

Low-Latency Solutions for Storage-Hungry Embedded Applications

“Flash-on-Ethernet?”

Dr. Endric Schubert, MLE

Dr. David Boggs, MLE

Dr. Ulrich Langenbach, Fraunhofer HHI

Dr. Peter Gregorius, Fraunhofer HHI

System-of-systems are loosely coupled Embedded Systems which greatly benefit from the high performance of modern SSD technology. Machine visioning, medical imaging, and advanced driver assist systems are among those storage-hungry applications. However, communication latency and bandwidth in between the systems have a significant impact on the overall robustness, cost and performance.

Current techniques based on fieldbuses such as CAN, Flexray, have begun to hit the bandwidth wall and are more and more replaced by multi-Gigabit Ethernet plus techniques for hardware-acceleration of networking protocol stacks.

We present a proof-of-concept implementation specifically targeted for storage-hungry System-of-systems. Integrated into a modern FPGA with multicore ARM CPUs to run Open Source Linux, single-chip solutions become possible which provide full compatibility with all relevant network and storage interface protocols and can reach userspace latencies within few microseconds.

Missing Link Electronics is ...

We are a Silicon Valley based technology company with offices in Germany. We are partner of leading electronic device and solution providers and have been enabling key innovators in the automotive, industrial, test & measurement markets to build better Embedded Systems, faster.

Our mission is yo develop and market technology solutions for Embedded Systems Realization via pre-validated IP and expert application support, and to combine off-the-shelf devices with Open-Source Software for dependable, configurable Embedded System platforms.

Our expertise is I/O connectivity and acceleration of data communication protocols, additionally opening up FPGA technology for analog applications, and the integration and optimization of Open Source Linux and Android software stacks on modern extensible processing architectures.

MLE is a technology partner of Fraunhofer Heinrich-Hertz-Institute, a Certified Xilinx Alliance Partner, a member of the Altera Design Service Network, and an active contributor to the Open Source software ecosystem.



Fraunhofer HHI is...

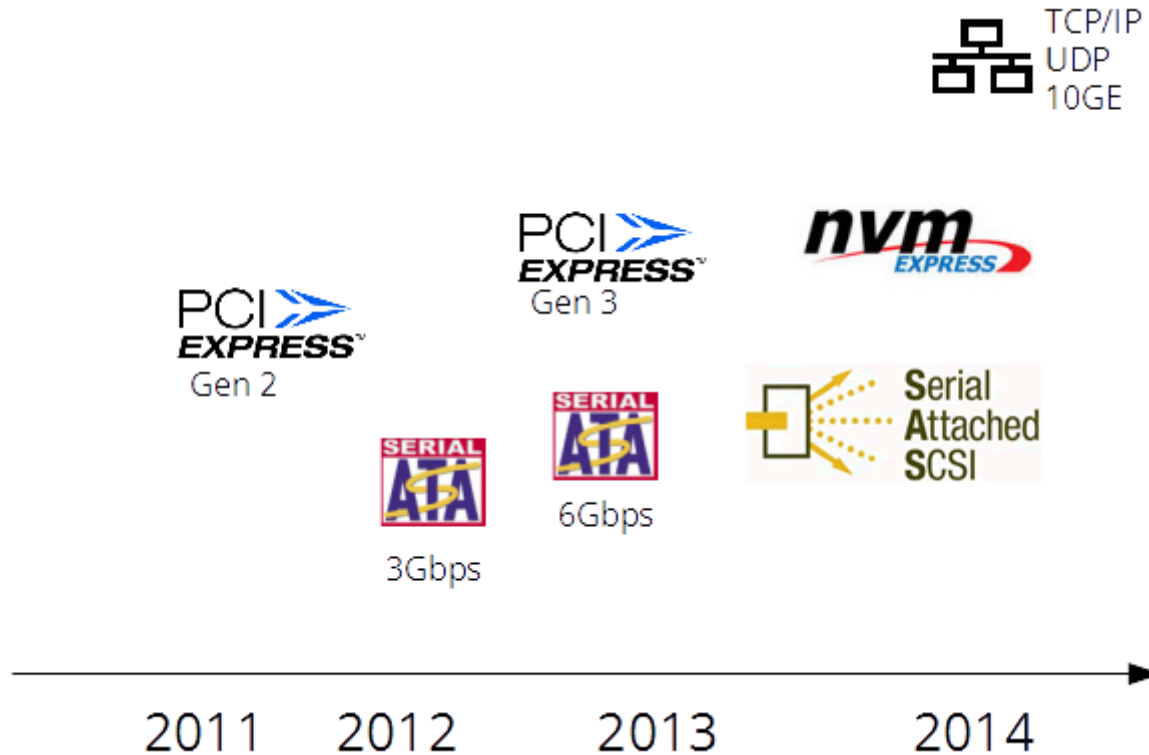
Founded in 1949, the German Fraunhofer-Gesellschaft undertakes applied research of direct utility to private and public enterprise and of wide benefit to society. With a workforce of over 23,000, the Fraunhofer-Gesellschaft is Europe's biggest organization for applied research, and currently operates a total of 67 institutes and research units. The organization's core task is to carry out research of practical utility in close cooperation with its customers from industry and the public sector.

Fraunhofer HHI was founded in 1928 as "Heinrich-Hertz-Institut für Schwingungsforschung" and joined in 2003 the Fraunhofer-Gesellschaft as the "Fraunhofer Institute for Telecommunications, Heinrich-Hertz-Institut,

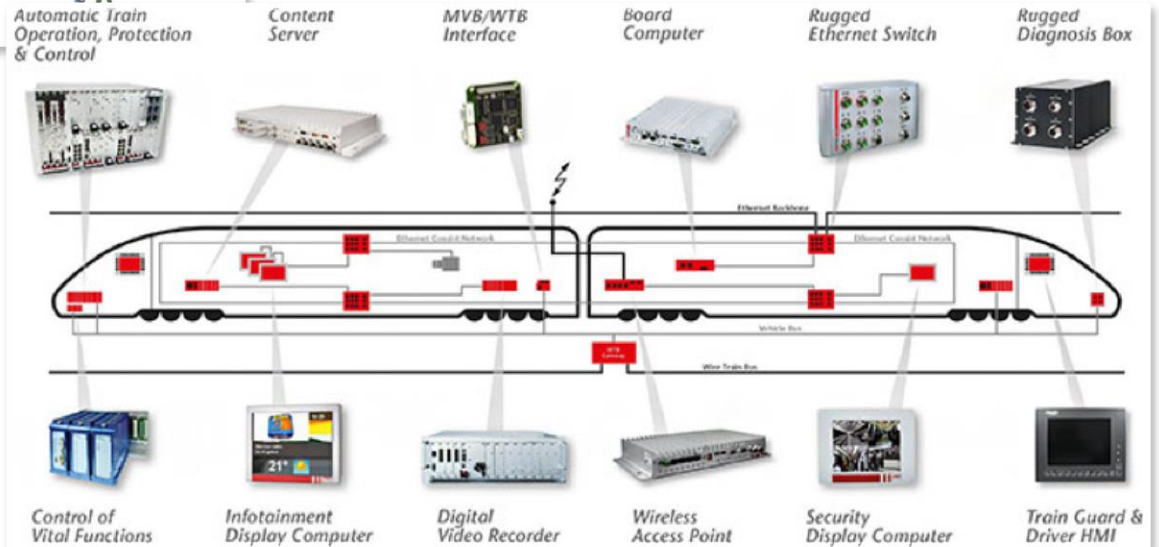

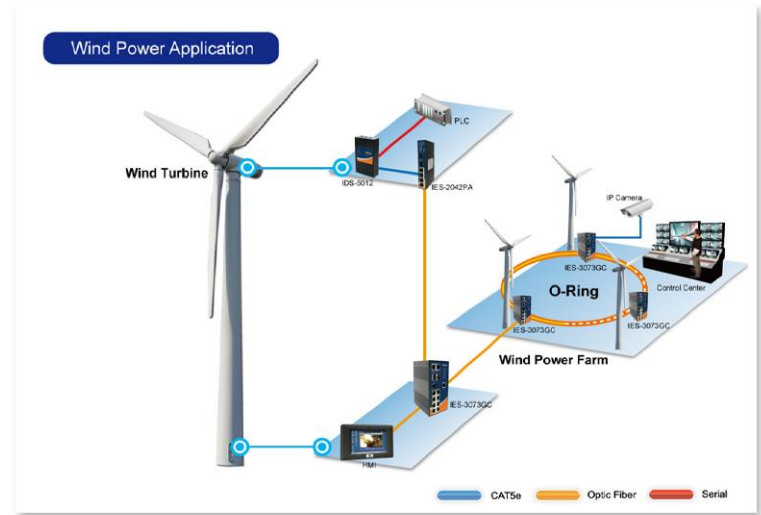

Today it is the leading research institute for networking and telecommunications technology, "Driving the Gigabit Society" .



