A High-Speed Optical Multi-drop Bus for Computer Interconnections

Michael Tan, Paul Rosenberg, Jong Souk Yeo, Moray McLaren, Sagi Mathai, Terry Morris, Joseph Straznicky, Norman P. Jouppi, Huei Pei Kuo, Shih-Yuan Wang, Scott Lerner, Pavel Kornilovich, Neal Meyer, Robert Bicknell, Charles Otis, and Len Seals

Hewlett Packard Company
E-mail: mike.tan@hp.com

Abstract

Buses have historically provided a flexible communications structure in computer systems. However, signal integrity constraints of high-speed electronics have made multi-drop electrical busses infeasible. Instead, we propose an optical data bus for computer interconnections. It has two sets of optical waveguides, one as a fan-out and the other as a fan-in, that are used to interconnect different modules attached to the bus. A master module transmits optical signals which are received by all the slave modules attached to the bus. Each slave module in turn sends data back on the bus to the master module. Arrays of lasers, photodetectors, waveguides, microlenses, beamsplitters and Tx/Rx integrated circuits are used to realize the optical data bus. With 1mW of laser power, we are able to interconnect 8 different modules at 10 Gb/s per channel. An aggregate bandwidth of over 25 GB/s is achievable with 10 bit wide signaling paths.

1. Introduction

Multi-drop buses find many applications in computer systems, particularly in the I/O and memory systems. They allow architects to create systems that are readily expandable, with uniform access times for all devices. The ability to expand a computer's main memory by adding memory DIMMs is good example of this. For applications such as front side buses, the ability to broadcast simplifies the design of coherency protocols. Increasingly however, buses are being displaced by networks of point to point links; for example PCI migrating to PCI Express. This is driven primarily by signal integrity considerations [1]. The impedance discontinuities inherent in multi-drop electrical buses have made it impossible to scale their data rates to track improvements in processor performance [2]. In order to keep the design and provisioning of systems flexible, we would like a high-speed alternative to traditional electrical buses.

The difficulty of creating a multidrop bus electrically is easily avoided by using optics. Optics provide numerous benefits over copper: 1) its high carrier frequency of 10^{15} Hz results in no signal degradation with increased modulation frequency 2) low propagation loss (< 0.1dB/km for glass fiber) 3) broadband impedance matching (i.e. anti-reflection coatings, beamsplitters) with low loss 4) EMI immunity and 5) low power consumption. Moreover, optical interconnects have higher density than copper interconnects at high data rates.

2. Multi-drop optical bus

A master-slave parallel optical bus is shown in Figure 1. It is comprised of two unidirectional signal buses to which four or more modules are attached. The master module broadcasts signals on the bus where they can be received by any module. Data is sent back to the master from any module. Two sets of parallel optical waveguides, one as a fan-out (10 bits wide) and the other as a fan-in (10 bits wide), are used to interconnect the different modules attached to the bus. The master module

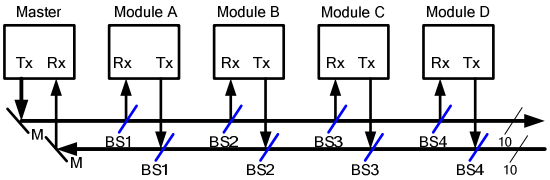

Figure 1: An optical bus with four slave modules shown.

transmits (Tx) optical signals which are received (Rx) by all of the slave modules attached to the bus. Optical beamsplitters (BS1… BS4) at each of the "T" junction interconnection nodes are used to tap power from the bus. Each beamsplitter has a different reflectivity R_i and transmissivity T_i, chosen so as to distribute the transmitted optical power equally among the interconnected modules. The relationship between the reflectivities of the beamsplitters for the case of *no loss in the beamsplitters* (i.e. $T_i=1 - R_i$) is given by

$$R_n = \frac{R_1}{\prod_{m=2}^{n}(1-R_{m-1})k^{m-1}} = \frac{R_1}{(1+R_1-nR_1)} \quad \text{for } n>1 \text{ and } k=1$$

where R_1 is the reflectivity of the first beamsplitter, n is the position of the tap along the bus and $k=e^{-\alpha L}$ takes into account the propagation loss of distance L between taps. Excess loss from the splitters can also be included in the k factor. The reflectivity of the first beamsplitter R_1 is

1550-4794/08 $25.00 © 2008 IEEE
DOI 10.1109/HOTI.2008.15

determined by power required by the receiver to maintain a reliable link (receiver sensitivity) with a reasonable link margin. Each slave module in turn sends data back on the fan-in bus to the master module by means of identical beamsplitter junctions. On the return path, beamsplitter ratios are selected to ensure that the same optical power is received from each of the slave modules no matter how many stages they are from the master. It is interesteing to note that the reflectivity values of the beamsplitters on the return path are the same as for the output path (see Figure 2 below).

