FirstLight: Pluggable Optical Interconnect Technologies for Polymeric Electro-Optical Printed Circuit Boards in Data Centers

Richard C. A. Pitwon, Kai Wang, *Student Member, IEEE,* Jasper Graham-Jones,
Ioannis Papakonstantinou, *Member, IEEE,* Hadi Baghsiahi, Bert Jan Offrein, *Senior Member, IEEE,* Roger Dangel,
Dave Milward, and David R. Selviah, *Member, IEEE*

Abstract—The protocol data rate governing data storage devices will increase to over 12 Gb/s by 2013 thereby imposing unmanageable cost and performance burdens on future digital data storage systems. The resulting performance bottleneck can be substantially reduced by conveying high-speed data optically instead of electronically. A novel active pluggable 82.5 Gb/s aggregate bit rate optical connector technology, the design and fabrication of a compact electro-optical printed circuit board to meet exacting specifications, and a method for low cost, high precision, passive optical assembly are presented. A demonstration platform was constructed to assess the viability of embedded electro-optical midplane technology in such systems including the first ever demonstration of a pluggable active optical waveguide printed circuit board connector. High-speed optical data transfer at 10.3125 Gb/s was demonstrated through a complex polymer waveguide interconnect layer embedded into a 262 mm × 240 mm × 4.3 mm electro-optical midplane. Bit error rates of less than 10^{-12} and optical losses as low as 6 dB were demonstrated through nine multimode polymer waveguides with an aggregate data bandwidth of 92.8125 Gb/s.

Index Terms—Optical planar waveguides, optical polymers, optical waveguides, optoelectronic devices.

I. INTRODUCTION

PREVAILING trends in the data storage industry [1], [2] are poised to severely impact the design of future data center subsystems. Over the last decade, the volume of data being cap-tured processed, stored, and manipulated as digital information has increased exponentially. The increase in demand for digital data storage capacity has, in particular, been fuelled by information from critical business applications, email communications, the Internet, and multimedia applications. The projected global digital demand will grow from 1.8 to 7.9 zettabytes (ZBs) from 2011 to 2015 of which the amount available for data storage will reach 3.3 ZB in 2015 [3].

However, the exponential increase in system bandwidth and density required to satisfy this demand will impose unmanageable cost and performance burdens on future data center technologies. In particular, the integration of more data intensive applications, such as high-performance computers, and the reduction in size and the increase in the number of high-speed ports of peripheral storage devices, such as hard disk drives, will cause the density of printed channels on the data storage midplane to increase. The increase in data communication speeds will further expose the system to some of the fundamental constraints incurred when higher frequency data are conveyed along electronic channels. Many of these constraints can be mitigated to some degree [4], however, at a mounting cost to the overall system design [5]–[8].

A. Modern Data Storage System Architectures

Data storage technologies, as shown in Fig. 1, form the crucial building blocks of modern data centers, wherein data storage arrays (typically magnetic or solid state media) are connected within systems of varying complexity and size ranging from simple high capacity storage racks to high performance computing data storage systems.

A generic data storage array system [see Fig. 1(a)] comprises an array of hard disk drives connected to one side of a passive midplane, while controller modules and power supplies are connected to the other side. As shown, the midplane and its peripherals are connected in a mutually orthogonal geometry, which, as described later, is advantageous for in-plane optical connections.

Fig. 1(b) shows a data storage integrated platform, which incorporates greater functionality such as server hardware into a data storage system. Fig. 1(c) shown a data storage rack in which assorted data storage subsystems, such as those shown in Fig. 1(a) and (b) are incorporated according to the required capacity and processing power of the rack and data center application.

Fig. 1. (a) Generic data storage system with storage media connected to one side of the midplane and controller modules and power supplies connected to the other. (b) Data storage integrated application platform. (c) 2 Petabyte data storage rack.

B. Data Center Interconnect Topologies

The interconnect topology of high-speed, high-reliability (enterprise class) data storage systems is typically defined by a passive dual star configuration, whereby each data storage device supports two duplex data links on the midplane, one to each of at least two separate controller modules. As a current example, the Storage Bridge Bay Specification [9], which defines mechanical, electrical, and low-level enclosure management requirements for an enclosure controller slot, allows for a maximum of 48 disk drives in a given data storage system enclosure, each supporting two duplex links to each of two controller modules, with a further 17 duplex links directly connecting the controllers to each other. The midplane of a 48-drive storage array would therefore have to accommodate 113 duplex links or 226 high-speed transmission lines. Furthermore, there is an enhancement of the drive interface standard in development [10], which allows up to four independent duplex links per drive scaling the number of high-speed links on the data storage midplane accordingly. The level of fault tolerance and scalability offered by these topologies forces increased complexity and cost into the midplane, particularly when interconnect protocols define serial data rates beyond 12 Gb/s.

C. Performance Constraints of High-Speed Electronic Links

Serial attached SCSI (SAS) is a point-to-point bus protocol, which defines the speed with which data are conveyed between the peripheral devices in a storage system environment. The SAS roadmap [11] currently defines a serial data rate of 12 Gb/s and is set to double to 24 Gb/s by 2016. PCI Express is also gaining prominence as an alternative drive interface, specifically for solid state drives. PCI Express 3.0, which is currently implemented, is specified to run at a serial data rate of 8 Gb/s.

There are a number of factors, which will limit the operational bit rate of the high-speed links on a commercial copper midplane. Dielectric absorption and skin effect are the key loss mechanisms on a copper trace [12], which cause an increase in signal attenuation with frequency, while electro-mechanical connectors introduce parasitic capacitance and inductance effects, and conductive vias between layers within the printed circuit board (PCB) can act as impedance stubs giving rise to partial reflections in the signal path [13], [14]. At signal data rates of 24 Gb/s and higher, additional design measures need to be

taken into account including the use of lower dielectric loss PCB substrates, skew and loss controlled electro-mechanical connectors and enhanced via control techniques such as back-drilling or buried vias [15]–[17]. In order to mitigate rising crosstalk, copper channels will need to be moved further apart; however, due to the spatial constraints of modern enclosure form factors for data communication systems, more complex routing patterns will be required and the number of high-speed layers in the midplane PCB increased. The available space on the data storage midplanes is further restricted by the need for milled access slots to allow sufficient air flow through the system, while space on controller and other peripheral cards is consumed by ever increasing component densities as functionality is scaled.

The exponential increase in capacity, processing power, and bandwidth density in future data storage systems can be addressed by incorporating high-density optical channels at the midplane and controller board level within data storage enclosures.

D. Implementation of Embedded Optical Interconnect

It is, therefore, proposed that the projected performance bottleneck in data storage systems is mitigated by incorporating electro-optical printed circuit board (OPCB) and interconnect technology on the midplane [18]–[23]. A number of key challenges to the commercial proliferation of optical PCB interconnect technology have yet to be overcome. These include the development of low-cost methods of waveguide fabrication and deposition, reduction in minimum in-plane waveguide bend radius to make PCB routing of optical channels viable, repeatable high-precision assembly of optical components, and a viable method of optically connecting line cards to the waveguide interface of an optical PCB.

In order to evaluate the viability of these technologies in a data center environment, a demonstration platform has been constructed, comprising four test line cards, which are plugged electrically and optically into an electro-optical midplane. An OPCB has been developed incorporating conventional ten copper layers for electrical and electronic interconnect, and an optical interconnect layer comprising polymer optical waveguides. In addition, pluggable optical connectors have been developed incorporating high-speed parallel optical transceivers and a passive alignment mechanism to ensure accurate dynamic optical engagement between the transceiver interfaces and the embedded polymer optical waveguides in the OPCB. The polymer waveguide layer includes self-alignment features to enable passive alignment and assembly of proprietary optical connector receptacle devices onto the polymer waveguides. A complex optical interconnect pattern was designed to meet exacting specifications to demonstrate how polymer waveguides would perform when subjected to the routing constraints expected within a conventional midplane form factor.