Hardware Acceleration of Protocols

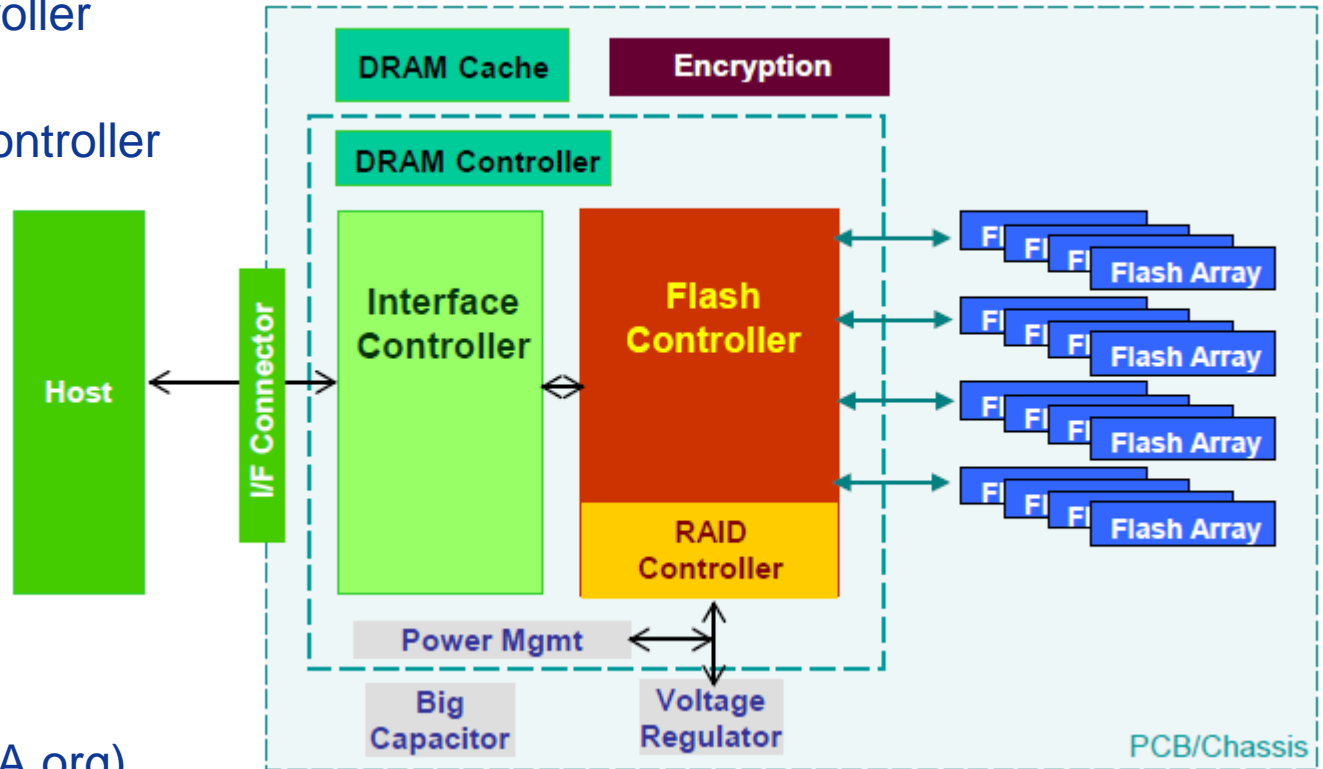


Systems-of-Systems



SSD Architecture Overview

- Others:
Flash Controller
- MLE:
Interface Controller



(Courtesy: SNIA.org)

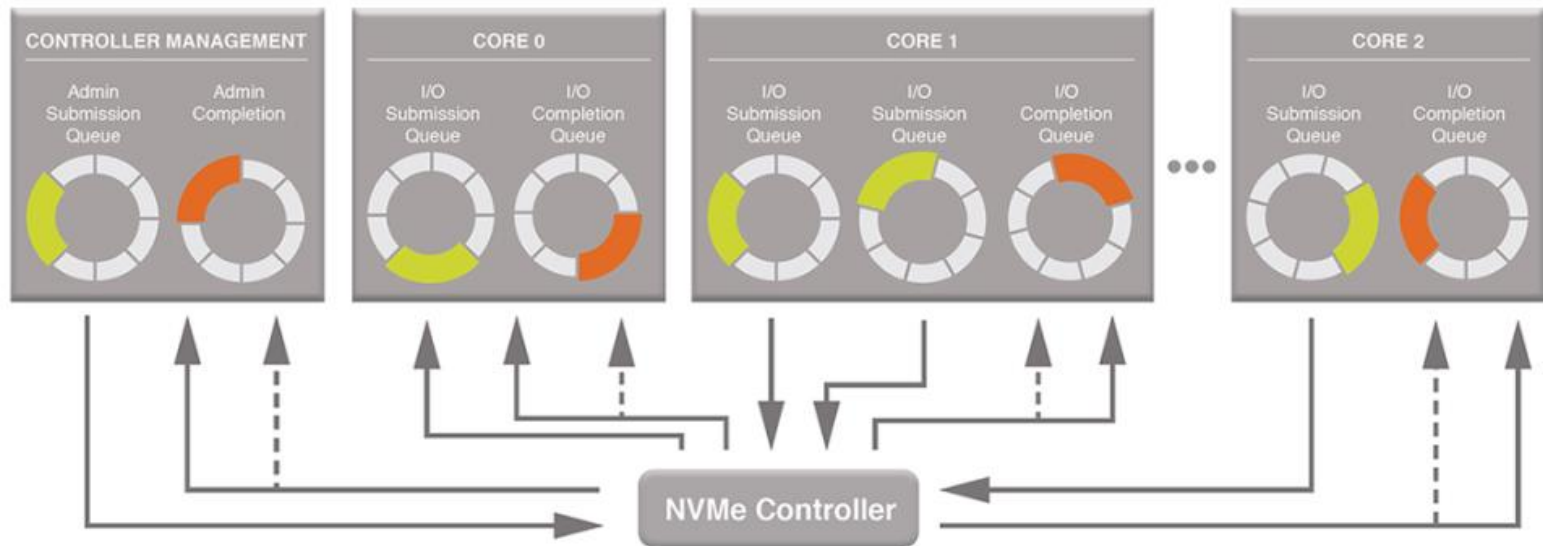
Flash-on-Ethernet

- Fast Network i/f
 - 10GbE / 40GbE
- Dependable Protocol
 - iSCSI
- Future-proof SSD i/f
 - NVMeexpress



