Swift40e™

Smart Data on the Move

Dr. Bodo Parady
Chief Technology Officer
Pentum Group, Inc.

Swift40e

Smart Data on the Move

Introduction

Pentum Group Inc. (PGI) is proud to introduce Swift40e, the world's first 40 Gbps (Giga bits per second) IP-based end-to-end system interface card. Swift40e incorporates standard communication protocols and connects directly to the OC-768c optical WAN infrastructure via PCI Express® Technology (PCIe). This allows quick deployment using today's communications infrastructure and communications switches to support wide-area network applications from across the campus to around the world. In addition to world-class networking performance, the Swift40e family provides a unique capability for customized data stream pre- and post-processing and can be configured to provide over 15 TeraOps/sec for custom applications.

With the continued advances in high-performance computing, data storage, and the growing volumes of data generated through increasingly sophisticated networks of sensors, demands on wide-area network bandwidth are growing dramatically.

Success in the battlespace of today and tomorrow depends upon data collected over widely deployed complex sensor systems. These large volumes of data must be rapidly collected, assessed and integrated to ensure efficient and effective application of resources to achieve military objectives.

In cyberspace, security threats are also rapidly evolving and demand rapid acquisition and integration of very large volumes of data to detect and defend from cyber attack.

The ongoing advances in network, computers and data storage technologies create new opportunities to advance business efficiencies through more effective product development and business process integration. Increases in video fidelity create new opportunities for widely distributed teams to efficiently collaborate on new products and services. Large remote data stores are increasingly core components of the enterprise for both day-to-day operational efficiency and disaster recovery.

All of these new technologies are driving demand for ever greater network bandwidth.

40 Gbps: The Next Step in WAN Performance

To advance the state of the art in wide-area networking (WAN) to the next level requires a choice of transport and framing protocols that can be supported in today's deployed telecommunications networks. The next increment in end-to-end, single-stream optical transport over WAN is 40 Gbps. Today, 40 Gbps is primarily used for data aggregate transport over long haul and is widely deployed by telecommunications providers. There are many hardware equipment manufacturers such as Cisco Systems, Juniper Networks, Nortel, Ciena, Alcatel-Lucent and others that provide switches, routers, network interfaces and other network components to support this application. To support standard IP transport protocols for 40 Gbps, POS (IP Packet Over SONET) provides a standards-based specification that can be universally carried across these deployed long haul backbones. For these reasons, Swift40e implements OC-768c POS IPv4 and IPv6 POS.

To support the widest range of today's computers requires high bandwidth and low latency standard based I/O interfaces. At this time the two best candidates for this demanding level of bandwidth are HyperTransport™ and PCI Express. The Swift40e development targeted PCI Express since it is the industry-standard and widely deployed.

By selecting these popular standards, the Swift40e enables the rapid and flexible deployment for tomorrow’s demanding data-intensive applications.

Swift40e Architecture

The basic system architecture of Swift40e in the standard configuration is illustrated in Figure 1. The core of the hardware logic is located on a motherboard that uses the Extended ATX form factor. The motherboard is housed in a standard rack-mountable 1U chassis. The Swift40e motherboard is connected to the host system with up to 2 PCIe adapter cards (X10D) with cables that can be up to 6 meters in length. The physical hardware is shown in Figure 2.

A key design objective of Swift40e was to support full OC-768c duplex bandwidth. Since PCIe was the targeted I/O platform interface, and COTS systems are available that support up to 16 lanes of PCI Express, Swift40e, in the standard configuration, incorporates dual PCIe X16 that, together, provide aggregate I/O bandwidth to support full bi-directional 40 Gbps data transport.

