

# Co-Development of Test Electronics and PCI Express Interface for a Multi-Gbps Optical Switching Network

C.E. Gray<sup>1</sup>, O. Liboiron-Ladouceur<sup>2</sup>, D.C. Keezer<sup>1</sup>, K. Bergman<sup>2</sup>  
Georgia Institute of Technology<sup>1</sup>, Atlanta, GA  
Columbia University<sup>2</sup>, New York City, NY

## Abstract

*This paper presents the design and performance characteristics of a system designed to interface between a PCI Express port and an optical packet switched network as well as provide inline test capability for the whole system. A single lane of PCI Express traffic is inverse multiplexed across eight parallel channels and retransmitted in a burst packet at aggregate data rates of 20 to 36 Gbps. Testing options include loopback self-test, data synthesis and substitution in-line with or in place of system data, and variable channel-to-channel skew. The I/O interfaces also support a range of variable analog parameters such as peak output amplitude, output amplitude swing, and common mode ranges for the input and output to evaluate the performance of or adapt to changes in the opto-electronic components. This design flexibility also allows for use of the system in more conventional electronic applications with little or no required modifications.*

## 1. Introduction

The latest trend for supercomputing systems has been for very large-scale shared-memory systems that can be constructed from commercial off the shelf processors. Such parallel systems, consisting of hundreds or thousands of processors, are most often bottlenecked with long latencies and delays on inter-processor or processor to memory messages [1]. Recently, the widespread introduction of high-bandwidth serial interfaces such as PCI Express and HyperTransport has opened up new design opportunities to overcome these limitations, though these technologies are more appropriate to chip-to-chip or board-to-board interconnections which limit the scalability of these implementations.

A common box-to-box interconnection approach is to implement the systems independently and connect them using various network or switching topologies. The interconnection overhead and latency from conventional networking of the individual elements limits the performance benefits and scalability of the overall system. Emerging all optical packet switching networks are a potential solution to this problem because of their fundamental differences from conventional networks. All optical networks reduce end-to-end propagation latency as a result of minimizing or eliminating the various optoelectronic conversions in a mixed format system and

eliminate the buffering necessary in electrical systems. Delays can be further reduced by leveraging dense wavelength division multiplexing (DWDM) to transmit an otherwise long sequence of data in a much shorter burst across a number of parallel wavelengths, decreasing the time required to transmit the data while also increasing the overall transmission capacity [2].

However, there are some fundamental challenges to interfacing a high-speed off-chip interconnection bus like PCI Express to an optical packet switched network, which is the focus of an ongoing joint project between research groups at the Georgia Institute of Technology and Columbia University. The co-developed test electronics and interfaces are intended to serve the dual purpose of system solution for transporting data originating from a PCI Express port across an optical packet switching network as well as system level test and verification. Many of the design features present in the test electronics provide inline test capabilities necessary to evaluate and characterize the system as a whole which is difficult or impossible using other test systems.

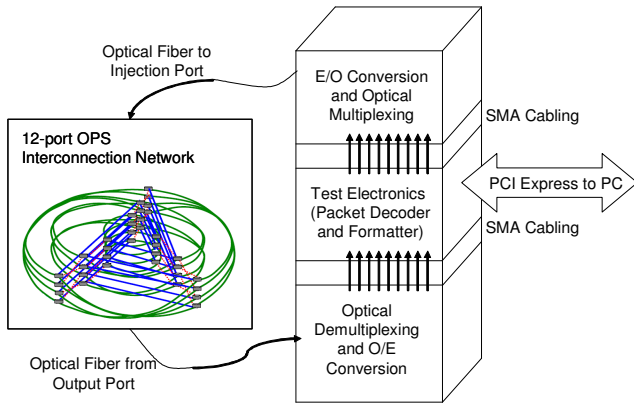
The current overall system is designed to convert a single lane of PCI Express traffic into a collection of eight parallel wavelengths at 2.5-4.0 Gbps each, as well as the corresponding Clock, Frame, and Routing signals contained by the packet. This results in an aggregate data rate of 20-36 Gbps, but additional payload signals could be handled on parallel wavelengths to extend the system capacity.

## 2. System Overview

The optical packet switching (OPS) network being utilized for this project is based on the Data Vortex topology [3] and is being developed at Columbia University. The network consists of concentric cylinders, each made up of circulating optical fibers and periodic switching nodes. The nodes are designed to respond to particular wavelengths which carry the routing information bits within the packet. Depending on the logic value of the routing bit, the data can be directed to a different cylinder or continue circulating within the present cylinder. Optical packets are injected at input nodes on the outer cylinder, and eventually exit at destination nodes on the inner-most cylinder.

