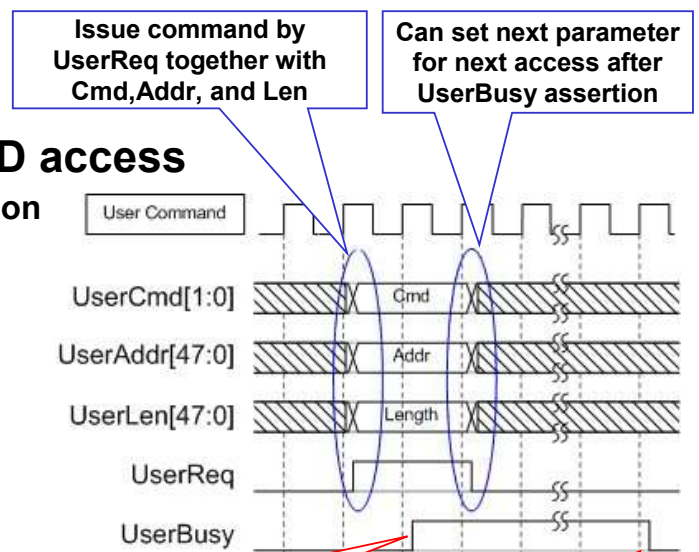
NVMe-IP Merit2: Command I/F



- Easy Connection User I/F
 - Set Command/Address/Length
 - Issue UserReq pulse
- Full Automatic control for SSD access
 - User only can wait UserBusy negation



Command I/F Signals



IP-Core asserts UserBusy='1' and start operation

UserBusy='0' when operation finish

Command I/F waveform

2-Nov-17

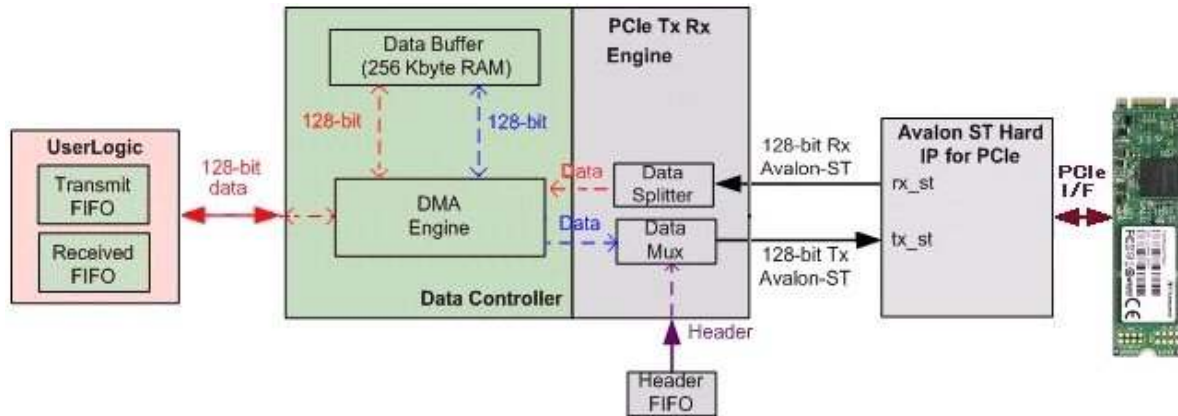
Design Gateway

Page 10

NVMe-IP Merit2: Data I/F



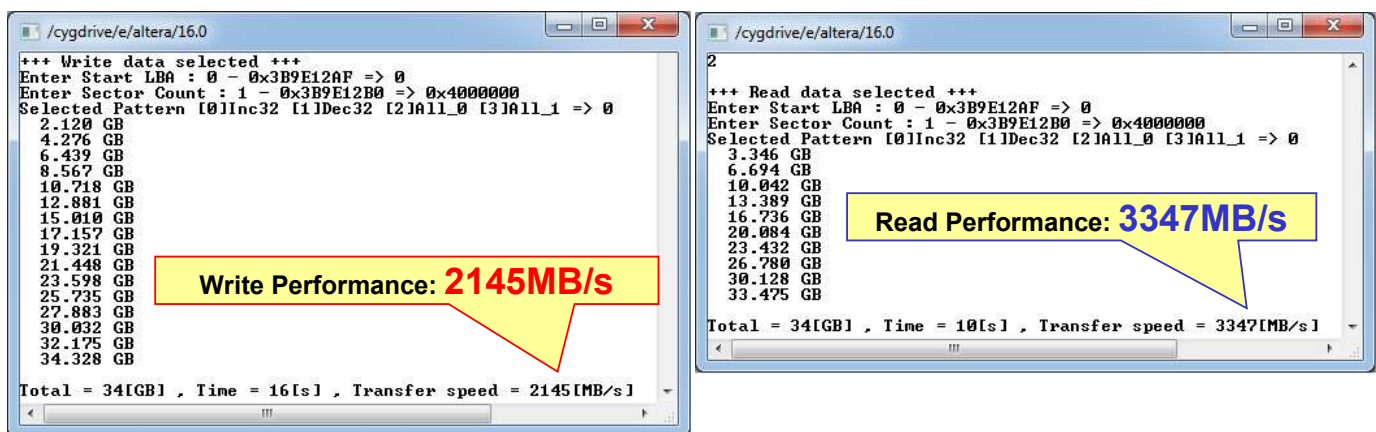
- Simple 128bit FIFO for each of read and write
 - General FIFO of standard Intel library
 - Data buffer using 256KByte Block memory in NVMe-IP