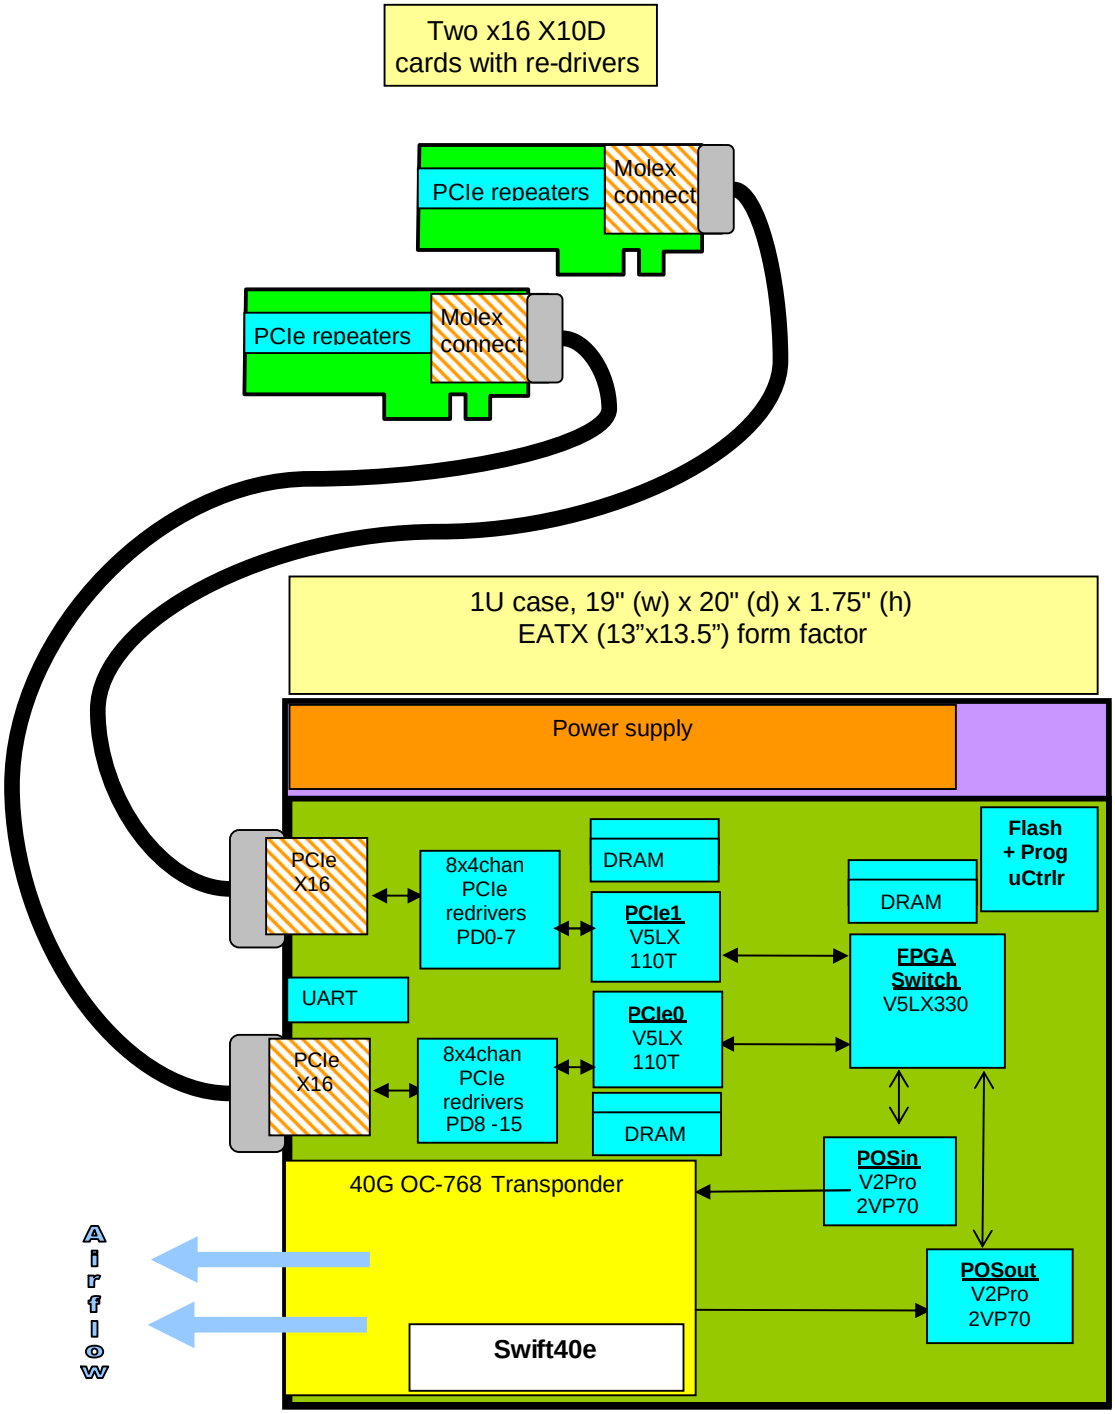

Figure 1. Swift40e Block Diagram Showing Data Flows

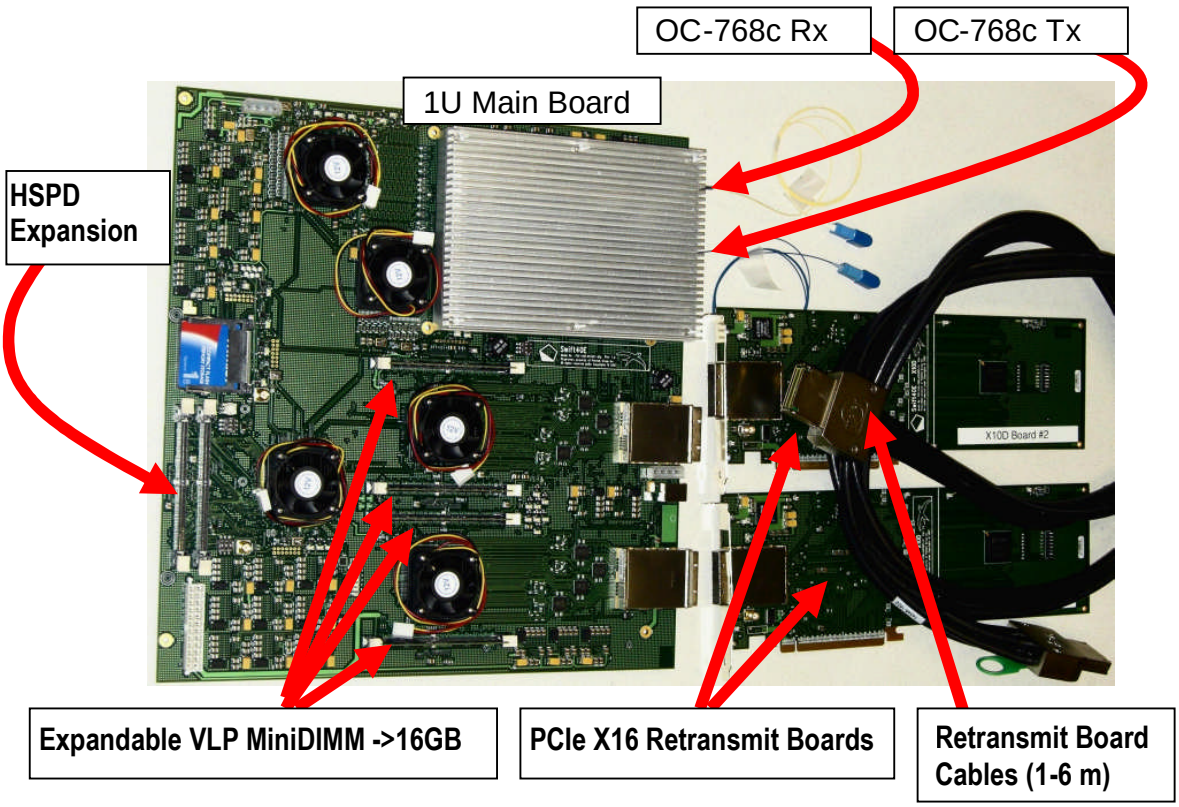

Figure 2. Swift40e Hardware Components

To support the integrated operation of the dual channels, PGI developed a scalable protocol to control 1, 2 or even more PCIe channels, seamlessly, and at full bandwidth. Inputs from each of the PCIe channels are processed through dedicated Xilinx Virtex5 Field Programmable Gate Arrays (FPGA) at PCIe0 and PCIe1, and are integrated together at the FPGA Switch. A number of load balancing options can be selected to ensure maximum system throughput and performance depending on the I/O requirements of a user's application workload.

The FPGA Switch separates the integrated data stream into distinct ingress or egress data streams for electro-optical conversion through the devices POSin and POSout, respectively. These, in turn, communicate with an optical transceiver through standard OIF-SFI-5 connecting via a 300 pin MSA

connector. FPGAs POSin and POSout FPGAs implement all SONET and POS protocol handling.

All communications between the FPGAs are based on a proprietary high performance communication protocol developed by PGI called HSPD™. The benefits of HSPD are both high bandwidth and low latency parallel communications instantiated in a lightweight core with a very small footprint on the FPGA. This leaves many available resources for additional processing tasks.

FPGAs PCIe0 and PCIe1 have associated attached DRAM that can be configured with a total on-board memory between 4-16 GB. The FPGA Switch can also be optionally configured with up to 4 GB of attached DRAM.