The network is designed to support the transmission of very short packets with low latency so PCI Express was selected as a complimentary technology, providing high

bandwidth and low latency access to a system processor. A high level system overview is shown in **Figure 1**.



**Figure 1.** System topology

Since the respective systems are constantly under further development, the majority of the opto-electronic interfaces are modular, utilizing differential PECL signals over flexible cables and SMA connectors. The test electronics are designed to serve as a bridge to a personal computer using a standard card edge connector to a single lane PCI Express slot, though the onboard FPGA is capable of fulfilling this role using synthetic data in the absence of this connection. The test electronics contain the components necessary to structure and transmit outgoing network packets as well as receive and process the incoming packets. Careful attention has been made to the selection of the electrical to optical (O/E) and optical to electrical (E/O) components to support the burst nature of the signals.

While **Figure 1** shows the system in a loopback configuration, a single transaction across the OPS network is one way as shown in **Figure 2**. A transaction originating at a personal computer is received by the test electronics, formatted into an aligned packet, augmented with the network control signals, converted to optical wavelengths and injected into a network input port. The packet circulates to the appropriate destination port where the Frame, Payload, and Clock signals are converted back to electrical form, deskewed (to partially compensate for any chromatic dispersion that may have occurred within


the network or part variation between channels), processed by the test electronics and transmitted to the destination personal computer (or similar system, such as a high-speed memory node).

If a response is required, such as a memory read request, a similar but independent transaction would occur with the destination and source nodes reversed. For testing purposes, the same board is often used as the source and destination but such a configuration would be unlikely in the target application. The test system has many built-in features, controlled through a standard USB port, that allow for the adjustment of performance parameters (such as channel-to-channel skew) for transmission and to analyze the patterns and relative timing of the returned signals from the OPS network. Timing accuracy, in particular, is a major concern since it must be controlled on a picosecond scale to ensure proper alignment and reception of the data across all the payload channels.

### 3. System Components

#### 3.1 Optical Packet Switching Network

The Data Vortex topology, upon which the target network is based, is an all optical packet switching network designed for low latency communications. The Frame and routing signals are used within the OPS network for routing decisions while payload wavelengths are transparently passed through the system without being processed. This architecture eliminates the need for complex optical buffering, and the parallel presence of the routing information eliminates the need to further process the packet payload to decode and evaluate routing decisions. A traditional electrical networking system would require buffering and retransmission of the signal at each routing node, as would older generation optical networking systems (requiring additional conversion of the signal from optical to electrical form and back again). By avoiding these unnecessary buffering and conversion steps, the optical packet switching network achieves considerably reduced latencies compared to store and forward networks. In addition, since there is no conversion, sampling, and buffering of the payload, the performance of the network is independent of the number and signaling rate of the payload channels. This permits



**Figure 2.** System data flow [12]

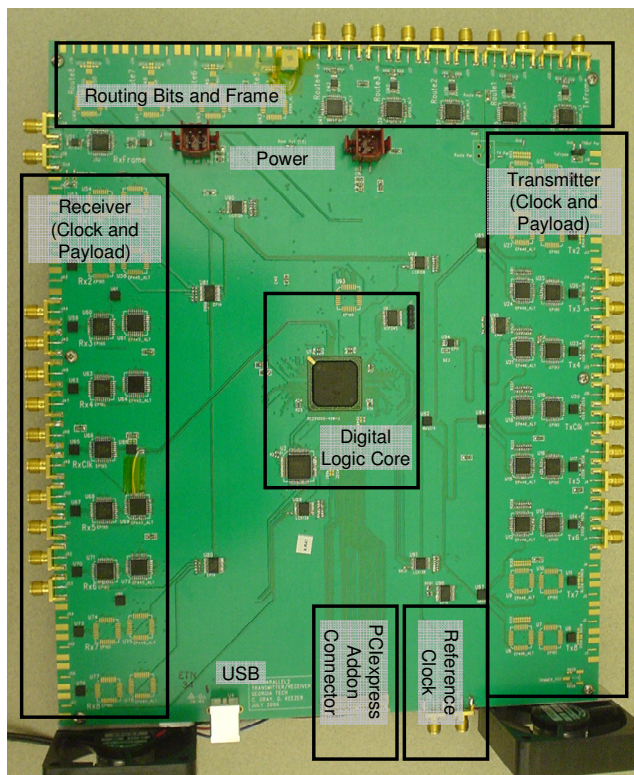
future expansion in both parameters without any necessary changes to the network itself.