In this paper, we introduce the FirstLight project, which focuses on four key technology enablers namely 1) design of compact optical waveguide layouts appropriate for data system midplane form factors; 2) fabrication of OPCBs with embedded polymer waveguides; 3) low-cost, high-precision techniques for passive alignment and assembly of optical interface components onto polymer waveguides; and 4) complete pluggable in-plane

Fig. 2. Schematic illustrating waveguide embedded in lower and upper claddings.


Fig. 3. Waveguide routing pattern with four optical engagement apertures for the FirstLight demonstrator. (a) Negative cascading waveguide. (b) Positive cascading waveguide.

connector solutions for polymer OPCBs. The FirstLight demonstration platform is presented, which brings together all these technology enablers into a high bandwidth density data communications enclosure.

II. ELECTRO-OPTICAL MIDPLANE WITH EMBEDDED POLYMER WAVEGUIDES

A. Polymer Optical Waveguide Layer

A passive electro-optical midplane was designed, which included ten electrical layers devoted to power distribution and low-speed bus communication and one optical polymer interconnect layer to convey high-speed (10.3125 Gb/s) serial data between peripheral connectors. An optical interconnect layer was fabricated on one surface of the midplane PCB stack-up from a proprietary acrylate/polyurethane based polymer exhibiting a propagation loss of 0.03–0.04 dB/cm at a wavelength of 850 nm. The optical layer stack comprised a 70 μm core layer, sandwiched between a 100 μm lower cladding and a 30 μm upper cladding layer (measured from above the top of the core layer), whereby the polymer in the guiding core layer had a higher refractive index than that in the bounding cladding layers (see Fig. 2).

The refractive index of the core material was $n_{core} = 1.5600$, while that of the cladding was $n_{cladding} = 1.5240$, giving rise to step-index multimode (MM) waveguides with a core-cladding index difference of $\Delta n = 2.3\%$ and a numerical aperture (NA) of 0.33. Rectangular channels were patterned using a vectorial laser direct imaging (LDI) writing process to define waveguides in the core layer with a cross section of 70 μm × 70 μm, which was suitable to meet the launch and capture tolerances on the optical transmit and receive elements [24], [25]. By deploying the high-precision passive alignment technique described in Section III, additional losses are anticipated to be no more than a maximum of 0.2 dB due to vertical cavity surface emitting laser (VCSEL) misalignment tolerance and of 0.4 dB [26] due to photodetector (PD) misalignment tolerance for the waveguides designed.

B. Optical Interconnect Design

In the demonstration platform, a complex polymer waveguide interconnect layout was designed to form the optical layer of a 262 mm × 240 mm × 4.3 mm electro-optical PCB. The design was guided by two key requirements: 1) in rack chassis systems, the peripheral cards are similarly oriented, so this was set as a constraint for the design of the system. All midboard optical engagement interfaces were oriented in the same direction to allow line cards to be pluggable in arbitrary slots and interchangeable. 2) The design needed to demonstrate the routing compactness and maneuverability, which was typical of a high-density data storage system. Therefore, waveguide structures are needed to be carefully designed according to meticulously calculated optical waveguide layout design rules [25] to minimize the optical loss in each waveguide segment and to ensure that the aggregate (total) insertion loss for each waveguide is such that the optical power at the exit of each waveguide falls within the receiver sensitivity threshold to allow bit-error-free signal transmission.

The optical interconnect layer, designed by UCL and IBM Research was defined by a complex routing pattern (see Fig. 3), which included four quasi-rectangular optical engagement apertures, multiple nonorthogonal crossings, and cascading 90° bends. The engagement apertures are interconnected by a point-to-point waveguide network, whereby each aperture maintains one bidirectional link (comprising two waveguides) to every other aperture, resulting in a total of 12 waveguides on the board. The sizes of the engagement apertures were determined by the form factor of the pluggable connector prototype, which is described in Section IV.

All waveguides, but two, have four cascaded 90° bends comprising a negative cascade followed by a positive cascade. A negative cascade occurs when one 90° bend is followed by another 90° bend, which curves in the opposite direction to the first bend, giving rise to an inflection point in the waveguide and a net waveguide angle change of 0° [see Fig. 3(a)]. A positive cascade occurs when one 90° bend is followed by another 90° bend, curving in the same direction as the first and thus, giving rise to a net waveguide angle change of 180° [see Fig. 3(b)]. To minimize bend losses, a radius of curvature of 17 mm [24] was applied on all bends as this gave the lowest loss. A number of waveguides intersect at one or more positions to accommodate

Fig. 4. Top view micrographs showing two waveguides intersecting with (a) a crossing angle of 90° and (b) a crossing angle of 145°.

TABLE I
OPTICAL LOSS OF EACH WAVEGUIDE SEGMENT

Combined input/output coupling loss at both ends (dB)	90° Bends		Propagation loss (dB/cm)	Crossing	
	Radius (mm)	Loss per bend (dB)		Crossing angle (°)	Loss per crossing (dB)
4.60 ± 0.50[1]	16.5	0.94 ± 0.07	0.08 ± 0.01	160	0.08 ± 0.03
	17.5	0.83 ± 0.07		150	0.05 ± 0.03
2.07 ± 0.50[2]	18.5	0.85 ± 0.07		145	0.05 ± 0.03
	19.5	0.89 ± 0.07		140	0.04 ± 0.03
	20.5	0.95 ± 0.07		130	0.03 ± 0.03
	22.5	0.97 ± 0.07			
	23.5	0.99 ± 0.07			
	24.5	1.03 ± 0.07			

[1] Average coupling loss without index matching fluid.
[2] Average coupling loss with index matching fluid.

space restrictions. The crossing angles chosen range from 130° to 160° (with respect to the relative propagation direction of light in the two waveguide branches) and the measured optical losses are from 0.03 to 0.08 dB per crossing [25] (see Fig. 4).

The optical layout was designed and the corresponding optical insertion loss calculated based on previous experimental measurements [24] carried out on photolithographically fabricated polymer acrylate test waveguide [27] samples, which were cut using a dicing saw. The test waveguide elements measured included 90° bends and crossings at a range of angles. The insertion loss of each test waveguide is listed in Table I. The test waveguide elements listed were deployed in the FirstLight optical board layout and were used to calculate the power budget.

UCL had previously found [24], [25] that the optical loss decreased as the bend radius increased; however, as the bend radius increases beyond a certain point, around 17 mm in our measurements, propagation loss becomes more significant than the combination of transition and radiation loss since the larger radius bends have longer lengths. This trend is shown in Table I. The least loss of 90° bends was measured for bend radii in the range of 16–22 mm. Therefore, we chose bend radii values within this range taking the polymer waveguide fabrication techniques into account. The power budget of each optical link depends on the optical transmitter output power and on the sensitivity of the receiver used in the system. A receiver sensitivity of −11.1 dBm [28], matching the transceiver devices used in the demonstrator, was selected as the threshold for receiving error-free signals at 10.3125 Gb/s for our system design. The output power of the optical transmitters deployed in the system was −1.48 dBm, which limited the maximum tolerable insertion loss in each optical link to a power budget of 9.62 dB.

The division of the waveguides into straight and right-angled bend segments is a useful tool for calculating optical link loss.

TABLE II
CALCULATED OPTICAL LOSS IN EACH WAVEGUIDE INTERCONNECT IN TERMS OF CASCADED LOSS FROM EVERY WAVEGUIDE SEGMENT

Waveguide	No. of crossings	No. of 90° bends	Length of straight section (cm)	Calculated optical loss (dB)	
				1	2
1	0	4	2.17	8.24	6.04
2	0	4	2.17	8.24	6.04
3	0	4	2.14	8.26	6.06
4	3	4	2.24	8.54	6.34
5	5	4	2.04	8.48	6.28
6	3	4	5.95	8.67	6.47
7	3	4	5.78	8.68	6.48
8	2	4	6.80	8.87	6.67
9	4	4	9.67	9.17	6.97
10	0	4	12.99	9.30	7.10
11	0	6	16.30	11.36*	9.16
12	0	6	20.19	11.81*	9.61

[1] calculation based on the coupling loss without index matching fluid.
[2] calculation based on the coupling loss with index matching fluid.
*insertion loss is higher than the threshold 9.62 dB due to routing constraints.