The DRAM connector on FPGA Switch at the core of the Swift40e board has the capability to interface to user-provided mezzanine cards. Through the use of the HSPD protocol, such a card can extend the computational and analytic capabilities of Swift40e for streaming data.

Advanced Processing on the Fly

Stream Processing

A key principle of the Swift40e design was the incorporation of Xilinx Virtex5 FPGAs for key processing tasks. The advantage of this approach is that Xilinx in the Virtex5 family offers a broad variety of processing capabilities and performance grades. This creates the opportunity for the Swift40e to host customized applications for processing data in flight on both ingress or egress in real-time.

The most demanding real-time data flow applications require coherent processing of the data stream sequence. As an example, many signal and image acquisition systems use tools such as FFT, FIR filters, or IIR filters to produce processed data. The inherent nature of these filters requires that the data be presented and processed in its natural order. Techniques such as splitting streams down multiple network channels often compromise the ability to apply these core filtering techniques. As an example, an FFT requires the fully assembled input vector before initiating butterfly operations. Similarly, the synchronization costs and/or the scale of parallel programming costs often make this class of problems extremely difficult to solve. These kinds of problems need two key technology features to support a solution; *very fast single stream network transport* coupled closely with *very high bandwidth streaming data processing resources.*

In the area of cyber security, examples of this kind of demanding streaming applications are high bandwidth data encryption/decryption, network monitoring, and networking auditing.