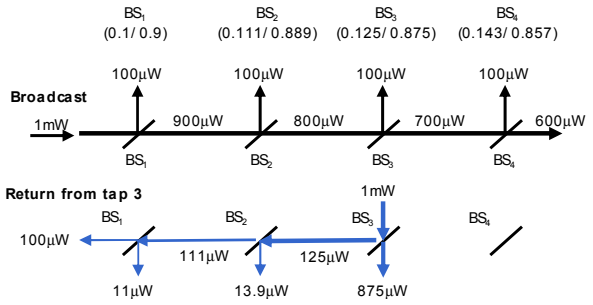

Figure 2: Mirror ratios for fan-in and fan-out for zero-loss case with four slave modules. The beam split ratios are given in parenthesis as (R_i/T_i).

The master module performs an electrical-to-optical conversion by means of a Tx transmitter IC which drives a single 1x10 850nm VCSEL array capable of 10Gb/s modulation. Light from the VCSEL array is collimated by means of a microlens array and coupled to the array of waveguides with a 90-degree turning mirror (M). The optical signal is broadcast to the different modules by means of the beamsplitters. At each of the slave modules, the tapped optical signal is focused onto to high speed GaAs PIN photodetector arrays by means of another identical microlens array and converted back to electrical signals via a Rx receiver IC. Data from each of the slave modules is sent back to the master module via an identical set of optical and electronic components. Only one slave can send data back to the master at a time. This can be ensured either by the master exactly scheduling all slave transfers, or through the addition of an arbitration mechanism between slaves. The maximum optical fan-out is determined by the Tx power, the Rx receiver sensitivity and the associated losses in the system. At 10Gb/s, the receiver sensitivity of commercially available transimpedance amplifier (TIA) receivers is on the order of -14dBm (40µW). The major system losses stem from waveguide propagation loss, coupling loss and beamsplitter losses. Table 1 shows a detailed power link budget for the optical bus. A waveguide loss of -0.1dB/cm with a propagation length of 3cm and an excess loss of about -0.3dB at the splitters were assumed. A receiver sensitivity of -14dBm was used. An additional 0.5dB of coupling loss into the receiver was assumed. With this amount of transmit power, receiver sensitivity, and loss the maximum fan-out is 7. Increasing the input power from +0.6dBm to +1dBm, will increase the link margin to ~3.45dB per tap. While decreasing the input power to -0.5dBm will decrease the link margin to ~1.95dB. If the propagation and excess loss per tap are reduced to -0.2dB and -0.25dB, respectively, we will be able to increase the optical fan-out to 8 with a link margin of 3dB. The beamsplitter power reflectivities for this case would be 0.082, 0.10, 0.123, 0.155, 0.204, 0.284, 0.439 and 0.867. Thus, with about 1mW of laser power, we are able to interconnect 7-8 modules at 10 Gb/s per channel. An aggregate bandwidth of over 25 GB/s is achievable with a 10 bit wide fan-in and fan-out bus.

Table 1 *Link budget for a 1x7 fan-out*

Tap number	1	2	3	4	5	6	7
BS Reflectance	0.09	0.11	0.15	0.20	0.28	0.45	0.96
Tap Input Pwr (dBm)	+0.6	-0.4	-1.5	-2.8	-4.4	-6.4	-9.7
Loss (dB)	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3
Excess loss (dB)	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3
Reflected Pwr (dBm)	-10.5	-10.5	-10.5	-10.5	-10.5	-10.5	-10.5
Transmit Pwr (dBm)	-0.4	-1.5	-2.8	-4.4	-6.4	-9.7	-25
Rx coupling loss	-0.5	-0.5	-0.5	-0.5	-0.5	-0.5	-0.5
Rx Pwr (dBm)	-11	-11	-11	-11	-11	-11	-11
Link Margin (dB)	3	3	3	3	3	3	3