The link losses on all 12 waveguides designed were calculated based on cascading the losses of all segments comprising a given waveguide. This neglects any mode mismatch loss occurring between different segments, which is expected to occur particularly between cascaded positive and negative bends and between straight and bend sections. The link losses are listed in Table II, which shows that all of the 12 waveguides were designed to have insertion losses lower than the error-free threshold of 9.62 dB when index matching fluid ($n = 1.5433 \pm 0.0005$ for 840.0 nm at 25 °C) was applied on the coupling interfaces, while 10 out of 12 waveguides fell within the desired threshold when no index matching fluid was applied. The predicted optical losses of the last two waveguides exceeded the desired threshold, due to routing constraints increasing both the length of and the number of bends in the waveguides, such that both material absorption and bend losses became prohibitive.

Though the relatively large minimum bend radii required at this stage would place significant routing constraints on future optical PCB layouts, these could be effectively mitigated by refinement of manufacturing techniques or novel structuring of the waveguide to reduce bend loss as demonstrated by Xyratex [29].

C. OPCB Fabrication Process

The electro-optical midplane fabricated by Varioprint and IBM Research-Zürich was built up of ten copper layers and one polymer layer. It is well known that the surface of such PCBs is not flat. Preliminary tests on the electrical layers showed significant thickness variation on the complete stack-up of the PCB. Experimental characterization of surface height variation of the section of the PCB substrate [the yellow rectangular insert in [see Fig. 5(a)] reserved for the optical waveguides was carried out through surface roughness scans at various locations, including an area containing copper fiducials, which served as alignment reference features for high-precision positioning of the waveguides onto the PCB. Scan 1 was taken on the FR4 PCB surface near the edge of the substrate section with no copper features along its path and yielded a measured
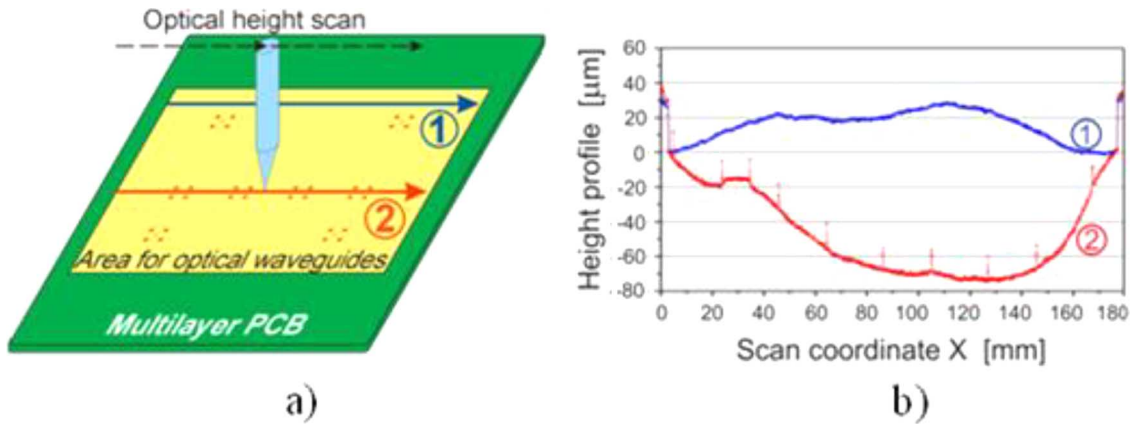
Fig. 5.   (a) Schematic of FR4 substrate surface scans performed on a PCB containing up to ten laminated copper layers. Scan 1 carried out on the reserved PCB substrate surface near the edge with no copper marks, height variation ± 12 μm; scan 2 traversed some copper marks, height variation ± 38 μm. (b) Experimental results of the two scans on the FR4 substrate surface showing the overall height variations of up to ± 50 μm.


Fig. 6.   Schematic depiction of polymer layer deposition using doctor blade method.

surface height variation of ± 12 μm. Scan 2 was taken on the FR4 PCB surface along the middle of the substrate, which traversed some copper features and yielded a measured surface height variation of up to ± 38 μm across the PCB [see Fig. 5(b)] giving us the overall surface height variation ± 50 μm on the substrate.

If the optical layer were to have been deposited directly onto the electrical PCB surface by using the well-known doctor blade method (see Fig. 6), the substrate surface variation would be reduced according to (1) which is based on the deposition technique and material used by IBM

$$d = 0.6 \times G \qquad (1)$$

where d is the polymer layer thickness and G is the distance between the doctor-blade tip and the substrate. The correction scaling factor 0.60 is specific to this material. If the optical lower cladding layer were deposited directly over the PCB board, the variations of lower cladding thickness would theoretically be reduced to ± 30 μm (± 50 μm $\times$ 0.60), while the core thickness variation would be reduced to ± 18 μm (± 30 μm $\times$ 0.60) accordingly. As these values are unacceptably high, it was decided that the optical layer should be fabricated on a separate flat flexible substrate, which would later be laminated to the electrical PCB (see Fig. 7) in order to minimize the thickness variation of the optical waveguide features.

A proprietary liquid cladding polymer was deposited onto a 100 μm thick polyimide substrate with a doctor-blading process


Fig. 7.   Separate flexible polyimide substrate with optical polymer layers deposited and patterned.

applied to control the thickness. The photosensitive polymer was uniformly cured with collimated ultraviolet light (365 nm) from a Mercury/Xenon (Hg/Xe) lamp to polymerize and cure the 100 μm thick lower cladding layer. A higher refractive index liquid core polymer was then deposited onto the lower cladding layer and doctor-bladed to a thickness of 70 μm. The core features were patterned using the LDI technique whereby the beam of a GaN ultraviolet laser diode operating in continuous wave (CW) mode at 372 nm wavelength was moved across the substrate to selectively cure those parts of the core layer, which would form the waveguides [see Fig. 8(a)]. By means of a subsequent wet-chemical development process step, the nonexposed parts were then removed. As the optical layers were fabricated

Fig. 8. (a) Waveguides measured by microscopic top view fulfill width specification of 70 μm within ± 2 μm. (b) Waveguide heights experimentally determined by optical surface scanning satisfy target height of 70 μm within ± 2 μm. Peripheral "dummy" waveguides are shown to provide mechanical self-alignment as described in Section III.

Fig. 9. Far-field of the 50/125 μm fiber used in the experiment with and without mode filling effect.

on the comparatively smoother polyimide substrate, the waveguide core thickness variation was reduced to within ± 2 μm [see Fig. 8(b)]. The two outer dummy waveguides were slightly thicker than the operational waveguide could be due to the increased substrate surface height toward the edge of the engagement slot.

Mechanical slots were milled into the polyimide backed optical layer and compliant pins assembled onto the electrical PCB in order to align both substrates together prior to a cold lamination process. Finally, the optical engagement apertures were milled out within the outline of the midplane. The midplane was 262 mm long, 240 mm high, and 4.3 mm thick. The thickness of the electro-optical midplane was made of the 4 mm thick electrical stack-up and the 300 μm polyimide and polymer layer stack-up.

D. Insertion Loss Measurements

The measurements of optical insertion loss of the optical interconnects were conducted by UCL and Xyratex. An ST connector packaged 850 nm VCSEL was connected to a standard 50/125 μm step-index MM fiber with NA_{fiber} of 0.22. The fiber core cross section and its NA were both smaller than the waveguide core cross section of 70 μm $\times$ 70 μm and the waveguide $NA_{waveguide}$ of 0.33, respectively. This reduced the fiber coupling loss during the butt coupling measurement. The fiber was 10 m long and was wound 20 times around a 38 mm diameter circular mandrel in order to fill the NA of the fiber and we confirmed through measurement that the NA of the fiber was fully filled with a large number of transverse modes (see Fig. 9). The fiber was then butt coupled to the waveguide launch facet.