### 3.2 Test Electronics

The current system is based on a design presented previously [4] [5] but has been upgraded to double the previous number of high-speed payload channels. The system now supports eight data channels as well as a ninth channel that is used as a parallel transmitted clock for data recovery at the destination. This implementation avoids the complicated clock and data recovery (CDR) circuits present in many other optical packet switching network implementations. Additionally, most CDR solutions would be impractical for the extremely short packets used in this application, requiring 100 or more [6] bit periods to lock to the incoming signal.

The total number of routing control signals was increased from four to eight, though the current network is only large enough to require four bits of routing information. The board also contains numerous design refinements to improve overall performance and flexibility to support the testability of the OPS network. This flexibility is also required to adapt to the ever evolving physical interface of the OPS network as newer and better performance optical components are evaluated and integrated into the design.

**Figure 3** shows a photograph of the custom test electronics with the prominent features labeled. The central element is a CMOS based FPGA, flash



**Figure 3.** Photograph of the test electronics

configuration PROM, and USB interface collectively referred to as a digital logic core (DLC) [7] [8]. This DLC, in conjunction with an add-on card being developed and is featured in the next section, allows for the interconnection of the board to a PCI Express expansion port on a personal computer.

The PECL circuits for the transmitting and receiving functions can be seen surrounding the DLC components. The high-speed output data channels, which include the source clock, are shown to the right. These channels operate at 2.5Gbps, serializing an 8-bit parallel bus per channel originating from the FPGA. This method has been used to achieve signal rates up to 4.0Gbps though additional multiplexing using similar techniques have been used to achieve rates as high as 6.4Gbps [9], providing a future upgrade path if necessary. The reference clock used for this serialization process can be found to the bottom right of the board and is distributed to the respective channels using a line matched distribution tree to ensure identical timing across all channels.

The frame and routing signals, operating at a much lower speed and therefore not requiring additional serialization, are present at the top of the board. Input logic for the frame and high-speed data/clock is located to the right. The inverse of the transmission process occurs for the received channels, using the parallel clock to sample and deserialize the data into 8-bit parallel words that can be read by the FPGA. Flexible coaxial cables and SMA connectors are used to interface differential PECL signals to removable electro-optic modules (not shown).


Each channel can be independently skewed in time to assure proper alignment or to characterize the system. Every outgoing and incoming signal, high or low speed, has the capability for the addition of up to 10ns of additional programmable delay in 10ps steps [5]. This delay can be used functionally to align outgoing channels or deskeew incoming signals as compensation for chromatic dispersion or part variation between channels. For testing, the adjustable delay allows for the simulation of various sources of timing variance and characterization of the system performance.

The output channels incorporate a variety of analog controls that allow the signal to be permuted across a range of peak amplitude values and amplitude swings [4]. These adjustments can be used to stimulate the network under non-ideal conditions or a combination of these adjustments can be used to vary the signal bias level. Should this adjustment range be insufficient to match the common mode range between the test electronics and the opto-electronic modules, the input and output buffers can also be isolated from the main board power system and operated from additional power supplies biased above or below the system supplies. Should even these measures be insufficient, external adapter modules can and have been constructed for the specific system needs. These provisions have been made because, due to the burst

nature of the signals, AC coupling cannot be used to match the common mode ranges between the systems without a severe impact on the signal quality.

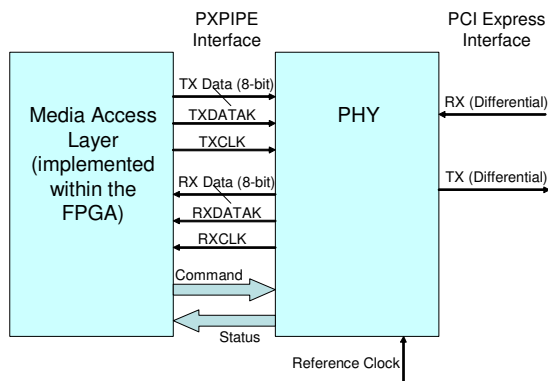
### 3.3 PCI Express Interface Card

A picture of the PCI Express interface card is shown in **Figure 4**. The card fits into a standard 1X PCI Express slot and utilizes a commercial single-lane PCI Express electrical physical layer (PHY) device to handle the low level protocol and signaling. The PX1011A serializes and deserializes a single PCI Express 2.5Gbps lane to an 8-bit data PXPIPE interface, which is a superset of the PHY Interface for the PCI Express (PIPE) specification.