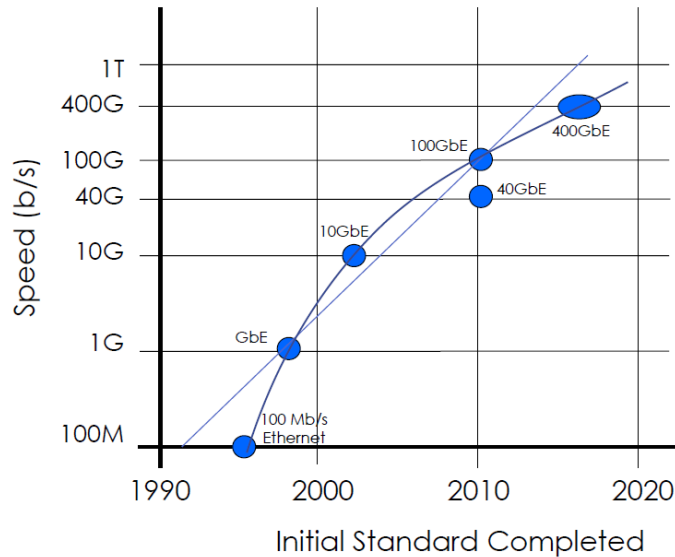
Benefits of NVMeexpress

- Built for PCIe and Flash
- Multi-Queue Facilitates Acceleration



Benefits of Ethernet

Ethernet Speeds – Log



© 2014 Ethernet Alliance


(Courtesy: Ethernet Alliance)



Distance vs Speed

Ethernet operates at different speeds over different distances depending on the media :

- backplanes up to 1m
- Twinax to 15m
- Twisted pair to 100m
- Multimode fiber to 5km
- Single-mode fiber to 40km