The optical power at the launch facet of the fiber was chosen to be -1.48 ± 0.02 dBm. The launch fiber was mounted on a group of motorized translation stages with submicrometer step resolution in three axes, $x, y,$ and z, to accurately align the fiber to the waveguide and to optimize coupling into the waveguide. A thin silicon PD with an 8 mm aperture was required to fit through the engagement aperture of the FirstLight test board and aligned to the output facet of the waveguide in order to capture the light received through the waveguide and measure the waveguide insertion loss. The detector was calibrated against the light condition in the dark room where the measurements were conducted. The loss measured for each waveguide included the coupling losses at the fiber–waveguide

interface and waveguide–PD interface, absorption by the propagation medium, bend transition losses due to modal mismatch between waveguide segments of differing radii of curvature, and losses incurred at crossing junctions. Though the FirstLight board contained 12 waveguides in total, the longest three waveguides were excluded from experimental characterization due to damage during assembly. Insertion loss measurements were made on the remaining nine waveguides both with and without index matching fluid applied. The measurement results are shown in Fig. 9. The waveguides in the optical interconnect layout were designed to never exceed an insertion loss of 9.62 dB, the threshold required to achieve communication at a bit error rate (BER) of less than 10^{-12}.

Fig. 10 includes the comparison of the calculated optical losses without index matching fluid (red column) and the measured insertion losses of waveguides without index matching fluid applied (blue column) which are on average 2.36 dB higher than the calculated predictions. The insertion losses of waveguides, i.e., 2, 4, 5, 6, 7, 8, and 9 are in excess of the 9.62 dB error-free threshold. The higher insertion losses are partially due to the higher scattering losses at the waveguide end facets, which depend on the surface roughness of the end facet in question. The surface roughness in turn depends on how the end facets were cut and polished [30]. In this case, the waveguide end facets are located at the edges of the milled midboard optical engagement apertures. Though the end facet roughness of these waveguides could not be measured without damaging the OPCB, waveguides cut using similar milling techniques to those deployed on this board exhibit RMS surface roughness values ranging from 183 to 350 nm [30]. The gray columns in Fig. 10 represent the calculated insertion losses with index matching fluid, the values for which are 2.57 dB lower than the waveguide losses calculated without index matching fluid applied. The green columns in Fig. 10 represent the measured insertion losses with index matching fluid applied. The measured insertion losses of the nine waveguides were on average 2.93 dB lower than those measured without applying the index matching fluid (blue columns). This discrepancy is of the order of the 2.57 dB discrepancy shown in measured results in Table I. In both cases, the measured loss is higher than the calculated

Fig. 10.   Predicted and measured values for insertion loss of a group of nine waveguides on the OPCB. The plots were arranged in order of increasing waveguide length, starting from the shortest waveguide 1, as listed in Table II. The longest three waveguides were excluded from experimental characterization due to damage during assembly. The dashed line shows the 9.62 dB threshold.

loss with and without index matching fluid. This may be due the fact that we ignored any mode mismatch loss occurring between different waveguide segments. The discrepancies in insertion loss shown in Fig. 10 (both calculated and measured) between waveguides with and without index matching fluid applied highlight the strong dependence of loss on end facet roughness. Index matching fluid, however, is not a practical means of reducing the roughness of waveguide end facets in applications involving repeatable connection to the waveguides as the fluid tends to dry out gradually and may accumulate dust, which in turn will cause the end facet scattering losses to increase in an unpredictable manner. Therefore, a more durable method using waveguide core polymer to smooth the end facets was invented and demonstrated to permanently reduce the end facet roughness. The deployment of this technique improved waveguide transmissivity by 0.49 dB on average compared to waveguides with index matching fluid applied [30].

III. PASSIVE ALIGNMENT AND ASSEMBLY METHOD

One crucial requirement for the commercial deployment of optical PCB technology is a low-cost technique for the high-yield assembly of optical interface components onto the optical layers. In order to enable high-volume assembly, it is preferable that such techniques be passive and repeatable. A proprietary fabrication technique and method of passively aligning and assembling parallel optical microlenses to embedded polymer waveguide arrays was successfully developed [24], [30], [31]. These form a critical part of the pluggable in-plane connection interface between arbitrary external optical devices, either passive or active, and a PCB embedded optical circuit.

A. Fabrication of Passive Alignment Features

The complete fabrication process for the passive alignment features on the optical layer is outlined in Fig. 11. The proce-

dure involves the fabrication of passive mechanical registration features in the core layer during the same process step in which the waveguide cores themselves are patterned using the LDI technique, whereby a 372 nm CW GaN ultraviolet laser diode was moved across the uncured wet film core layer to selectively cure the salient core features. These features included the signal waveguides as well as additional "dummy" waveguides [see Fig. 8(b)], which are positioned on either side of the signal waveguides and as a result, their positional accuracy with respect to the signal waveguides is comparable to that of the signal waveguides to each other. A layer of upper cladding polymer was then spun-coat over the core features; however, instead of uniformly curing the upper cladding, it was selectively cured to ensure that the central signal waveguides were completely clad, while the mechanical registration waveguides were not. For this, a photolithographic mask was aligned over the optical substrate, the mask allowing all areas of the upper cladding to be cured except the mechanical registration waveguides. Thus, during the subsequent wet chemical development step, a clearance is created around the mechanical registration waveguides. This clearance allows for direct mechanical registration of arbitrary components to the waveguides. In addition, the fabrication tolerances required to pattern the upper cladding for this purpose are far lower than those required to pattern the waveguides themselves. It is only important that the outer edges of the registration waveguides, which form the mechanical datum, be mechanically exposed. Preferably, the upper cladding should partially cover the registration waveguides in order to provide structural reinforcement and reduce the risk of the registration waveguides delaminating under the strain. However, this is not strictly necessary and as shown in Fig. 12, the registration waveguides in the FirstLight OPCB were left completely uncovered without any adverse effects. This technique can be implemented using most waveguide fabrication processes. The positional tolerance of the mechanical registration features with respect to

Fig. 11. Fabrication process for the passive alignment features on the optical layer.


Fig. 12. Upper cladding opening showing waveguides for signal transmission as well as connector alignment features (dummy waveguides).

each other has been measured to be ±3 μm for lateral misalignment in-plane and ±4 μm normal to the PCB plane [26], [32].

B. Passive Assembly Procedure

A lensed connector receptacle was designed and developed by Xyratex and Samtec Inc., which comprised a custom moulded receptacle and commercial Omron geometric convex microlens array. The custom receptacle included compliant structures to allow it to mechanically engage with the registration waveguides on the board and a recess to accommodate a standard high performance "mechanical transfer" connector (MT) compliant lens array. The MT compliant interface on the
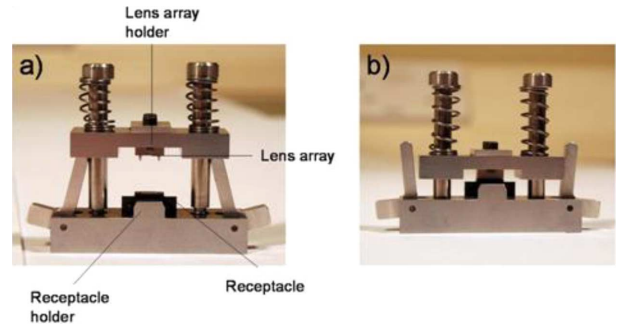

Fig. 13. Mechanical lens assembly jig designed to align the lens array into the custom receptacle. (a) Open position. (b) Engaged position.

lens array included two 0.7 mm pin slots and 12 microlenses arranged between the slots on a 250 μm pitch between lenses. In order to ensure that the lens array was accurately aligned within the receptacle, the receptacle also included two MT compliant pin slots, the dimensions of which matched those in the lens array. The lens array was fastened to the receptacle with a UV curable low shrinkage optical adhesive, Dymax OP-21.