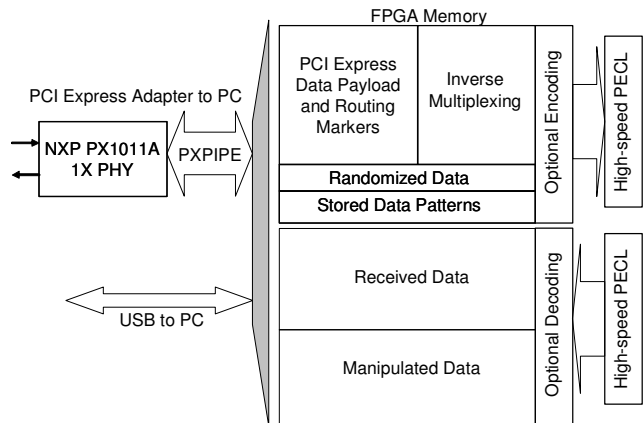
**Figure 4.** PCI Express interface card

The PHY greatly reduces the design complexity required for a system to interface to PCI Express, performing data and clock recovery from the serial stream, serialization and deserialization, 8b/10b encoding and decoding, etc. The PXPIPE interface provides a pair of 8-bit source-synchronous data busses and control signals at logic levels and data rates compatible with the FPGA, which is responsible for implementing the Media Access Control (MAC) layer. **Figure 5** shows an overview of the interface signals.



**Figure 5.** PXPIPE signal diagram

PCI Express data packets received from the computer are buffered briefly within the FPGA before being formatted into an eight payload wide packet that will be injected into the network. This buffering through the PHY and the FPGA, though it introduces additional unwanted latency and extra processing, does present a valuable opportunity for testing of the system. The incoming data stream from the PC via the PCI Express port is mapped into a region of memory within the FPGA that can also be accessed by other sources (currently defined as a pseudo-random data generator or stored data patterns configurable over the USB link). Assuming a loopback configuration, this allows a number of tests to be performed, including end-to-end verification from the PC and automated self-test from the FPGA using random or stored patterns. The traffic can also be monitored over the USB link (not in real-time) or manipulated by logic within the FPGA to simulate various network failures, such as corrupted or lost packets (see **Figure 6**).



**Figure 6.** Alternate data sources and encoding configurations

Also, since the current network implementation is designed with burst communications in mind, the outgoing data is transmitted without any additional encoding that would otherwise be used to ensure an average power in the resulting signal. Should the need arise to explore optional encodings, the FPGA can be reprogrammed to evaluate a wide variety of encoding methods with little additional overhead required.

The following data in this paper was gathered from experiments utilizing synthesized traffic from the DLC. The experimental results using data transmitted across the PCI Express connector were not complete at the time of the submission of this paper for publication. These results will be in the presentation slides.

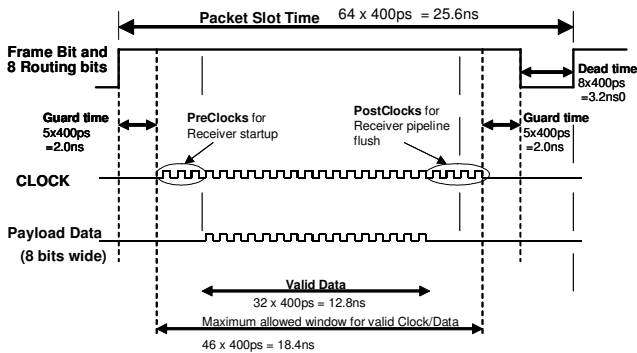
## 4. Packet Composition

The primary motivation for the development of this test system has been to support the testing and evaluation of the optical network in the context of a complete, end-to-end communications system. Such testing is difficult or impossible to perform using alternative methods because

of the burst timing and packet nature of the signals. Payload channels have been independently verified over fixed propagation paths through the network [3], but to study the system along dynamic routing paths a custom test system is required to capture and evaluate the signals at the packet level.

The current test electronics and the preceding version have been used to evaluate various packet structures that are compatible with the OPS network. The packet structure shown in **Figure 7** has been very successful so far though it has evolved slightly from its original introduction [4]. Eight data signals (each 32 bits in length) are decoded from the PCI Express packet or synthesized locally. These are precisely aligned in time with a source synchronous reference clock which is used to sample and recover the data at the destination. The length of this clock signal is slightly longer due to the setup and pipeline flush requirements of the deserializer used in the receiver. The presence of this clock allows for sampling of the data without the use of a clock and data recovery circuit which would not be usable with the short burst packets required by the network. Also, due to the source synchronous nature of this signaling, data recovery can be performed regardless of the propagation path through network as the clock and data timing relationship remains fixed.