3. Competing Technology

For bus applications at high date rates, one option in the electronic domain is to use daisy chains of point-to-point links (p2p). The FBDIMM memory-module standard is an example of this technique applied to memory systems. This approach has two significant disadvantages. First, the access latency increases proportionally to the number of "hops". Although the memory modules are designed to have a very low pass-through latency, the round-trip delay starts to approach the base memory-access time for systems with 4 or more FBDIMMs [3]. Second, there is a significant power penalty for having to retransmit the data at each intermediate FBDIMM module. Power is an increasing concern in the design of all classes of computer systems[4]. The Advance Memory Buffer (AMB) in the FBDIMM module adds 6W to the power of each memory DIMM irrespective of whether DRAMS on that

DIMM are accessed. It would be possible to implement the FBD protocol using point to point optical links at each stage of the outbound and return daisy-chains. However, although this would allow longer lengths it would also increase the overall delay, since the delay of the optical-to-electrical and electrical-to-optical conversions would occur on every hop, as opposed to once in each direction for the outbound and return communications with the broadcast structure. This arrangement also requires twice the number of optical-electrical converters and twice the power for an equivalent bandwidth. Modifying the bus protocol to use a single ring would yield a solution that used the same number of optical-electronic converters as the optical bus, but in this case the maximum throughput would be halved as there are no longer independent outbound and return datapaths.

The IBM/Agilent Terabus [5] is a complete optical system-interconnect solution, consisting of optical transcievers, package-to-board connectors, and waveguide technology that can be integrated with standard FR4 board-manufacturing processes. However, Terabus links are point-to-point, and the relatively high losses of the polymer-waveguide technology employed means that it would be extremely difficult to add a broadcast capability with a reasonably high fan-out [6]. Optical buses for computer interconnections have been studied previously [7]. However, most results to date have been limited to low data rates and or low number of fan-out, fan-in nodes due to high system loss [8,9,10,11].

4. Novel Low-Cost Technologies

Two novel technologies are used to realize the optical bus. These are 1) hollow metal waveguides (HMWG) [12] and 2) non-polarizing pellicle beamsplitters (NPPBS). These components will be described here in detail. HMWGs have several interesting properties that make them ideal candidates for use in intra-board interconnections. They are: 1) low propagation loss < 0.05dB/cm, 2) ease of fabrication, 3) low numerical aperture NA< 0.01, and 4) an effective index of ~1. The unity effective index yields zero skew between waveguide channels and also the lowest latency (0.033ns/cm) per unit length. These HMWGs are air-core light pipes with a rectangular cross section about the size of a human hair, having a high reflectivity coating in the interior. Unlike conventional millimeter-wave waveguides where the electromagnetic radiation is guided along the conductive metallic walls, the light is guided in the interior cross section of the hollow metal pipe, as the metal no longer acts as good conductor but as a high index layer. These modes can be described as rays which bounce back and forth at near grazing incidence to the metal walls as they propagate down the waveguide. A lens must be used to collimate (NA<0.02) the input beam launched into the waveguide so as to excite only these low loss modes.

The low numerical aperture of the hollow metal waveguide allows the insertion of taps or beamsplitters in the HMWG with little excess loss. This is done by inserting the beamsplitters at 45 degrees into slots or gaps cut into the waveguides. In Figure 3 is shown a Monte-Carlo ray-trace simulation of the excess loss as a function of gap separation between two 150μm core HMWGs. A 1mm gap between two 150μm core HMWG introduces an excess loss of about -0.34dB.

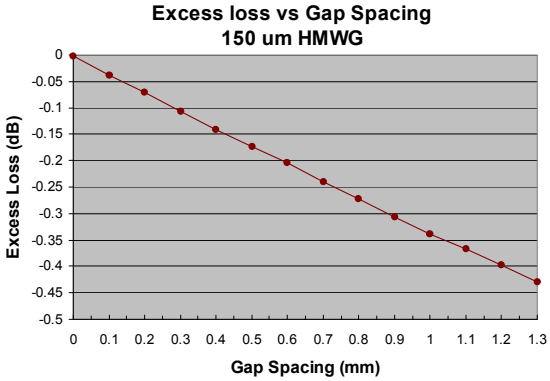

Figure 3 *Calculated excess loss dependence with gap spacing between two 150 μm HMWGs.*