Fig. 13 shows a lens assembly mechanical jig designed to align the lens array into the custom receptacle and hold it tightly in place during the UV curing process. While the lens assembly jig is in the open position [see Fig. 13(a)], the lens array is mounted onto compliant pins in the jig and the receptacle is mounted into a compliant recess. When the jig is closed [see Fig. 13(b)], the lens array is pressed into the lens receptacle and held under a strong spring tension. A UV source was subsequently applied to cure the adhesive between the lens and receptacle. The lens receptacle included recesses to contain and channel the adhesive away from sensitive areas such as the MT pin holes and the microlenses, while maximizing contact between those areas of the lens plate and receptacle which are neither in the signal path nor in the mechanical registration path. It was decided to fasten the lensed receptacle to the optical PCB using a UV curable optical adhesive. As alignment of the receptacle to the waveguides is critical, it was important that the receptacle lay flat on the smooth exposed lower cladding and that none of the adhesive seeped underneath it.

Fig. 14 shows a mechanical receptacle assembly jig designed to hold the passively aligned lensed receptacle tightly in position onto the optical PCB while the adhesive is applied as fillet bonds around the edges of the receptacle and subsequently cured. As with the lens assembly jig described previously, the receptacle assembly jig contained recesses to hold the passively aligned lensed receptacle in place over the waveguides under strong spring tension. The shape of the receptacle assembly jig was customized to be snap-fit into the connector aperture on the board.

Once the lensed or "secondary" receptacle is assembled onto the OPCB, a larger primary receptacle is aligned accurately to the secondary receptacle by means of alignment stubs protruding from the secondary receptacle and subsequently bolted over the optical engagement aperture. Fig. 15(a) shows an optical engagement aperture with just a secondary receptacle assembled and one with both a secondary and primary receptacle

Fig. 14. Receptacle assembly jig required to register lensed receptacle with waveguide alignment features and hold it in place during adhesive curing process.


Fig. 15. (a) Assembly of primary receptacle and secondary receptacle in optical engagement slots, (b) Pluggable optical connector prior to engagement with OPCB. Parallel MT compliant optical interface is visible as connector starts to engage with connector receptacle.

assembled. The purpose of the primary receptacle is to provide coarse mechanical alignment of the OPCB connector during the mating process, while the secondary receptacle provides the precise optical alignment of the connector interface with the waveguide interface. Fig. 15(b) shows an OPCB connector plug during the mating process about to engage with the primary connector receptacle mounted on the opposite side of the board.

IV. PLUGGABLE ACTIVE ELECTRO-OPTICAL MIDPLANE CONNECTOR

A. Optical Midplane Board Connection

As the midplane and its peripherals are connected in a mutually orthogonal way, Xyratex developed an in-plane pluggable connector technology and connection scheme [33] whereby the optical interfaces of optical transceiver modules housed on the mating edge of the peripheral cards can be butt-coupled to optical channels embedded on the midplane. This builds on the connection methodology demonstrated in our previous research initiatives [34] in which the optical axis of the peripheral transceiver module is collinear with the OPCB embedded optical channels, thus, eliminating the need for right-angled mirrors and minimizing the number of boundaries incurring optical loss. Xyratex developed a prototype active pluggable connector to allow optical connection between the peripheral line cards and the optical layer embedded in the midplane. The connector comprised a parallel optical transceiver, connector housing, and a pluggable engagement mechanism.

B. Parallel Optical Transceiver

Xyratex developed a quad parallel optical transceiver circuit to accommodate the proprietary connection technique. Fig. 16 provides an annotated view of the quad transceiver circuit mounted on a partially flexible and rigid substrate. The transceiver circuit comprised three sections: a base section [see Fig. 16(e)], a flexible bridge section [see Fig. 16(d)], and a moveable optical platform [see Fig. 16(c)]. The circuit was constructed on a flexible laminate substrate, which was reinforced with rigid FR4 layers in the base section and optical platform leaving the intermediary bridge section flexible. The base section allowed for the electrical connection of the transceiver to the peripheral line card by means of a high-speed electronic array connector provided by Samtec capable of supporting data rates up to 11 Gb/s.

A microcontroller [see Fig. 16(f)] provided a two-wire control interface through which various parameters of the transceiver could be externally controlled or monitored including laser bias and modulation currents, receiver squelch, signal detect, and IC temperature read back. The design allowed for multiple transceivers to operate as slave devices on a single two-wire communications bus.

The moveable optical platform contained a quad VCSEL array [see Fig. 16(i)], a VCSEL driver array [see Fig. 16(k)], a PIN photodiode array [see Fig. 16(j)] and a quad transimpedance amplifier/limiting amplifier (TIA/LA) [see Fig. 16(l)]. The uncooled VCSELs emitted at a nominal wavelength of 850 nm with a full beam divergence of 24° giving rise to an NA of 0.21, which is 78% of the NA of the midplane waveguides. The PIN photodiodes were responsive to the same wavelength and had a circular receive aperture of 70 μm diameter, which was chosen to be large to maximize misalignment tolerances and to reduce modal or speckle noise. The PIN photodiodes had a nominal responsivity of 0.62 A/W and a bandwidth of 7.8 GHz. Zarlink provided the quad VCSEL arrays, VCSEL driver arrays, TIA arrays, and PIN photodiode arrays. Two MT pins were assembled into the optical interface platform. The VCSEL and photodiode arrays were attached to a lead frame, which included two MT compliant guide slots. These were used to guide the lead frame over the pins and thus, accurately align the VCSEL and photodiode arrays to the protruding MT pins. An MT compliant 12 lens microlens array [see Fig. 16(h)] with two MT guide holes was aligned along the MT pins protruding from the optical platform, over

Fig. 16. Parallel optical transceiver circuit and optical interface (a) Front ramped plug. (b) Optical platform guide pins. (c) Optical interface. (d) Flexible bridge. (e) Rigid base section. (f) Microcontroller. (g) MT compatible alignment pins. (h) Microlens array. (i) Quad VCSEL array. (j) Quad PIN photodiode array. (k) Quad VCSEL Driver array. (l) Quad TIA/LA array.

the active emitting and receiving apertures of the VCSEL and photodiode arrays and attached to the lead frame with UV curable adhesive. The VCSEL and photodiode arrays were spaced apart on the lead frame such that the VCSEL emitting apertures were aligned to the four microlenses on the left-hand side and the photodiode receiving apertures were aligned to the four microlenses on the right-hand side (as seen from above), with the central four lenses remained unused. Thus, eight channels (four transmit and four receive) are sustained on this single row MT form factor. In future, it would be easy to incorporate higher channel numbers into the same form factor such as 12 (single row), 24 (double row), or 48 (quad row) channel arrays.

Each of the eight channels was capable of sustaining a data rate of 10.3125 Gb/s giving rise to an aggregate bandwidth of 82.5 Gb/s. The optical connection interface comprised a collimating 1×12 microlens array and a pair of mechanical registration pins, designed to be compliant with MT style parallel optical interfaces. This allowed for standalone testing of both the transceiver and the midplane waveguides with an MT terminated fiber-optic patch-cord. The flexible bridge section allowed the optical platform to "float" relative to the peripheral device, thus, ensuring that when coupled to the midplane, the optical connection remained relatively impervious to transient movements and vibrations in the system. In order to maximize the mechanical flexibility of the flexible bridge, slots were cut in the substrate between the high-speed differential pairs [see Fig. 16(d)]. It also has the benefit of reducing the electromagnetic crosstalk between the tracks. The circuit was mounted into a metallic connector housing [see Fig. 18(b)], which included grooves to enable the required movement of optical interface during the pluggable connection process described in the following. The transceiver included guide pins on the sides of the


Fig. 17. Dual geometric lens coupling interface between transceiver and OPCB.

Fig. 18. (a) Optical transceiver circuit mounted on a flexi-rigid substrate. (b) Connector module housing. (c) Grooves to enable required movement of optical interface during mating process.

optical platform [see Fig. 16(b)], which engage with slots in the connector housing to support the connection process described later. The transceiver also included a ramped plug [see Fig. 16(a)] on the front of the optical platform, which engages with the primary receptacle in the midplane in the first part of the connection process.

C. Dual Lens Coupling Solution

As misalignment tolerance is important to minimize the connection cost, UCL recommended that an expanded beam connector design be used. The microlens array attached to the transceiver formed one half of a dual lens arrangement, and the second microlens array was part of the secondary receptacle, assembled over the waveguide interface on the midplane (see Fig. 17). When coupled, the transceiver lens array in combination with the midplane lens array served to image the VCSEL output into the midplane waveguide and the waveguide output into the photodiode aperture.