The Frame bit signals when the data is valid and the eight header channels carry the routing address data, which is used within the data vortex along with the frame to transparently route the message to the desired port.

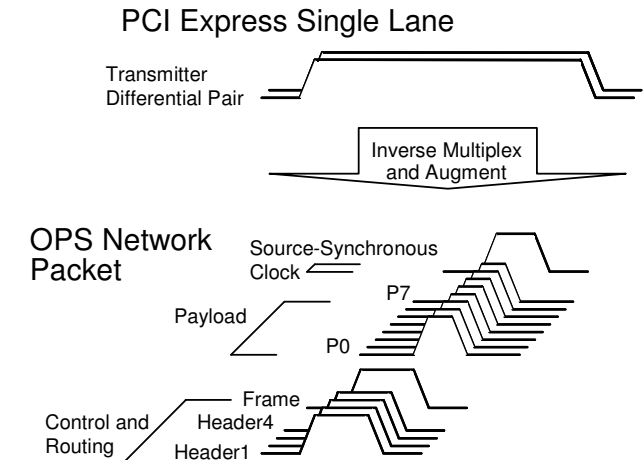


**Figure 7.** Basic packet structure

The individual fiber-optic interconnections of the OPS network are sized to match a packet injection once every 25.6ns or 64 bit periods at 2.5 Gbps. The payload data occupies 32 of these bit periods, with the remaining 32 bits allocated as pre-clocks to condition the deserializer, post-clocks to flush the internal pipeline of the deserializer, guard times to either side of the clock and the enclosing edges of the frame/header signals, and a dead time between all packets. The first two features are accommodations for the electronic circuitry used to sample and recover the signal at the test electronics receivers while the later two features accommodate the particular requirements of the OPS network. The sizable guard times

are required due to the routing nature of the network which does not buffer packets. Decisions are made at each routing node and the packet is routed through one of two output ports. Should the destination node be occupied, the packet is deflected to the other port. A virtual buffer is therefore created, allowing a packet to circulate around a cylinder until the exit port or routing path is available. Unfortunately, most commercially available optical components are not designed for such rapid switching and short data bursts. As a result, there is the possibility for a slight trimming at the extreme edges of the packet as each routing decision is made which these dead times are designed to accommodate.

By comparison, the length of a PCI Express packet is quite large, with even the relatively short 128ns packet used for this project being approximately five times longer than the slot injection time. Due to the limited payload space within a network packet, this information must be inverse multiplexed and distributed across eight parallel data channels, effectively time compressing the signal to fit within the available data window (**Figure 8**). The packet is further augmented with a ninth channel that is utilized as a sampling clock at the destination. Routing information is also extracted from the PCI Express packet and translated into the appropriate routing bits used within the OPS network itself.

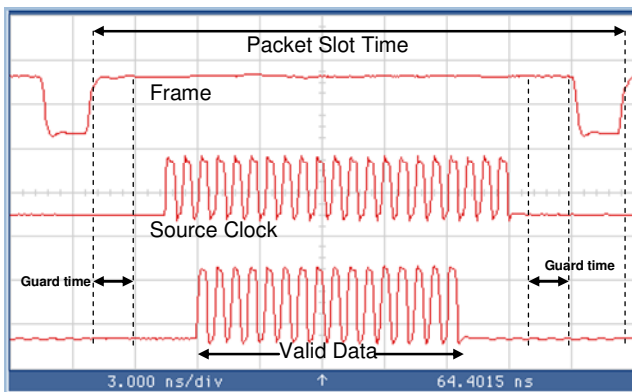


**Figure 8.** Time compression and relative mapping of a PCI Express lane to an OPS network packet

Due to the decoding necessary to process the data, extract the routing information, and map the data onto the eight channels, the signal must be buffered temporarily and increases slightly the amount of latency seen by the system. While this is unfortunately unavoidable, it is also the last time until the packet arrives at the receiver that this occurs. Once the packet is passed through the E/O conversion modules, no additional processing or buffering occurs to the signals, which would otherwise be required in an electrically based network [10].

The system is presently configured for a single lane of 2.5 Gbps PCI Express traffic, but the parallel nature of PCI Express lanes can be leveraged to add additional payload channels to a packet. Just as additional lanes of PCI Express signals can be broadcast in parallel, additional payload channels can be added using a collection of eight optical data channels for every additional lane supported. Since the network implementation is independent of the payload, no additional Frame or routing signals are required.