The theoretical loss of these hollow metal waveguides was estimated by Marcatili[12] for circular metallic waveguides. For square waveguides whose core dimension a is much larger than the wavelength λ, the attenuation can be approximated by [13]

$$\alpha = \frac{\lambda^2}{(2a)^3}\,\mathrm{Im}\left(\frac{1+n_{clad}^2}{2\sqrt{1-n_{clad}^2}}\right)$$

where n_{clad} is the refractive index of the metal cladding layer. From this expression, we see that the loss can be made small by choosing the waveguide dimension to be much larger than a wavelength. For a 150μm square waveguide with silver cladding (n_{clad}=0.15+i5.68), the theoretical loss is around -0.0015dB/cm @850nm. The modal dispersion of these waveguides can be simply estimated using the ray model. The difference in propagation delay between the on-axis ray and the grazing ray is given by

$$\Delta T = \frac{Ln}{c}\left(\frac{1}{\cos(\theta_{\max})} - 1\right)$$

where L is the propagation length, n is the effective index, c is the speed of light and θ_{max} is the maximum angle the ray makes with the waveguide axis. For a path length of 30cm and an NA of 0.01 (θ_{max}~1.1°), the differential delay is ~20 fsec. Also, since the refractive index of air changes with temperature by only about 1ppm/°C, the propagation delay of these waveguides will change about 1 part in 10^4

for a 100 degree change in temperature. Thus, we conclude that HMWGs are a good choice for intra-board chip-to-chip interconnections which can scale with increased bandwidth.

The other essential component of the optical bus is the optical tap (the direct analogue to an electrical stub in a transmission line). The purpose of the optical tap is to

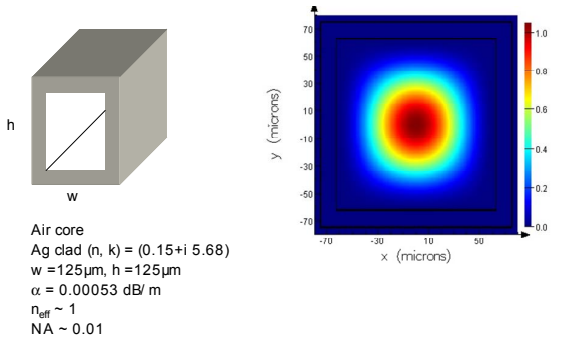

Figure 4: *Properties of Hollow Metal Waveguides for EH_{11} mode*

reflect a predetermined amount of optical power out of the waveguide while transmitting the rest thru with little excess loss. This is achieved by inserting a beamsplitter of known reflectivity and transmissivity at a 45 degree angle to the axis of the waveguide. The reflectivity of the taps must increase in a predetermined manner so that an

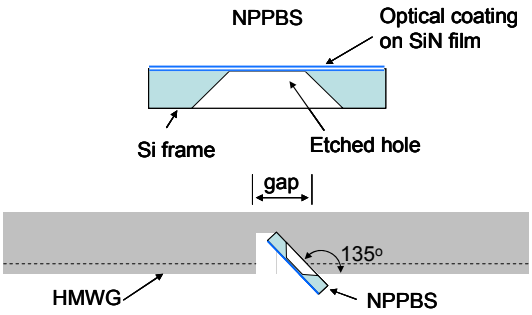

Figure 5 *Construction of non-polarizing pellicle beamsplitters and its insertion to the HMWG.*

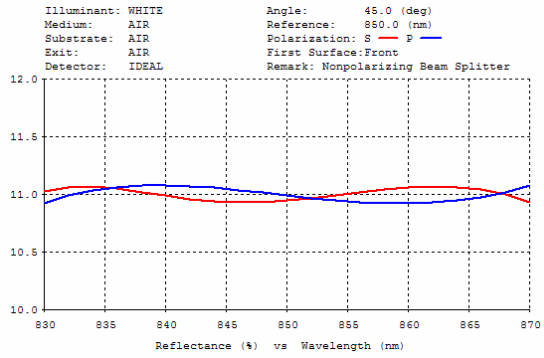

Figure 6 *S and P reflectance of a NPPBS with 15 layers of SiO_2, TiO_2 deposited on a 250nm thick SiN film.*

equal amount of power is received at each tap. Additionally, the tap should introduce little or no perturbation to the optical mode in the waveguide. Pellicle beamsplitters < 10 μm thick are preferred over other types of beamsplitters since etalon or ghosting effects (from back reflection) and, more importantly, beam walk-off are minimized. Beam walk-off is due to the refraction of the optical beam as it exits the finite thickness of the beamsplitter, resulting in the displacement of the beam from its original path. Used in a HMWG, the displaced beam would hit the metal walls introducing loss. The beam walk-off at 45 degrees for a 100μm thick BK7 glass beamsplitter is about 30μm. A non-polarizing beamsplitter is also required due to the random polarization state of the VCSEL.