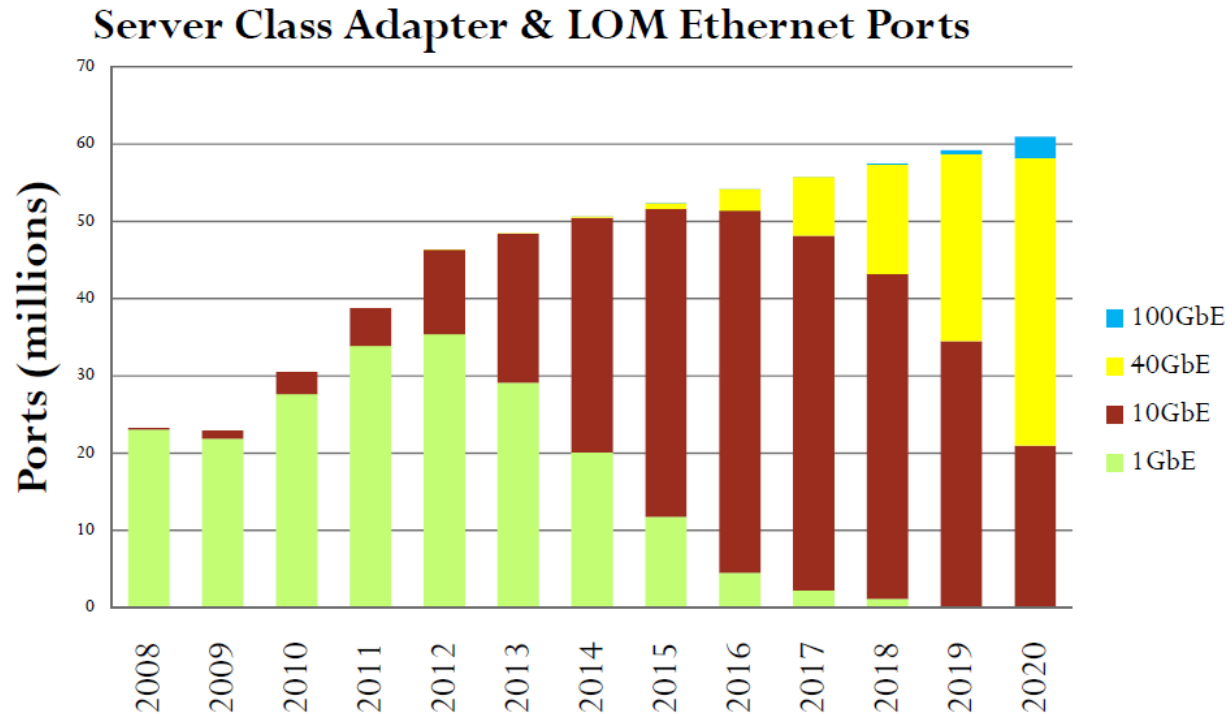


© 2014 Ethernet Alliance

5

Benefits of Ethernet

Add Some History and Map it to Port Volume



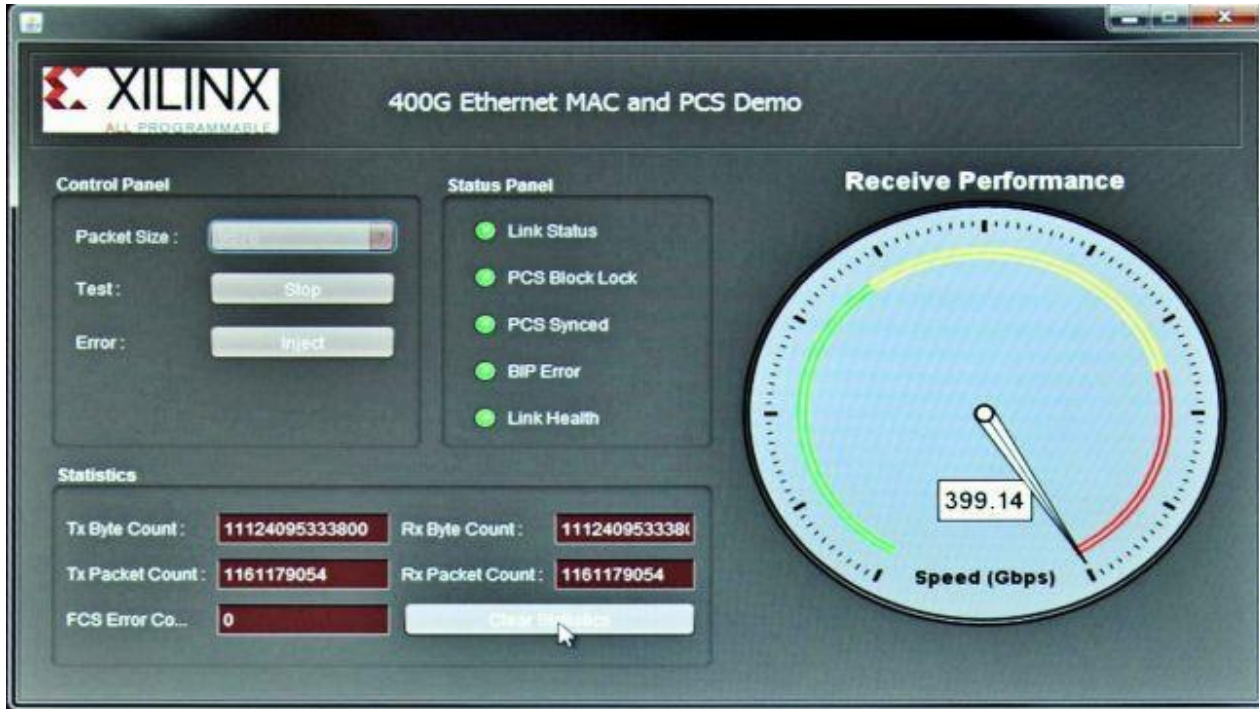
Source data: Crehan Research, 2012

15

IEEE 802.3 Higher Speed Ethernet Consensus Ad Hoc

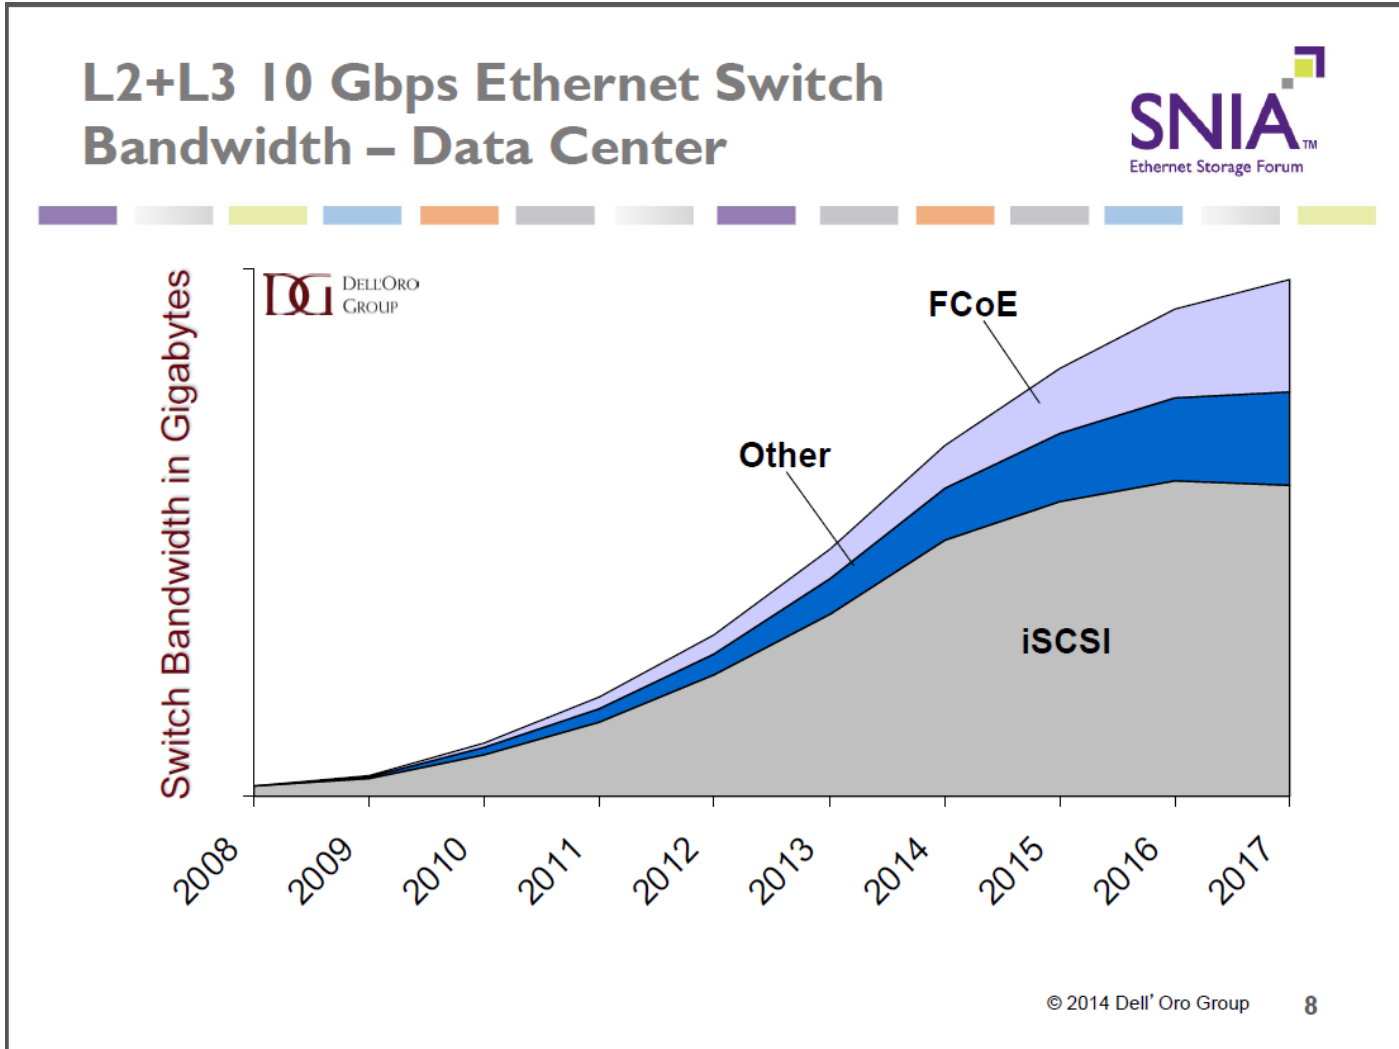
September 2012

400 Gigabit/s Ethernet



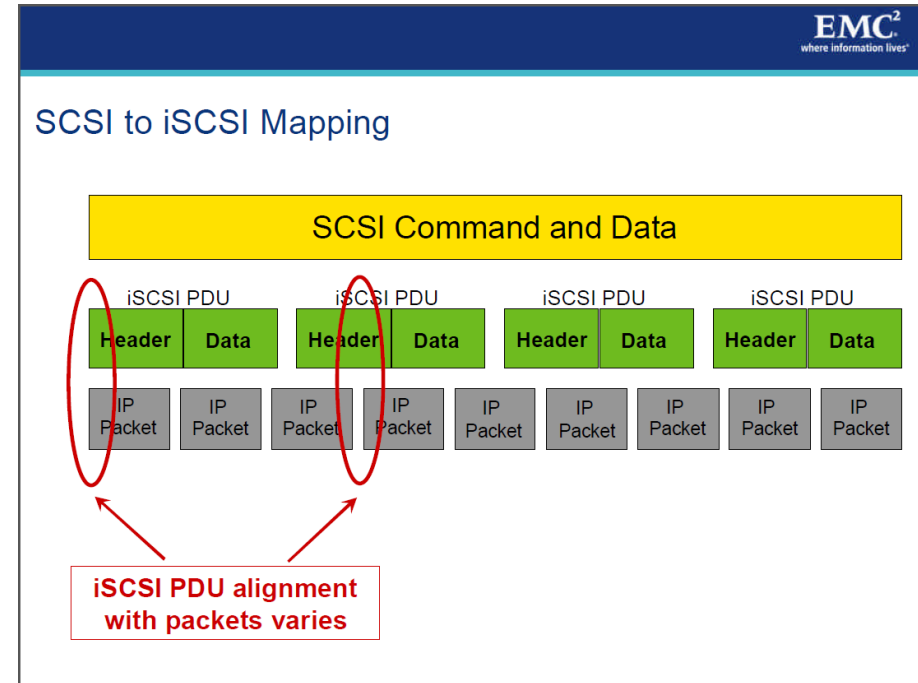
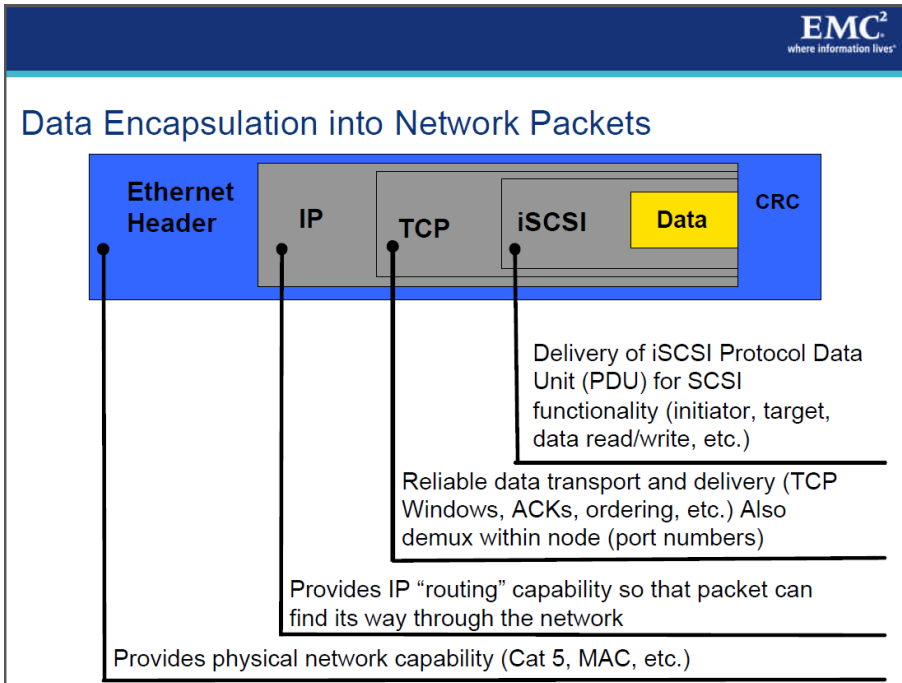
WDM & Next Generation Optical Networking 2014 conference in Nice, France (courtesy Xilinx)