Simulations were carried out by Omron using Zemax ray tracing software to select the optimum configuration of coupled microlenses required to image the output of the transceiver VCSEL onto the launch facet of the waveguide and the output of the waveguide onto the receiving aperture of the photodiode. It

Fig. 19. Optical transceiver and OPCB with connector receptacle. (a) Open position. (b) Engaged position.

Fig. 20. Peripheral test card and pluggable optical connector attached to an OPCB. Inset is a waveguide illuminated with 635 nm laser for aid of visualization.

was determined that the best way to achieve this with minimum insertion loss was to deploy a microlens array P1L12A-F1 over the transceiver interface and a microlens array P1L12A-C1 over the waveguide interface as shown in Fig. 17. A total insertion loss of 0.72 dB was calculated over the dual lens interface between the VCSEL and the waveguide, while a total insertion loss of 1.11 dB was calculated between the waveguide and the photodiode.

The free space distance between the VCSELs, photodiodes, waveguides, and their respective lens arrays was chosen to ensure that, at the point of interface between the two lenses, the optical beam was expanded to a width many times that of the source width, whether the source was the VCSEL or the waveguide. An expanded beam width of 105 μm was predicted at the optical mating interface between the VCSEL and waveguide, while a 195 μm beam width was predicted at the optical mating interface between the waveguide and the photodiode. One crucial benefit of this arrangement was to make the connector far less susceptible to contamination as any stray contaminants that settle on the lens interface would block a smaller proportion of the expanded beam than they would a beam of similar size to the sources.

D. Pluggable Connector Mechanism

As shown in Fig. 18, the transceiver circuit [see Fig. 18(a)] was assembled into a connector housing [see Fig. 18(b)] of size 89.5 mm × 34 mm × 12 mm, wherein the two lateral guiding pins, which form part of the transceiver optical platform, were slotted into compliant grooves in the housing [see Fig. 18(c)]. This enabled the controlled movement of the optical platform relative to the housing as required during the two-stage pluggable connection process, which is described as follows: as the peripheral line card is first inserted into the midplane, the ramped plug at the front end of the transceiver is funneled into the larger primary midplane receptacle and the transceiver lens array moved into position under the midplane lens array housed in the secondary receptacle [see Fig. 19(a)].

As the connector is then pushed further into the larger midplane receptacle, the lateral guiding pins on the optical platform are guided along the grooves in the connector housing, which are angled such as to move the transceiver lens array toward the midplane lens array. The MT pins on the optical platform then engage with the MT compliant slots in the secondary receptacle aligning both lens arrays to each other with a high degree of precision [see Fig. 19(b)]. When the peripheral line card is extracted, the connection process is reversed.

V. Electro-Optical Midplane Demonstration Platform

To evaluate the viability of these technologies, Xyratex constructed a demonstration platform, which comprised a 10 U (445 mm) high Compact PCI chassis with a single board computer, an electro-optical midplane, and four peripheral test cards, each housing a pluggable optical connector.

A. Peripheral Test Boards

The peripheral test cards (see Fig. 20) were designed to relay external 10.3125 Gb/s 10 GbE LAN test data to each other optically across the midplane through the pluggable connectors. Each test card included a reconfigurable crosspoint switch to map test data from four commercial 10 Gigabit Small Form Factor Pluggable (XFP) ports on the front edge to the transceiver housed in the connector on the peripheral test card (back) edge. The switch also supported multicasting, whereby test data on any of its inputs could be copied to multiple outputs. This way, one external test stream could be mapped to all four VCSEL transmitters in the connector simultaneously allowing it be characterized while fully stressed. A field-programmable gate array was present on the board to allow user communication with the XFPs, crosspoint switch, and FirstLight transceiver. A PCI bridge chip allowed a user communications interface to be established between the single board computer and all the line cards via the electrical Compact PCI bus and connectors on the electro-optical midplane.

A GUI was developed to run on the operating system on the single board computer and provide selective user access to each peripheral line card connected to the midplane. The GUI allowed the user to configure the XFP and crosspoint switch parameters on each line card, in addition to giving full control over the parallel optical transceiver parameters.

B. High-Speed Data Transmission Across Demonstration Platform

Fig. 21(a) shows the fully assembled demonstration platform. An external Xyratex proprietary 10 Gb Ethernet LAN traffic source was arranged to convey a 10.3125 Gb/s test data stream along a fiber-optic cable to one of the commercial XFP devices

Fig. 21. (a) Demonstration platform fully assembled. (b) Eye diagrams showing 10.3125 Gb/s optical signals received from demonstration platform on nine waveguide links under test.

on the front edge of a peripheral test card in the demonstration platform. The XFP device converted the input optical data stream to a serial electronic data stream on the test card, which was then mapped by the crosspoint switch to one of the VCSEL transmitters in the connector attached to that card and reconverted into an optical data stream. As the connector was optically engaged to the midplane, the optical data stream was launched into a waveguide and conveyed to the receive element of another connector on a different test card in the chassis. The data were then converted to a serial electronic data stream, mapped to an XFP port on that test card and reconverted to an optical data stream on the output of the XFP device. Finally, a fiber-optic cable was connected between the XFP output port and a Tektronix CSA8000B communication signal analyzer where the test data were characterized.

In total, nine waveguides were tested as described and 10.3125 Gb/s test data were successfully conveyed between all test cards and their prototype connectors with an acceptable level of signal recovery. Fig. 21(b) shows the eye diagrams corresponding to the nine waveguides under test. The average total peak-to-peak jitter on the nine links including front end clock and data recovery through the exit XFP was measured to be 28.217 ps or 0.29 unit interval (UI), which is within the jitter thresholds for both transmitter input (0.61 UI) and receiver output (0.363 UI) as specified by the XFP MSA [35]. In addition, we carried out BER tests on the nine waveguides directly using an Anritsu signal analyzer MT1810 and an SFP+ driver and receiver unit and applied index matching fluid at both the launch and exit facets of the waveguides. The SFP+ unit [36] had the same sensitivity as the XFP model employed in the system, but excluded the clock recovery unit, so that the quality of the raw signal could be measured. A pseudorandom binary sequence $2^{12} - 1$ pattern length Ethernet LAN traffic 10.3125 Gb/s bit rate was generated by the MT1810 and was used to used to modulate the SFP+ transceiver. The optical signal was guided by a 50/125 μm step-index MM fiber with NA of 0.22, which was butt coupled to a waveguide channel on the prototype midplane. The output from the waveguide was captured by another MM fiber, which was connected to the receiver port of the SFP+ unit on the BER tester. We were able to test individual channels with this arrangement. An error rate of less than 10^{-12} was achieved through each of the nine waveguides under test.

VI. DISCUSSION AND FUTURE WORK

The FirstLight connector is the first active pluggable connector demonstrated to establish a midboard optical connection to MM waveguides in an OPCB. Though the form factor of the connector prototype developed is large, consuming an impractical amount of space on both the edge of the plugging line card and on the midplane, this prototype could be used to satisfy the interconnect requirements of a basic 24 drive storage array system [see Fig. 1(a)] in the following way. Due to the use of an MT compliant interface, 12 singlex channels and their drive electronics could be substituted for the four duplex channels deployed in the FirstLight prototype without changing the size of the optical interface. A maximum of four connectors (two singlex transmitter connectors and two singlex receiver modules) could fit on the edge of a standard data storage controller module in addition to the standard connectors required for power and low speed signals, allowing 48 duplex channels to be supported per controller and thus, 96 duplex channels on the midplane. This in turn would be sufficient to provide full interconnection to 24 drives in an enclosure. In its current form, however, such a form factor would be too large to lend itself to the scalability required in modern data storage systems. Therefore, research and development efforts into active and passive pluggable connector schemes are currently underway to significantly increase the density of optical links in new edge and midboard OPCB connectors. A data rate of 10.3125 Gb/s was demonstrated prior to the availability of 12 Gb/s SAS compliant components. The photonic subcomponents assembled in the FirstLight connector were operational up to 11.6 Gb/s and, therefore, the connector could be deployed into data storage systems to carry SAS 2 data operating at 6 Gb/s. Furthermore, subcomponents are now available matching the same form factor and power dissipation as those deployed in the FirstLight connector, which can accommodate serial data rates of over 12 Gb/s; therefore, the connector can be easily upgraded to accommodate SAS 3 data operating at 12 Gb/s. By the end of 2012, compliant photonic subcomponents able to convey data rates over 25 Gb/s are expected to be commercially available.