Another example is the movie industry's evolving plan to replace film with all-digital content and use electronic transport for all production tasks from basic post-production through final theater distribution.

By coupling the fastest single network interface with a powerful streaming data processing capacity, the Swift40e enables a quantum leap to support the most demanding data flow analysis and processing.

Swift40e FPGA TeraOps/sec Processing Power

Swift40e offers several options for increasing stream processing capacity. The base Swift40e board has two Xilinx Virtex 5 LX110T FPGAs PCIe0 and PCIe1 that can be upgraded through either FPGA model and/or speed grade.

Both of the LX110T FPGAs can be upgraded to the LX220T, or LX330T. As shown in Table 1, this adds significantly to the user's capacity to program custom applications. In Table 1, it is assumed that 20 percent of the LUTs on the LX110Ts and 70 percent of the LX330 currently on the board are available for user application development.

Table 1. Swift40e FPGA configuration options

Virtex 5	LUTS	User LUTS	User Top/s	DSP Slices	DSP GMACs
LX330 (Standard)	331,776	232,243	5.8	192	105
LX110T (Standard)	110,592	22,118	0.6	64	35
LX220T	221,184	110,592	2.8	128	70
LX330T	331,776	199,066	5.0	192	105
Total(Standard)		276,479	6.9	320	175
Total Max		630,375	15.8	576	315

To calculate the number of teraoperations per second (Top/s), we use the rule of thumb that it takes 20 LUTs to perform a simple bit-wise operation at a 500 MHz clock. Overall, that means that Swift40e as it is shipped can perform over 5 Top/s, and it can be upgraded to support over 15 Top/s.

To support advanced signal and image processing, the Xilinx Virtex 5 incorporates DSP slices to accelerate multiply-accumulate operations, otherwise known as a MAC function. These DSP slices are found in both the LX330 and the LX110Ts, standard components on the Swift40e. As shipped the Swift40e has a total of 320 DSP slices which, based on estimates

published by Xilinx (352 GMACS for 640 DSP slices[1] where GMACS are billions of executed MACs per second), yields 175 GMACS at 550 MHz. Fully expanded, the Swift40e can reach 315 GMACS, making the Swift40e a powerful platform for signal and image processing.

The Swift40e can also be optionally configured to support a user-supplied expansion card that interacts with FPGA Switch via HSPD communications. This interface would support 32 Gbps full duplex communications (total 64 Gbps of bandwidth).

Combining the elements of powerful FPGA-based processing, high-speed chip-to-chip HSPD interconnects, and leading edge 40 Gbps network interface enables Swift40e to support a whole new class of applications targeting high-speed coherent datastream analysis. With a full speed standard-based interface into computer systems, the Swift 40e can be seamlessly integrated into today's advanced HPC architectures. As an example, the SGI® Altix® family is architected on a modular approach tied into a unique high bandwidth, low latency computer memory fabric. Altix customers can mix and match Intel® Architecture microprocessors, special purpose FPGA processor accelerators (RASC), and a broad array of high speed disk options that, when included with the Swift40e, offer extraordinary opportunities to work with large volumes of data that must be acquired over long distances and processed at speed to enable real-time response.

Summary

With the Swift40e, Pentum Group, Inc. presents the first commercially available WAN support at 40 Gbps. Swift40e advances the state-of-the-art in high-performance WAN with a standards-based architecture that maps well with today’s deployed telecommunications infrastructure. With this, users can now develop and effectively deploy next-generation data intensive applications.

By including powerful and configurable processing options, the Swift40e enables new classes of network-centric applications based on a stream processing paradigm. Applications in advanced weapons systems, biology and medicine, scientific research, media creation and delivery, advanced security (both physical and cyberspace) and high-speed remote data access will all benefit from the Swift40e design.

[1] “XtremeDSP Selection Guide”, Xilinx, http://www.xilinx.com/publications/prod_mktg/pn0010944-3.pdf

About Pentum Group, Inc.

Pentum Group, Inc. provides core technology and senior technologists to help clearly define the impact of technology and its ability to improve and accelerate critical organization missions.

Dr. Bodo Parady (CTO of Pentum Group, Inc.) is a 25-year veteran in computer system technology development. He currently is Chief Technology Officer for Pentum Group, Inc. and leads development of the Swift40e.

Trademark Notices

All other product names are for informational purposes only and may be trademarks or registered trademarks of their respective companies.