Benefits of iSCSI



Challenges with iSCSI

Computational complexity, Unpredictable latency

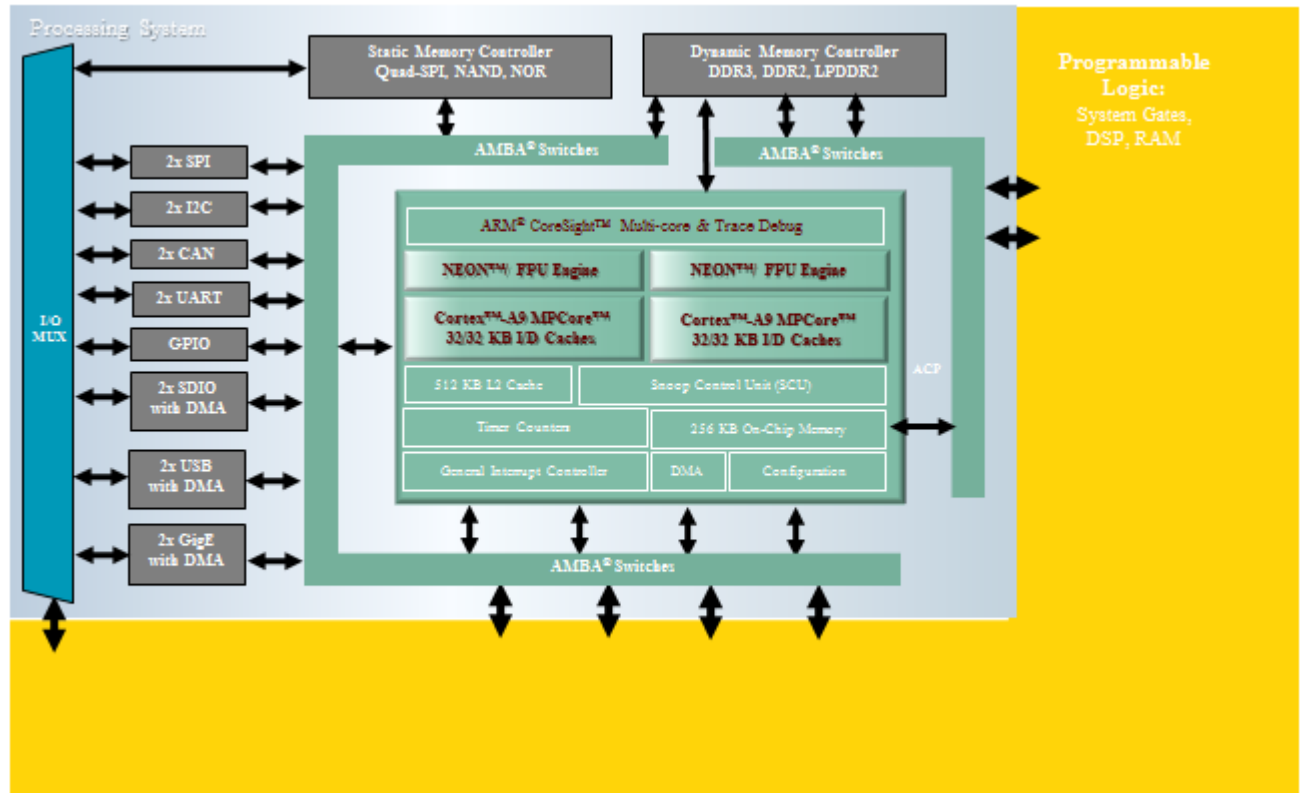


Hardware Acceleration Enabled by Modern All-Programmable SoC

Programmable I/O

Programmable Software

Programmable Logic

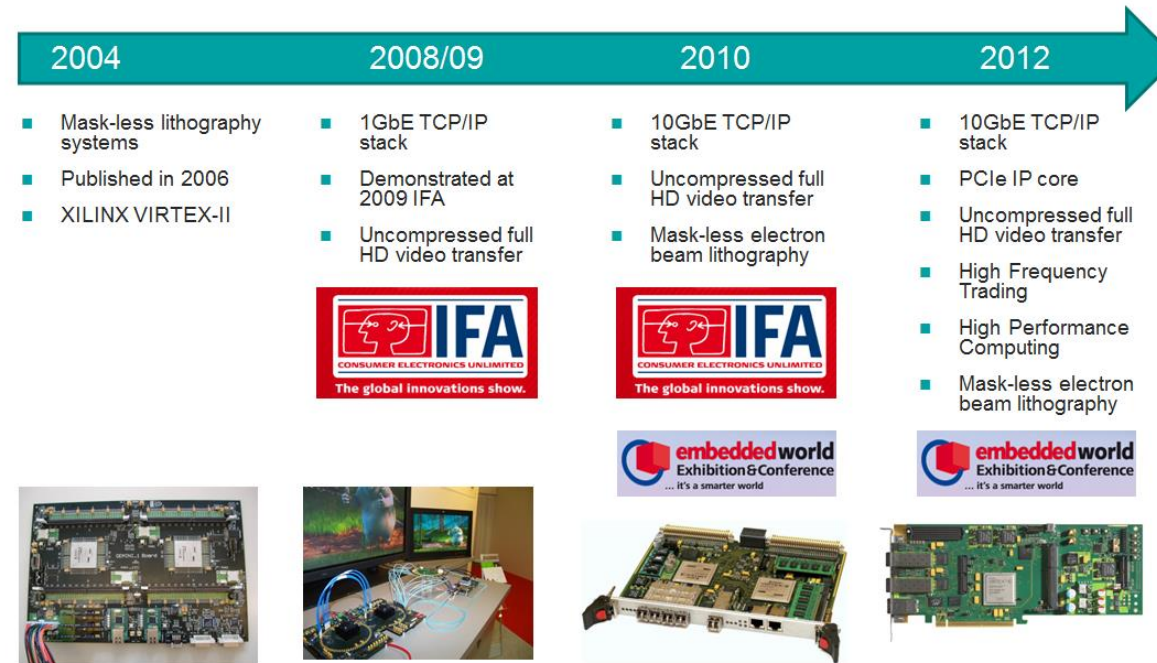


“Put the processing burden where it belongs!”