**Figure 9** shows the Frame, Clock, and a single Payload channel from a fully formatted packet. Since the payload (including the clock) is passed through the OPS network unprocessed, the exact positioning within the packet is variable provided the clock-to-data relative timing remains consistent. It has actually been observed that moving these signals earlier into the packet, eroding the guard time at the left of the image, works better than centering the data in the middle of the packet slot time.



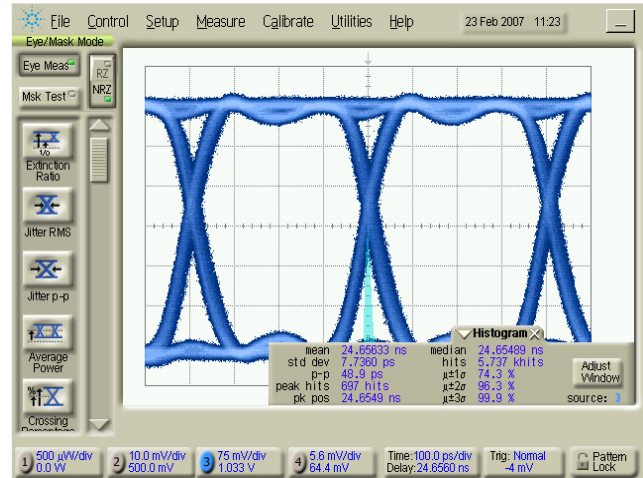
**Figure 9.** Test stimuli signals used for the Optical Test Bed application

## 5. System Performance

Thus far the majority of effort and design focus has been in the evaluation and verification of the test systems electrical characteristics in a standalone state and while interfaced to the optical network. The results gained from the previous design indicated, as expected, that the most critical parameter affecting overall system performance was timing precision. As a result, a lot of effort was placed into the redesign of critical features of the new system, such as the devices used in the signal path and the physical layout of the reference clock and received clock distribution trees.

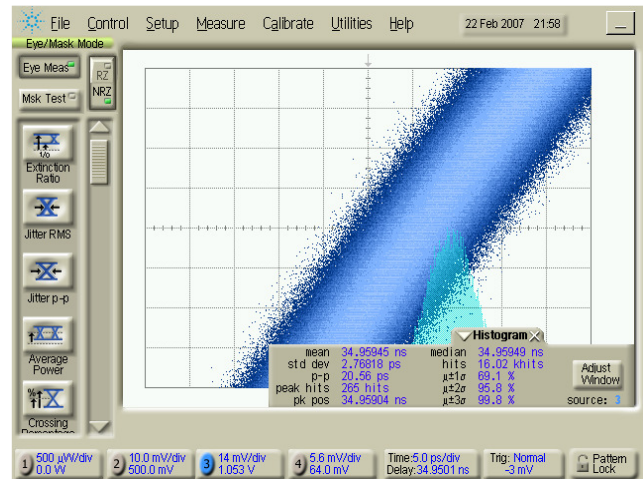
**Figure 10** shows a 2.5 Gbps eye diagram from a payload channel on the newest test electronics board. For this test, the output signal is a pseudo-random bit sequence generated by the FPGA using a linear-feedback shift register (LFSR). The signal has approximately 49ps of peak-to-peak jitter measured at the crossover point. Though this value is up slightly from the 47ps measured on the previous board, the difference can be accounted for

by a longer measurement time and increased data samples. Qualitatively, the new signal appears better with less overshoot and ripple as well as more uniform eye openings, which remained at approximately 0.88 unit intervals (UI).



**Figure 10.** 2.5 Gbps eye diagram

An individual repeated edge was measured as having 21ps of peak-to-peak jitter, down from the 24ps measured previously (**Figure 11**). With the silicon germanium (SiGe) buffers configured for full amplitude swing, the 20-80% rise and fall times single edges were found to be in the range of 70 to 75ps.



**Figure 11.** Jitter measurement on a single rising edge

Another timing issue that affects the optical signals is dispersion. Since the individual signals within the packet are on separate wavelengths, and the speed of propagation depends on the wavelength, the signals will drift slightly over time, proportional to the distance traveled through the network. While there are provisions for adjustable time delays on all of the channels, these are designed for testing purposes and to compensate for small variations in part delays or fiber length mismatches. They are inappropriate for and incapable of real-time adjustment to compensate


for dispersion effects of incoming packets. Any such adjustment would require knowledge of the number of hops taken through the system, which is not available due to the lack of a centralized control system, nor are the parts capable of adjusting at the rates required. It is therefore necessary to characterize the effect of this dispersion on the performance of the system.