In-plane bends are an inextricable requirement in complex waveguide layouts on OPCBs; however, bend radii as large as the minimum bend radius of 17 mm deployed on the First-Light midplane would be difficult to accommodate within a high density PCB layout, especially where high numbers of optical waveguides are involved. Research into novel waveguide structures allowing a reduction in bend loss and corresponding reduction in minimum bend radius is currently underway to address this issue [29].

The timescale of commercial deployment of OPCB technologies will be gated by the comparative cost of either implementing high-speed interconnects electronically or optically. In particular, the commoditization and associated cost reduction of optical transceiver subassembly technology will play a key role in determining whether a cost transition point can be reached during the lifetime of, for instance, the 12 Gb/s SAS protocol, which is currently being introduced or whether it will have to wait for the transition to 24 Gb/s SAS in 2016.

VII. Conclusion

The burgeoning demand for faster and more compact data storage systems is fuelling the need for embedded optical interconnect solutions. To address this, we have developed and successfully demonstrated an active pluggable optical PCB connector solution, which will allow peripheral devices to plug into and unplug from an electro-optical midplane with embedded MM polymer waveguides. We have developed and successfully implemented a technique to passively align optical devices to such waveguides with high precision. Our results have shown that a complex optical interconnect pattern of polymer waveguides can be deployed across a densely populated 6U board with error-free transmission.

References

[1] SearchStorage.com, "Data storage trends 2011: Predictions of hot data storage technologies," 2011 [Online]. Available: http://searchstorage.techtarget.com/Data-storage-trends-2011-Predictions-of-hot-data-storage-technologies

[2] C. Barrera and S. Wojtowecz, "Five storage trends for 2011," Computer World UK 2011 [Online]. Available: http://www.computerworlduk.com/in-depth/infrastructure/3258578/five-storage-trends-for-2011/

[3] J. Gantz and D. Reinsel, "Extracting value from chaos," *IDC iView*, pp. 1–12, 2011.

[4] H. Cho, K.-H. Koo, P. Kapur, and K. C. Saraswat, "The delay, energy, and bandwidth comparisons between copper, carbon nanotube, and optical interconnects for local and global wiring application," in *IEEE Int. Interconnect Technol. Conf.*, 2007, pp. 135–137.

[5] F. Gisin and G. Dudnikov, "State of the art of electrical high speed backplanes in industry today and the transition to optical interconnects," in *Proc. Biophoton./Opt. Interconnects VLSI Photon./WBM Microcavities, Dig. LEOS Summer Top. Meet.*, 2004, vol. 1, pp. 7–8.

[6] W. Heirman, J. Dambre, I. Artundo, C. Debaes, H. Thienpont, D. Stroobandt, and J. V. Campenhout, "Predicting the performance of reconfigurable optical interconnects in distributed shared-memory systems," *Photon Netw. Commun.*, vol. 15, pp. 25–40, 2008.

[7] J. A. Conway, S. Sahni, and T. Szkopek, "Plasmonic interconnects versus conventional interconnects: A comparison of latency, crosstalk and energy costs," *Opt. Exp.*, vol. 15, no. 8, pp. 1069–1071, 2007.

[8] K.-H. Koo, H. Cho, P. Kapur, and K. C. Saraswat, "Performance comparisons between carbon nanotubes, optical, and Cu for future high-performance on-chip interconnect applications," *IEEE Trans. Electron Devices*, vol. 54, no. 12, pp. 3206–3215, Dec. 2007.

[9] "Storage bridge bay (SBB) specification," Storage Bridge Bay Working Group Inc. 2008 [Online]. Available: http://www.sbbwg.org

[10] "Enterprise SSD form factor," SSD Form Factor Working Group 2011 [Online]. Available: http://www.ssdformfactor.org

[11] "Serial attached SCSI master roadmap," 2011 [Online]. Available: http://www.scsita.org/sas_library/6gbs-sas/roadmap-1

[12] W. T. Beyene and C. Yuan, "An accurate transient analysis of high-speed package interconnects," *Analog Integr. Circuits Signal Process.*, no. 35, pp. 107–120, 2003.

[13] M. M. Pajovic, J. Yu, Z. Potocnik, and A. Bhobe, "GigaHertz-range analysis of impedance profile and cavity resonances in multilayered PCBs," *IEEE Trans. Electromagn. Compat.*, vol. 52, no. 1, pp. 179–188, Feb. 2010.

[14] H. W. Johnson, "High-speed backplane connectors," in *2011 IEEE Int. Symp. Electromagn. Compat.*, Aug. 2011, pp. 612–618.

[15] W.-T. Huang, C.-H. Lu, and D.-B. Lin, "Suppression of crosstalk using serpentine guard trace vias," *Progr. Electromagn. Res.*, vol. 109, pp. 37–61, 2010.

[16] R. Ramzan, J. Fritzin, J. Dabrowski, and C. Svensson, "Wideband low-reflection transmission lines for bare chip on multilayer PCB," *ETRI J.*, vol. 33, no. 3, pp. 335–343, Jun. 2011.

[17] H. F. Lee, C. Y. Chan, and C. S. Tang, "Embedding capacitors and resistors into printed circuit boards using a sequential lamination technique," *J. Mater. Process. Technol.*, vol. 207, no. 1–3, pp. 72–88, Oct. 2008.

[18] F. E. Doany, C. L. Schow, C. W. Baks, D. M. Kuchta, P. Pepeljugoski, L. Schares, R. Budd, F. Libsch, R. Dangel, F. Horst, B. J. Offrein, and J. A. Kash, "160 Gb/s bidirectional polymer-waveguide board-level optical interconnects using CMOS-based transceivers," *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 345–359, May 2009.

[19] L. Schares, J. A. Kash, F. E. Doany, C. L. Schow, C. Schuster, D. M. Kuchta, P. Pepeljugoski, J. M. Trewhella, C. W. Baks, R. A. John, L. Shan, Y. H. Kwark, R. A. Budd, P. Chin,walla, F. Libsch, J. Rosner, C. K. Tsang, C. S. Patel, J. D. Schaub, R. Dangel, F. Horst, B. Offrein, D. Kucharski, D. Guckenberger, S. Hegde, H. Nyikal, C.-K. Lin, A. Tandon, G. R. Trott, M. Nystrom, D. P. Bour, M. Tan, and D. Dolfi, "Terabus: Terabit/second-class card-level optical interconnect technologies," *J. Quantum Electron.*, vol. 12, no. 5, pp. 1032–1044, Sep./Oct. 2006.

[20] R. Dangel, C. Berger, R. Beyeler, L. Dellmann, M. Gmur, R. Hamelin, F. Horst, T. Lamprecht, T. Morf, S. Oggioni, M. Spreafico, and B. Offrein, "Polymer-waveguide-based board-level optical interconnect technology for datacom applications," *IEEE Trans. Adv. Packag.*, vol. 31, no. 4, pp. 759–767, Nov. 2008.

[21] S. Uhlig, "Micro-processing: Applied to optical interconnects and high-frequency packaging,".

[22] P. Lukowicz, J. Jahns, R. Barbieri, P. Benabes, T. Bierhoff, A. Gauthier, M. Jarczynski, G. A. Russell, J. Schrage, W. Sullau, J. F. Snowdon, M. Wirz, and G. Troster, "Optoelectronic interconnection technology in the HOLMS system," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 624–635, Mar. 2003.

[23] N. Bamiedakis, J. Beals, R. V. Penty, I. H. White, J. V. DeGroot, and T. V. Clapp, "Cost-effective multimode polymer waveguides for high-speed on-board optical interconnects," *IEEE J. Quantum Electron.*, vol. 45, no. 4, pp. 415–424, Apr. 2009.