Data path of NVMe-IP

NVMe-IP Merit3: Performance

- Automatic PCIe SSD access by pure hard-wired logic
 - Intelligent state machine for complete read/write command execution
 - Minimum over head and best performance by synchronized circuit



Performance Evaluation Result (Arria10SoC)

(SSD: Samsung MZ-V6P512BW)

NVMe-IP Merit3: Compact Size


- **Optimized size with minimum resource consumption**
 - Implements dedicated and optimized logic for NVMe SSD control
- **Block memory for data buffer**
 - Internal block memory can minimize access overhead

Family	Example Device	Fmax (MHz)	Logic utilization (ALMs)	Registers	Block Memory bit ¹	Design Tools
ArriaV GX	5AGXFB3H4F35C4	212	1175	2133	2,162,688	QuartusII 16.0
Arria10 SX	10AS066N3F40E2SGE2	280	1144	2120	2,162,688	QuartusII 16.0

NVMe-IP Core standalone resource usage

NVMe-IP Merit4: Environment


- **Real operation check with Intel evaluation board**
- **Free sof-file for evaluation before IP-core purchase**



NVMe-IP evaluation environment

NVMe-IP Merit4: Development Tool

- Adapter board for FPGA board evaluation (Part#: AB16-PCIeXOVR)
- Connect FPGA board to PCIe socket on component side
- Connect PCIe SSD to PCIe socket on solder side
- SSD R/W access via adapter board from NVMe-IP in FPGA




2-Nov-17

Design Gateway

Page 15

NVMe-IP Merit4: Reference Design

- QuartusII/Qsys project is attached with NVMe-IP deliverables
- Full source code (VHDL) except IP core
- Can save user system development duration
 - Confirm real board operation by original reference design.
 - Then modify a little to approach final user product.
 - Check real operation in each modification step.