The non-polarizing pellicle beamsplitters are comprised of a free standing membrane onto which a multilayer coating is deposited. The size of the pellicles is chosen to be much larger than the waveguide core requiring only precise placement of the pellicle at a 45 degree angle along the axis of waveguide. The complexity is placed on the optical coating to yield the correct reflectivity and transmissivity. NPPBS are fabricated by depositing a multilayer dielectric stack onto a thin (250nm) supporting membrane of Si_3N_4 on Silicon. The layer indices and thicknesses are designed so as to yield a desired reflectivity and transmissivity for both *s* and *p* polarizations at a 45 degree angle over a spectral bandwidth of 40nm. The coating+Si_3N_4 is released by patterning a hole on the backside of the Si wafer and etching the wafer in KOH, resulting in a free-standing ~2μm thick pellicle beamsplitter. An anti-reflection coating on the back side is obviated since the supporting Si_3N_4 film is part of the coating stack. In Figure 6 is shown the Reflectance of as a function of wavelength for both the *s* and *p* polarization of an 11% beamsplitter designed for a 45 degree angle. The layer thicknesses were obtained using a commercial optical thin-film design program [14]. The reflectance for both *s* and *p* polarizations are matched to within < ±0.5% over a 40nm wavelength range centered on 850nm. A 40nm wavelength spectral bandwidth is desired so as to relax the VCSEL requirements. The designs require accurate control of the indices and thicknesses of optical layers and the base SiN film. The residual stress of the optical coating must also be controlled so that the released pellicle remains flat. This is readily accomplished by adjusting the deposition condition to yield a stress balanced coating by alternating between compressive- and tensile-strained layers. The accuracy needed to achieve the correct beamsplitter ratio can be relaxed by providing for a larger power budget per tap.

5. Construction of the Optical bus

The optical bus is composed of a 12.5cm long 1x4 array of HMWGs with a set of optical taps spaced at

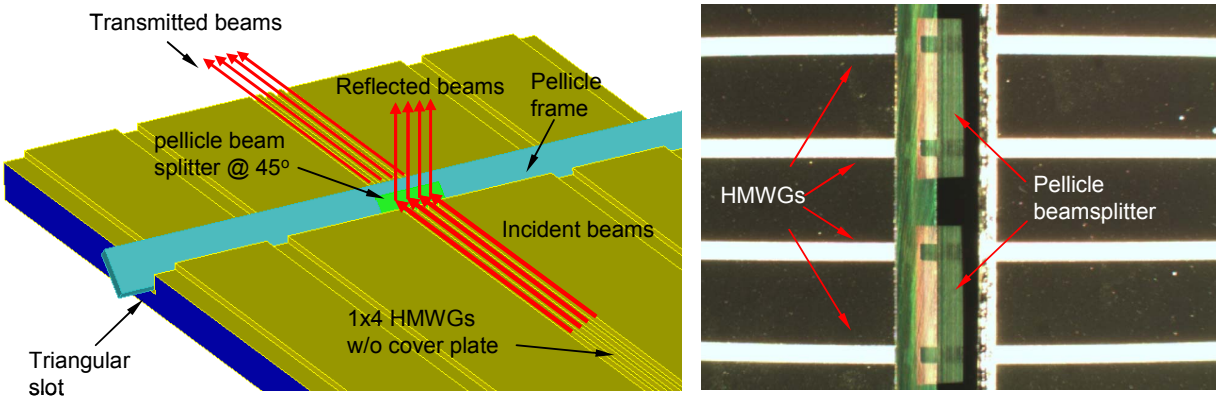

Figure 7 *Drawing and photograph of the 1x4 HMWG (no cover) with inserted PBS at 45 degree.*