Network Protocol Acceleration Technology from FhG HHI

Hardware Accelerated Internet Protocol

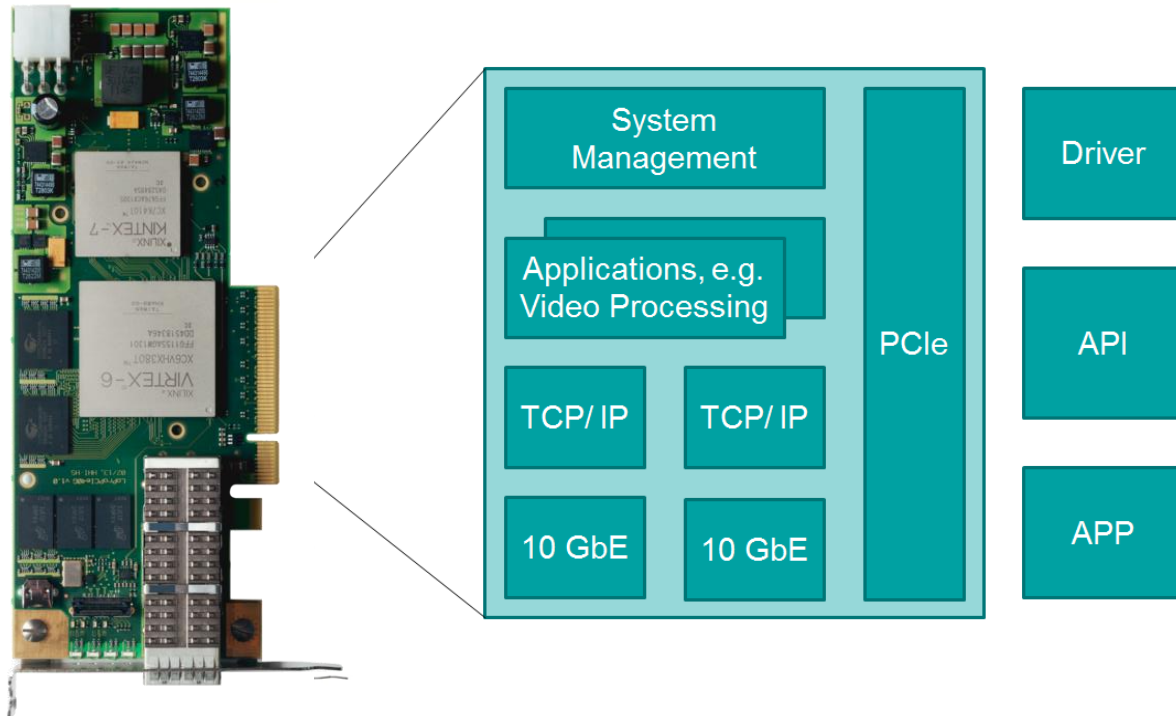
High Speed Hardware Architectures



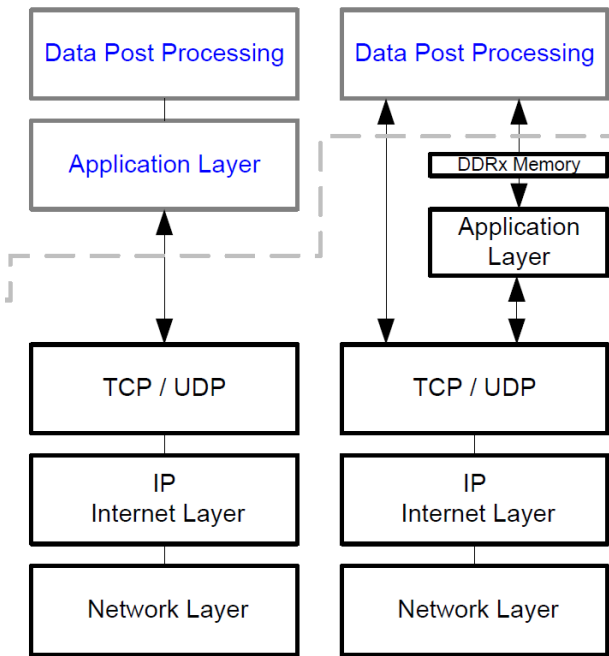
Network Protocol Acceleration Technology from FhG HHI

Hardware Accelerated
Internet Protocol

High Speed
Hardware Architectures



Network Protocol Acceleration More Than TOE

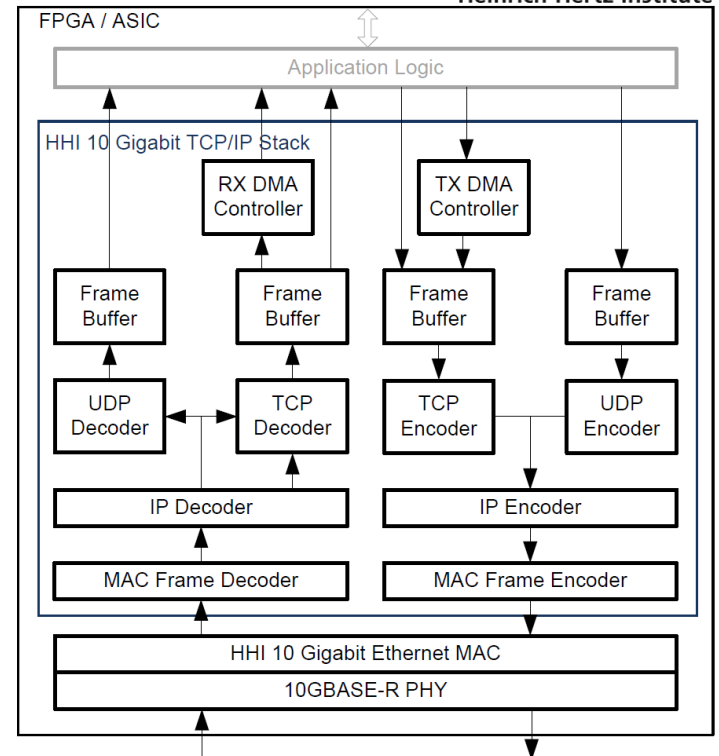


Fraunhofer HHI

- Entire TCP / UDP protocol processing inside FPGA
- Option to run Application Layer processing in HW, too!