Of the 350ps usable eye opening, the available sampling window in which the deserializing circuitry will properly function was experimentally measured as 270ps due to the setup and hold times required by the device. This window was measured by using the tunable delay, setting it to the minimum value that would allow the signal to be properly decoded across its full length of 32 bits. The range was gradually increased until the decoded signal became invalid. These results were obtained in the ideal case, operating the system in a loopback configuration with a cable going directly between the SMA ports of the board. Passed through the opto-electrical components and routed through five hops in the network, the signal was reduced to a 150ps usable timing window. It is believed that this reduction is a cumulative effect as the signal passes through the distributed printed circuit board, E/O, and O/E components as well as additional dispersion across the WDM bit-parallel message [11].

Though the system uses a source synchronous clock generated in parallel with the data, this clock is distributed to eight destinations at the receiver, introducing an additional source of timing variance in mismatched line lengths or part variation in the distribution tree. Combined with the reduced sampling window due to E/O and O/E conversions and dispersion as measured above, a diffused clock across eight or more channels would greatly reduce the usability and scalability of the system. Dispersion, since it is linearly proportional to the distance of travel, affects the number of hops that a packet can take which limits the total possible size of the network. Dispersion effects also grow larger as the total separation of wavelengths grows larger which would occur as additional payloads are added. Though these effects can be mitigated through design decisions and calibrated for, to an extent, even perfectly aligned signals at the input will vary in time at the destination. With a common clock across all of the channels, the final sampling window in which all the signals can be recovered is going to be smaller than that of an individual channel (**Figure 12**).

Additional variance in the clock distribution across channels will only increase the disparity. The total clock variance across the eight data channels was measured as approximately 70ps, but the difference in two clusters of four (channels one through four and five through eight) showed the clock to be within 20ps, suggesting possible part variation in the two halves of the distribution tree. Fortunately, since the data on each of the channels can be time delayed independently, this clock-to-data disparity


can be mitigated through calibration and achieve  $\pm 10$ ps alignment.



**Figure 12.** Cumulative effect of jitter and clock distribution on the data sampling window

The clock-to-data relationship on the transmitter side however is a fixed value, depending on the layout and parts used. The reference clock distribution across the nine high-speed channels was measured as approximately 20ps and the data from the FPGA varied up to 170ps among a random sampling of the 72 signals, though these 8-bit busses operate at only 312.5Mbps and can tolerate such values.

The amount of dispersion that a group undergoes is related to the particular set of wavelengths and the original optical spectrum was not selected with these effects in mind so the arrangement was redesigned. The physical network is tuned to a specific set of routing and frame wavelengths, so only the payload and clock signals were changed. The previous and updated spectrums are shown in **Figure 13**. The revised spectrum has been designed to reduce the effect of chromatic dispersion by centering the clock between the payload channels and condensing the overall range down to a tighter region. The nine wavelengths (eight payload + one clock) are distributed from 1543.73 nm to 1550.12 nm with 0.8 nm spacing between each channel, effectively reducing the timing skew between the outermost two channels. Additionally, by selecting the clock wavelength midway between the shortest and longest wavelength, the maximum clock-to-data timing skew is reduced to 58 ps/km. A very large data vortex network, having 10k x 10k input and output nodes, would incorporate approximately 200m of interconnected fiber. This results in a maximum clock-to-data timing skew of less than  $\pm 12$  ps (for SMF28 fiber with dispersion = 18 ps-nm/km) [3].



**Figure 13.** Optical packet spectrum from the previous design (left) and current (right). Wavelengths labeled BP are payload signals, F is Frame, and CLK is Clock. Unlabeled wavelengths are for routing. [11]

The O/E and E/O interfaces have also been physically optimized for compactness, flexibility, and mobility in a tower configuration. The number of components increases linearly with the number of channels used, and is therefore much more bulky than the test electronics system (which is less than double the previous size, due mostly to the space constraints of the SMA connectors). The O/E receiver was also redesigned to accommodate true burst-mode reception, as the incoming signal is at best active for only 50% of the time and is not encoded to maintain an average power.

## 6. Limitations and Future Work

Though this system meets many of our current design needs, there are a number of areas for improvement or future enhancement. For example, newer FPGAs natively implement SerDes blocks that could minimize or eliminate completely the need for external supporting PECL circuits, though it remains unknown if the strict timing requirements of this application can be satisfied. Supplementary experiments are ongoing to explore this particular question. The PCI Express standard is also being upgraded, supporting data links at 5.0 Gbps per lane. To support this higher speed, the payload channels would have to be upgraded to a higher rate, additional channels added, or a shorter PCIe packet used from the PC. Any of these solutions (or a combination of) would be non-trivial, i.e. requiring a redesign of the transmitters/receivers, improved support for wider payloads, etc.