Short-term development is possible without big turn back

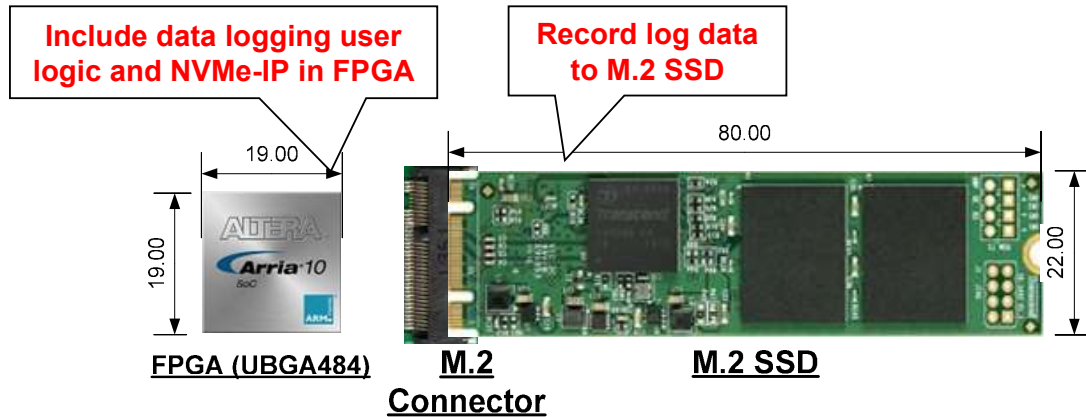
2-Nov-17

Design Gateway

Page 16

NVMe-IP Application Example 1

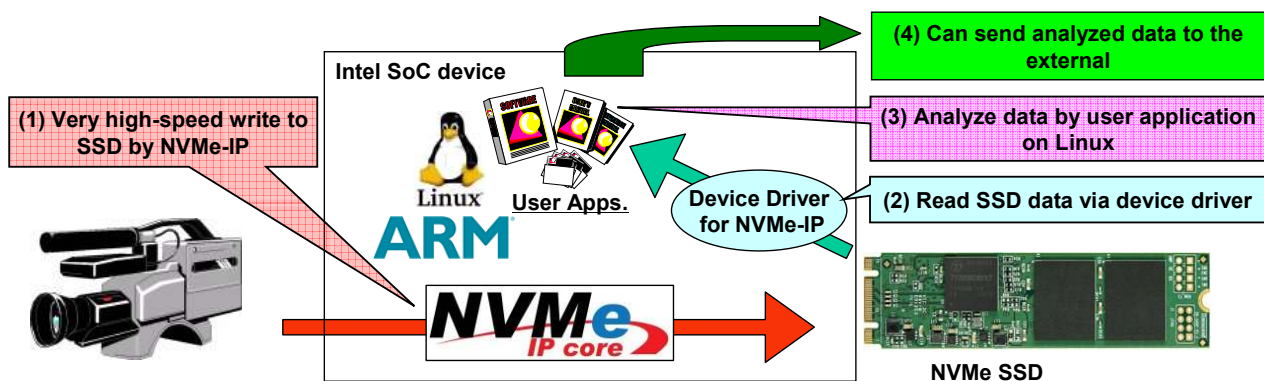
- Space-Saving FPGA data logging system
 - LatestFPGA+M.2 type SSD



System space image by UBGA 484 FPGA and M.2 SSD (unit: mm)

NVMe-IP Application Example 2

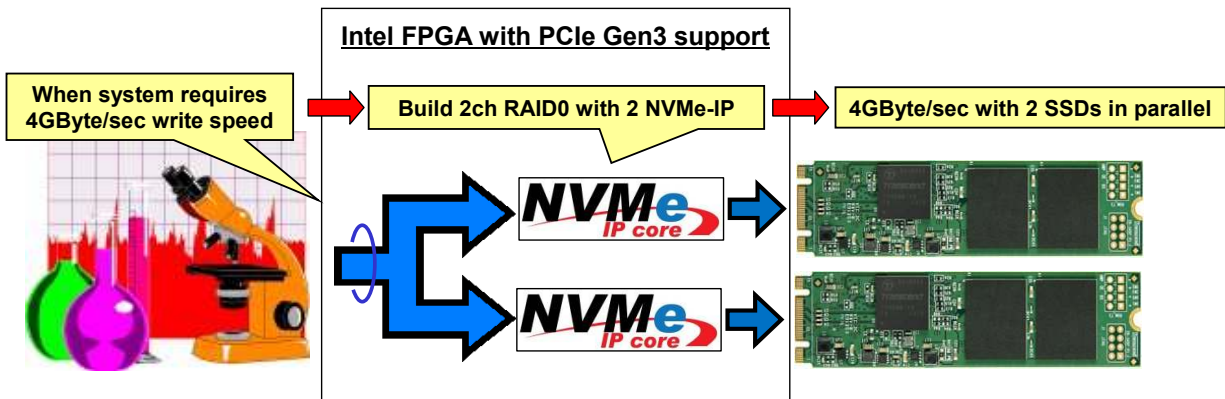
- Recording and Analysis system on Linux
 - Mount Linux and user analysis application on SoC device
 - Very high-speed data recording to SSD via NVMe-IP core
 - Data read from SSD via device driver and analyze by user application



Recording and Analysis system on Linux (device driver and reference design available)

NVMe-IP Application Example 3


- Ultra High-Speed Recorder
 - Double write speed with multiple SSDs RAID0 configuration
 - Provide RAID0 reference design with 2 NVMe SSDs



NVMe RAID system supporting 4GByte/sec recording rate

For more detail

- Detailed technical information available on the web site.
 - http://www.dgway.com/NVMe-IP_A_E.html
- Contact
 - Design Gateway Co., Ltd.
 - sales@design-gateway.com
 - FAX: +66-2-261-2290



Revision History

Rev.	Date	Description
0.1E	4-Aug-16	English Temporary Version (Ver0.1E)
1.0E	10-Aug-16	First release with resource usage information
1.1E	25-Aug-16	Modify page17 because only one x16 DDR4 device can keep NVMe SSD performance
1.2E	21-Dec-16	NVMe-IP core improvement by removing external DDR chip for data buffer
1.3E	23-May-17	Performance improved by internal PCIe bridge in NVMe-IP core
1.4E	6-Jun-17	Data buffer size fixed to 256KByte
1.5E	2-Nov-17	Added Linux driver application and 2ch RAID0 reference design