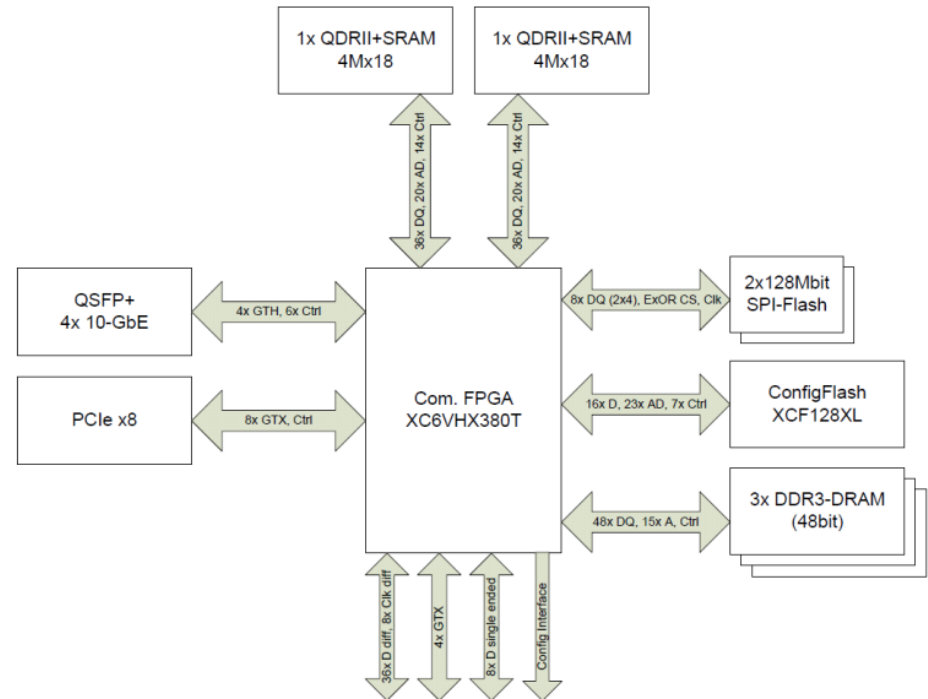
State-of-the-Art

- Software-only
- TCP Offload Engine (TOE)
- requires CPU



Network Protocol Acceleration Processing Devices

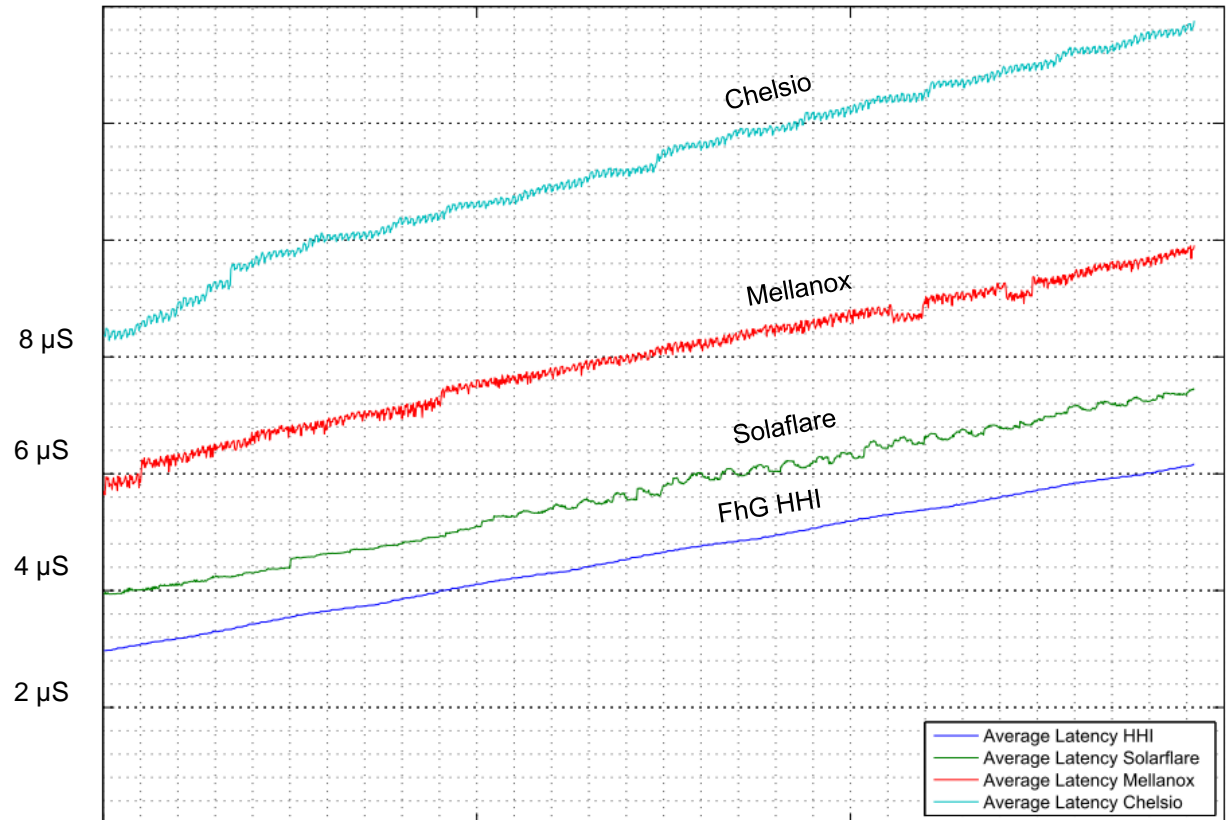
- Built on top of leading FPGA technology, MLE offers a complete integrated device solution encapsulating Network Protocol
- Acceleration for 1GbE, 10GbE, and beyond, plus Application Specific Protocol handling in hardware and/or software.
- Example FPGA technology
 - Altera: Stratix-5, Arria-10, Stratix-10
 - Xilinx: Virtex-6, Virtex-7, Kintex-7, Zynq-7, Kintex-Ultrascale, Virtex-Ultrascale



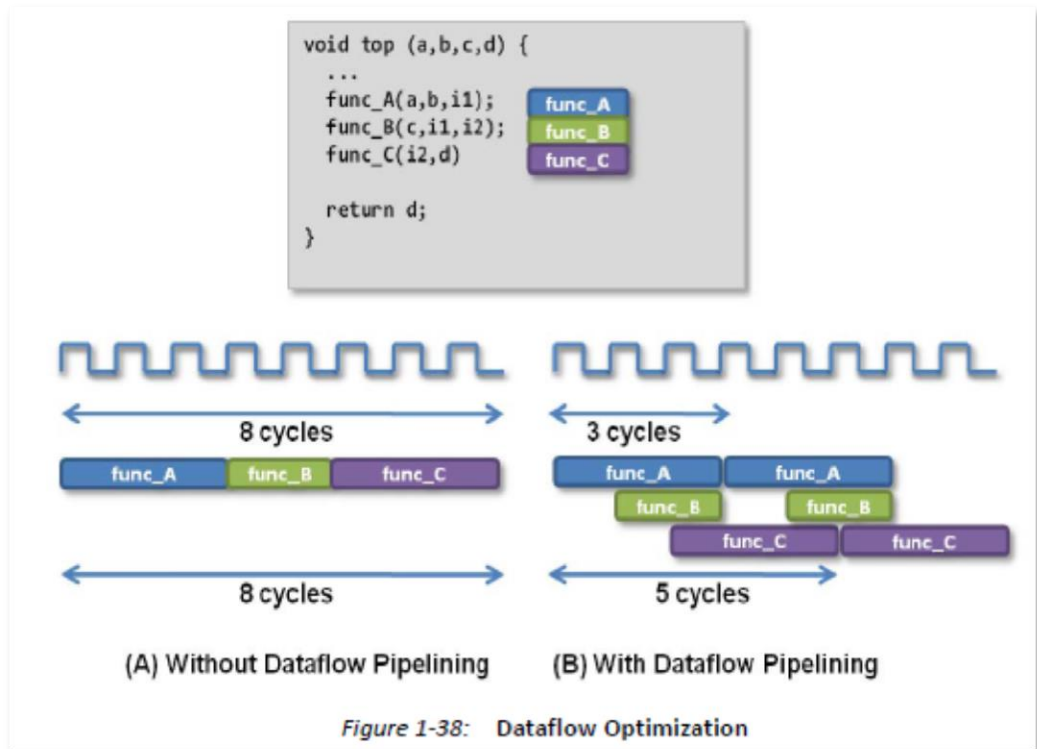
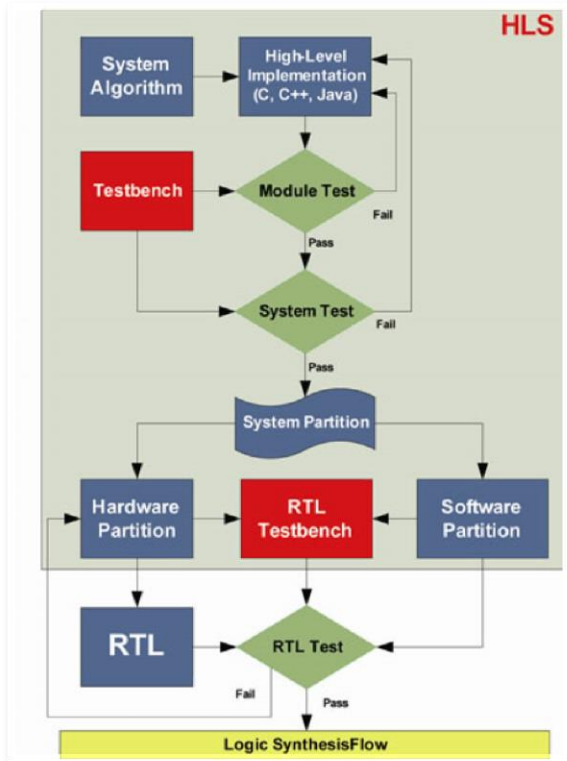
Network Protocol Acceleration Best in Class Performance

- Stand-alone TCP/IP & UDP/IP stack
- Point-to-point 1GbE or 10GbE
- Full line rate of TPRmax = 9.5896 Gbps
- TCP R/W latency of $TTR(W) \geq 1.4 \mu\text{s}$
- UDP R/W latency of $TUR(W) \geq 0.75 \mu\text{s}$
- Round trip time of $RTT_{\text{min}} \geq 2.25 \mu\text{s}$

(2013 benchmarking data from Fraunhofer HHI)