[24] I. Papakonstantinou, K. Wang, D. R. Selviah, and F. A. Fernández, "Transition, radiation and propagation loss in polymer multimode waveguide bends," *Opt. Exp.*, vol. 15, no. 2, pp. 669–679, 2007.

[25] K. Wang, I. Papakonstantinou, D. R. Selviah, H. Baghsiahi, and G. Yu, "Design rules for polymer multimode waveguide interconnects and design of waveguide layout for a highly constrained electrical-optical backplanes," *Opt. Exp.*, to be published.

[26] I. Papakonstantinou, D. R. Selviah, R. C. A. Pitwon, and D. Milward, "Low-cost, precision, self-alignment technique for coupling laser and photodiode arrays to polymer waveguide arrays on multilayer PCBs," *IEEE Trans. Adv. Packag.*, vol. 31, no. 3, pp. 502–511, Aug. 2008.

[27] K. Wang, D. R. Selviah, I. Papakonstantinou, H. Baghsiahi, and F. A. Fernandez, "Photolithographically manufactured acrylate polymer multimode optical waveguide loss design rules," in *Proc. 2nd Electron. System-Integr. Technol. Conf.*, 2008, pp. 1251–1256.

[28] "Datasheet: 10 Gbps dual rate 850 nm multi-mode XFP transceivers," MRV Communications, Inc. 2010 [Online]. Available: www.mrv.com

[29] R. C. A. Pitwon, C. Smith, K. Wang, J. Graham-Jones, D. R. Selviah, M. Halter, and A. Worrall, "Polymer optical waveguides with reduced in-plane bend loss for electro-optical PCBs," in *Proc. SPIE*, 2012, vol. 8264.

[30] H. Baghsiahi, K. Wang, W. Kandulski, R. Ferguson, and D. R. Selviah, "Reduction of the end facet roughness of photo lithographically fabricated polyacrylate multimode optical waveguide," *Opt. Exp.*, to be published.

[31] I. D. Johnson, R. C. A. Pitwon, D. R. Selviah, and I. Papakonstantinou, "Optical printed circuit board and manufacturing method," U.S. Patent WO2007/0101842006.

[32] D. R. Selviah, F. Fernandez, I. Papakonstantinou, K. Wang, H. Bagshiahi, A. C. Walker, A. McCarthy, H. Suyal, D. A. Hutt, P. P. Conway, J. Chappell, S. S. Zakariyah, and D. Milward, "Integrated optical and electronic interconnect printed circuit board manufacturing," *Circuit World*, vol. 34, no. 2, pp. 21–26, 2008.

[33] R. C. A. Pitwon, K. Hopkins, and K. McPherson, "Optical connector, a communication system and a method of connecting a user circuit to an optical transceiver," U.S. Patent WO2006/129069.

[34] R. C. A. Pitwon, K. Hopkins, D. Milward, M. Muggeridge, D. R. Selviah, and K. Wang, "Passive assembly of parallel optical devices onto polymer-based optical printed circuit boards," *Circuit World*, vol. 36, no. 4, pp. 3–11, 2010.

[35] 10 Gigabit Small Form Factor Pluggable Module , INF-8077i, 2005.

[36] "Datasheet: 10 Gbps 850 nm multi-mode SFP+ transceiver," MRV Communications, Inc. 2010 [Online]. Available: www.mrv.com

Richard C. A. Pitwon received the B.Sc. (Hons.) degree in physics from the University of St. Andrews, St. Andrews, U.K., in 2000.

He is currently a Lead Photonics Engineer at Xyratex, Hampshire, U.K., with more than 12 years experience in the design and development of high-speed photonic interconnect technologies including passive and active optical connectors, optical printed circuit boards, optical interconnect interfaces, and transceivers. He holds 21 patents in the field of embedded optical interconnect and has authored numerous conference, journal, and commercial publications in this area.

Dr. Pitwon is a Chartered Engineer with the Institute of Engineering and Technology and represents the U.K. on International Electrotechnical Commission optical interconnect standards committees.

Kai Wang (S'02) received the B.Sc. degree in electronic engineering and telecommunications from Shenzhen University, Shenzhen, China, in 2001, and the M.Sc. degree in microelectronic systems and telecommunications from Liverpool University, Liverpool, U.K., in 2002.

In 2003, he began Ph.D. research in the Department of Electronic and Electrical Engineering, University College London (UCL), where he became a Research Fellow in 2004. His current research interests include computer modeling in LCDs, backlights, multimode waveguides, and design optical printed circuit boards.

Jasper Graham-Jones received the Ph.D. degree from Southampton University, Southampton, U.K.

He has previously worked with the Open University, Reading University, and Portsmouth University, where he was involved in multiple research and consultancies. He is currently an Associate Professor in mechanical engineering with Plymouth University, Plymouth, U.K. His research interests include failure mechanisms and manufacturing.

Ioannis Papakonstantinou (S'02–M'09) received the Ph.D. degree in optical backplane interconnections from University College London (UCL), London, U.K., in 2008.

He later joined Sharp Laboratories of Europe, Oxford, U.K., to investigate nanoimprinted subwavelength gratings for multifunctional LCD backlights. From 2009 to 2011, he was with the Electronic Systems for Experiments Group, Department of Physics, European Organization for Nuclear Science. In 2011, he became a Lecturer in the Department of Electronics and Electrical Engineering, UCL. His research interests include nanophotonics for solar cells and medical applications.

Hadi Baghsiahi received the Bachelor's degree in physics from Shiraz University, Shiraz, Iran, in 2002, and the M.Sc. degree in photonics and optical telecommunications from Shahid Bahonar University of Kerman, Kerman, Iran, in 2005. He is currently working toward the Ph.D. degree at University College London, London, U.K., working on optical printed circuit boards.

He is also involved in the design of a light engine for a 3-D laser-based display system.

Bert Jan Offrein (SM'12) received the M.Sc. degree in applied physics and Ph.D. degree from the University of Twente, Twente, The Netherlands, in 1990 and 1994, respectively. His Ph.D. dissertation examined the measurement of optical nonlinearities and the design and realization of an all-optical switching device.

In 1995, he joined the Optical Networking Group, IBM Research, Zurich, Switzerland, where he was involved in the design and characterization of optical components for dense wavelength division multiplexing networks in SiON glass-on-silicon waveguide technology. In 2002, his research topic changed to optical interconnects for next-generation server systems. Since June 2004, he has been managing the Photonics Group, IBM Zurich Research Laboratory, addressing optical interconnects, electro-optic assembly, silicon photonics, nonlinear optical materials, and nanophotonic devices.

Roger Dangel received the Diploma degree in physics and the Ph.D. degree from the Swiss Federal Institute of Technology in Zurich, in 1991 and 1997, respectively.

From 1998 to 2000, he was with the Swiss Center for Electronics and Microtechnology. At the end of 2000, he joined the Photonic Networks Group, IBM Zurich Research Laboratory, where he was involved in the design and characterization of optical telecom components in SiON technology. In 2002, he became part of the newly formed Optical Interconnect Team, IBM Zurich, where he is responsible for design, fabrication, and optical characterization of the polymer-based datacom waveguides.

Dave Milward received the B.Eng. degree in electrical and electronic engineering from Liverpool University, Liverpool, U.K.

He currently serves as the Director of Skills and Technical Vitality at Xyratex Technologies Ltd., Hampshire, U.K.

Mr. Milward is a Chartered Engineer and member of the Institution of Engineering and Technology.

David R. Selviah (M'01) received the M.A. degree in physics and theoretical physics and the Ph.D. degree in optical engineering from Trinity College, Cambridge University, Cambridge, U.K.

He is currently a Reader in optical devices, interconnects, algorithms and systems in the Department of Electronic and Electrical Engineering, University College London, London, U.K. He has more than 200 publications and led the £ 1.3 million Innovative electronics Manufacturing Research Centre Optical Printed Circuit Board Flagship Consortium of three universities and eight companies.

Dr. Selviah serves on the International Electrotechnical Commission standards committees.