At the system level, the present communication protocols call for one-way transactions. This leaves the system vulnerable to dropped, corrupt, or out of order packets. There are no current plans to investigate these problems, but the test infrastructure is in place to assist in the development of higher layer protocols to address these issues and increase the system reliability and usability as a whole.

It has also been noted that though this system has been developed with a specific opto-electronic application in mind, many of the features and design elements can be applied to direct testing of electrical components and systems. The receiver logic is currently very specific to

the packet and burst nature of the signals, but some preliminary work has already been made to explore the modifications necessary to adapt the system to more general purpose usage. Some experiments have already been conducted and are being planned utilizing the system as a high-speed, multi-channel programmable data source to stimulate various test circuits.

## 7. Conclusions

High-performance shared-memory computing systems are very reliant upon high-bandwidth and low latency interconnections. The system presented enables system-to-system transport of PCI Express originating traffic across an optical packet switched network designed for low latency communication. The system contains many test and design features necessary to bridge these complimentary technologies while providing for in-line and loopback functional tests, characterization and performance evaluation, and the flexibility to adapt the system to a wide range of future needs.

## 8. References

1. D. Dai, D.K. Panda, "How much does network contention affect distributed shared memory performance?," *Proc. Int. Conf. Parallel Process.*, Aug. 11-15, 1997, pp. 454-461.
2. C. Hawkins, D.S. Wills, "Impact of Number of Angles on the Performance of the Data Vortex Optical Interconnection Network," *Journal of Lightwave Technology*, Vol. 27, Issue 9, pp. 3288-3294, Sept 2006.
3. A. Shacham, B.A. Small, O. Liboiron-Ladouceur, K. Bergman, "A Fully Implemented 12x12 Data Vortex Optical Packet Switching Interconnection Network," *J. Lightwave Technol.* 23 (10) 3066-3075 (Oct 2005).
4. D.C. Keezer, C. Gray, A. Majid, N. Taher, "Low-Cost Multi-Gigahertz Test Systems Using CMOS FPGAs and PECL," pp. 152-157, *Proc. of Design, Automation, and Test in Europe*, March 2005.
5. C. Gray, D.C. Keezer, O. Liboiron-Ladouceur, K. Bergman, "Multi-Gigahertz Source Synchronous Testing of an Optical Packet Switching Network," International Mixed-Signals Test Workshop 2006, Edinburgh, Scotland (June 2006).
6. A. Li, J. Faucher, D.V. Plant, "Burst-Mode Clock and Data Recovery in Optical Multiaccess Networks Using Broad-Band PLLs," *IEEE Photonics Technology Letters*, Vol. 18, NO. 1, pp73-75, Jan. 1, 2006.
7. J.S. Davis, D.C. Keezer, "Multi-Purpose Digital Test Core Utilizing Programmable Logic," *Proc. of the Intl. Test Conf.*, pp. 438-445, October 2002.
8. J.S. Davis, D.C. Keezer, K. Bergman, O. Liboiron-Ladouceur, "Application and Demonstration of a Digital Test Core: Optoelectronic Test Bed and

- Wafer-level Prober," *Proc. of the Intl. Test Conf.*, pp.166-174, Sept./Oct. 2003.
9. A.M. Majid, D.C. Keezer, "An Improved Low-Cost 6.4 Gbps Wafer-Level Tester" *Proc. of the 7th IEEE Electronics Packaging Technology Conference (EPTC)*, pp.814-819, December 2005.
  10. G. I. Papadimitriou, C. Papazoglou, A.S. Pomportsis, "Optical Switching: Switch Fabrics, Techniques, and Architectures", *Journal of Lightwave Technology*, Vol. 21, Issue 2, pp 384-405, Feb 2003.
  11. O. Liboiron-Ladouceur, C. Gray, D.C. Keezer, and K. Bergman, "Bit-Parallel Message Exchange and Data Recovery in Optical Packet Switched Interconnection Networks," *IEEE Photonics Technology Letters*, Vol. 18, No. 6, pp 779-881, Mar. 2006.
  12. C.E. Gray, O. Liboiron-Ladouceur, D.C. Keezer, K. Bergman, "Test Electronics for a Multi-Gbps Optical Packet Switching Network," *Proc. of the 8th IEEE Electronics Packaging Technology Conference (EPTC)*, pp 373-378, December 2006.