FPGA Implementation via High-Level Synthesis



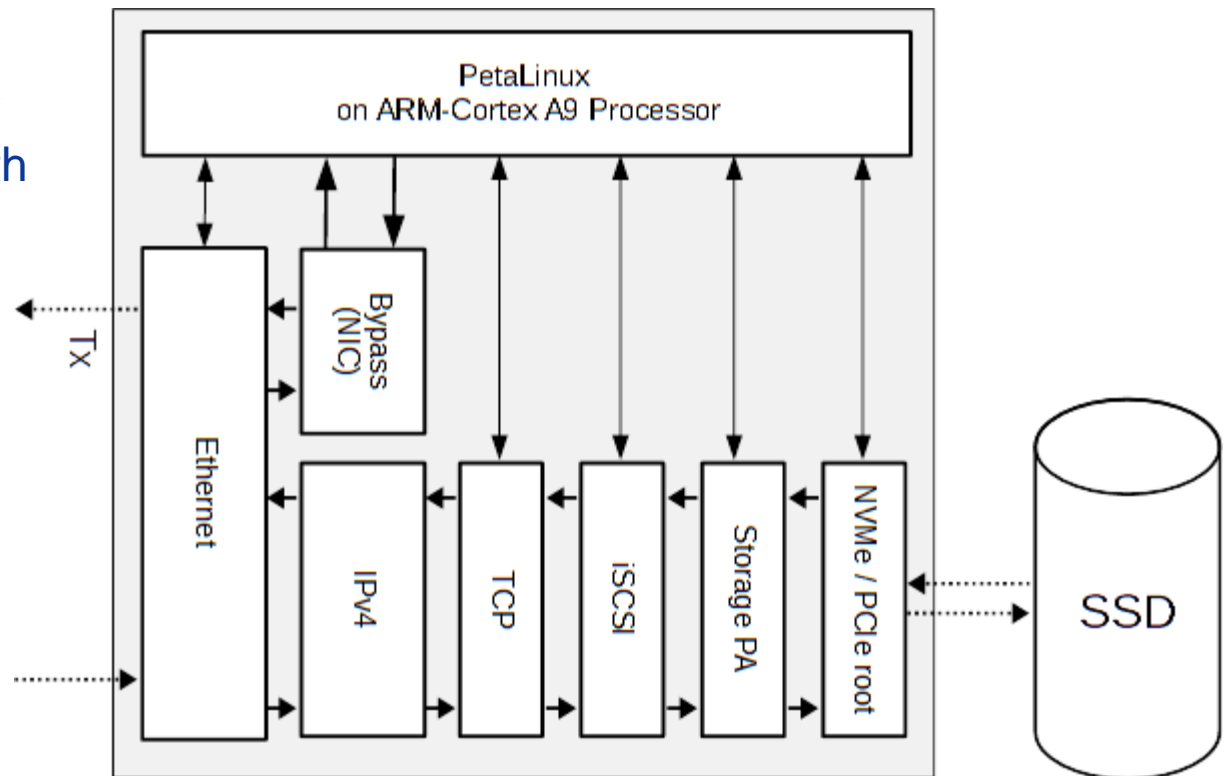
- Xilinx UG902 – Vivado Design Suite User Guide, High-Level Synthesis
- Xilinx XAPP1209 – Designing Protocol Processing Systems with Vivado HLS

Flash-on-Ethernet Architecture

Configurable, elastic system


Balance data rates for Latency and Bandwidth

- SSD
 - 200k IOPS
 - 800 MB/s
 - PCIe Gen2 x2
 - 10GbE
- SSD
 - 800k IOPS
 - 3 GB/s
 - PCIe Gen3 x4
 - 40 GbE



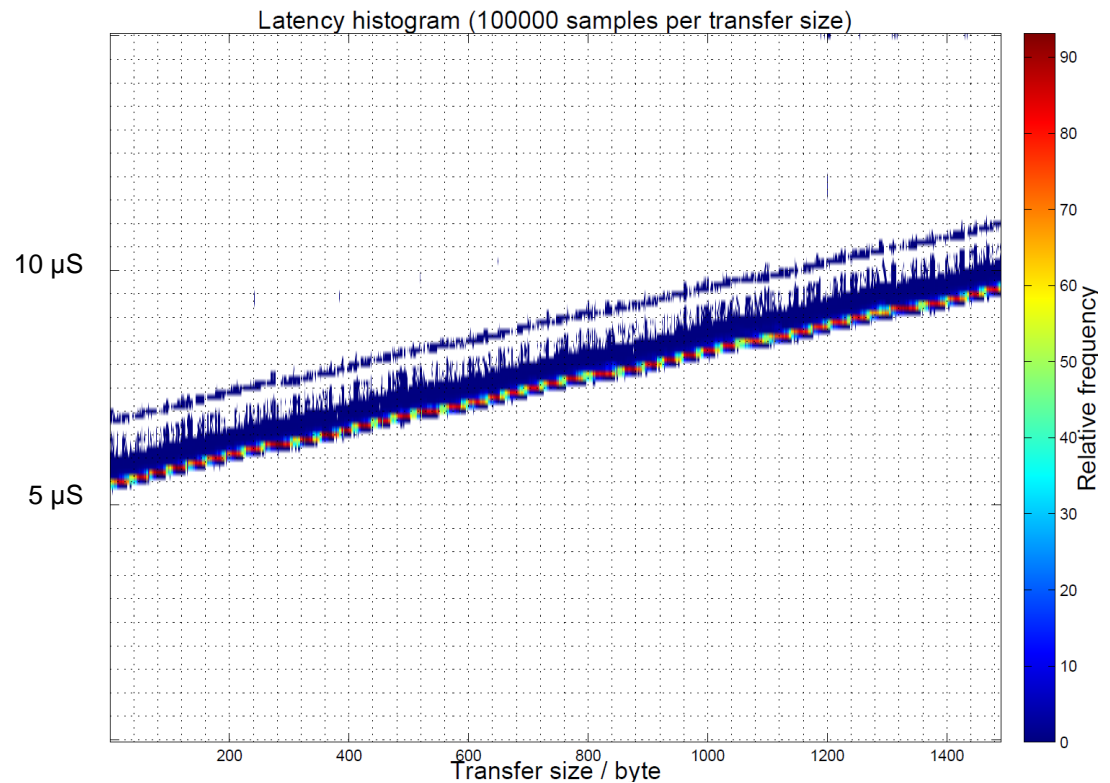
Flash-on-Ethernet Lab Setup at MLE

- Avnet Mini-ITX
- Xilinx Zynq 7045
- PetaLinux
- AHCI SSD via PCIe
- NPAP
- 10GbE




Preliminary Results

- Good determinism
- Reasonable Latency



- Embrace Faster Ethernet: 25GbE, 40GbE



(Courtesy:
Brad Booth,
Microsoft,
25G Ethernet CFI)

- [1] E. P. Markatos, “Speeding up tcp/ip: Faster processors are not enough,” in In 21st IEEE International Performance, Computing, and Communication Conference. IEEE, 2002, pp. 341–345.
- [2] P. E. McKenney, Is Parallel Programming Hard, And, If So, What Can You Do About It? Corvallis, OR, USA: kernel.org, 2011, available: <http://kernel.org/pub/linux/kernel/people/paulmck/perfbook/perfbook.html>.
- [3] Z.-Z. Wu and H.-C. Chen, “Design and implementation of tcp/ip offload engine system over gigabit ethernet,” in Computer Communications and Networks, 2006. ICCCN 2006. Proceedings.15th International Conference on, 2006, pp. 245–250.
- [4] G. Regnier, D. Minturn, G. McAlpine, V. Saletore, and A. Foong, “Eta: experience with an intel reg; xeon trade; processor as a packet processing engine,” in High Performance Interconnects, 2003. Proceedings. 11th Symposium on, 2003, pp. 76–82.
- [5] R. Braden, “Requirements for Internet Hosts - Communication Layers,” RFC 1122 (INTERNET STANDARD), Internet Engineering Task Force, Oct. 1989, updated by RFCs 1349, 4379, 5884, 6093, 6298, 6633, 6864. [Online]. Available: <http://www.ietf.org/rfc/rfc1122.txt>
- [6] A. Salman, M. Rogawski, and J. Kaps, “Efficient hardware accelerator for ipsec based on partial reconfiguration on xilinx fpgas,” in Reconfigurable Computing and FPGAs (ReConFig), 2011 International Conference on, 2011, pp. 242–248.
- [7] U. Langenbach, A. Berthe, B. Traskov, S. Weide, K. Hofmann, P. Gregorius, "A 10 GbE TCP/IP Hardware Stack as part of a Protocol Acceleration Platform", 2013 IEEE 3rd International Conference on Consumer Electronics
- [8] E. Schubert, D. Boggs, P. Gregorius, S. Voss, U. Langenbach, “Low-Latency Networking for Systems-of-Systems”, 2014 Embedded World Conference



Contact Info

Missing Link Electronics
www.missinglinkelectronics.com

Endric Schubert
endric@MLEcorp.com
Phone US: +1 (408) 320-6139
Phone DE: +49 (731) 141149-66