1.4cm. A 1x4 850nm VCSEL array plus driver IC on a Zarlink evaluation board was used as the laser source. A corresponding 1x4 GaAs PIN optical receiver (~20μW sensitivity) on a separate Zarlink evaluation board was used as the receiver. An Agilent Parallel BERT was used to drive the VCSELs at a data rate of 5Gb/s. The output of the receiver was displayed on an Agilent Infinium DCA. The 1x4 array of HMWGs with a 1mm pitch was fabricated on a suitable substrate material (Silicon was used) by using a dicing saw to create 150μm square channels 12.5cm long. Insertion of the optical taps is accommodated by cutting 0.8mm wide by 0.59mm deep triangular slots into the waveguides using a specially designed dicing saw blade. The blade is designed to create a reference 45±0.2 degree surface onto which the pellicle beamsplitter is mounted, see Figure 7. The clear aperture of the pellicle is approximately 0.5mm x 4mm which is wide enough to encompass all 4 waveguides. For ease of handling, the finished pellicle will remain in a Silicon frame which is 14mm long x 1mm wide and 0.25mm thick. This orthogonal cut intersects the entire 1x4 waveguide array allowing a single beamsplitter to be used for the 4 waveguides. After the channels and slots are defined, a high reflectivity coating of Ag was blanket deposited over the whole substrate, coating the channel walls. A thin, Au coated, metal sheet with 0.95mm wide and 4.6mm long openings at the tap spacing was used as the cover layer. These openings allow light to be coupled out from the pellicle beamsplitters into the receiver. The first slot of the bus has a 100% mirror rotated 90 degrees with respect to the optical taps and serves as the input coupler to the bus. The last slot in the bus also has a mirror which couples all the remaining light out of the waveguides. After the mirrors and taps are inserted, the two pieces of the HMWG are clamped together in a metal frame. The finished waveguide assembly is then mounted onto the demo board together with the optic and Tx , Rx evaluation boards. Figure 8 shows the demo board assembly with the HMWG and the Tx and Rx evaluation boards. The Rx evaluation board is on a sliding rail so as to facilitate the translation of the Rx board under each tap.

A three lens bulk optic was designed to reduce the VCSEL beam NA from 0.26 to 0.065, a 4X magnification. At this magnification, the 1x4 VCSEL array with a 250μm pitch will be imaged to four spots spaced at 1mm apart, setting the waveguide pitch. The non-optimal collimation of the input beam will result in a substantial power loss due to the excitation of the higher order modes in the HMWGs. This power loss will limit the number of taps we can have along the waveguide.

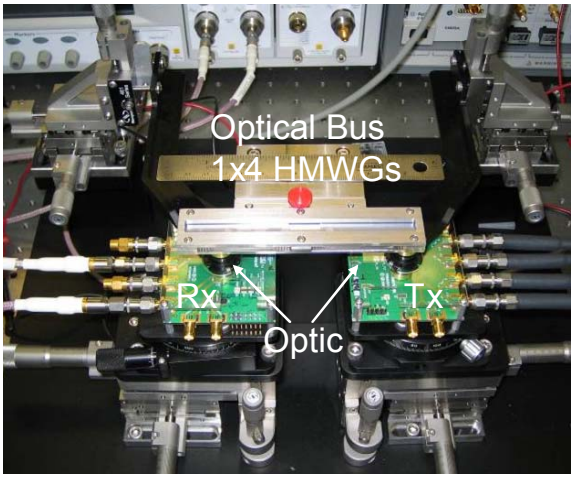

Figure 8 *Photograph of the demo board.*

Before the optical taps were inserted into the HMWGs, the propagation loss through the 12.5cm long waveguide with the 9 (0.8mm wide) slots (but without any mirrors or taps inserted) was measured to be around -0.1dB/cm. Straight waveguides without slots exhibit propagation losses of around -0.05dB/cm. These numbers include the coupling loss from the 500μm ball lens into the 150μm HMWGs. This corresponds to an excess loss per gap of around -0.15dB, assuming that each gap introduces the same loss. This value is lower than the -0.27dB predicted by ray tracing. In Figure 9 is shown a 10Gb/s eye received at the output of this multi-slotted waveguide. This clean eye opening indicates that there are no

detrimental effects from the gaps introduced into the HMWG. Polymer waveguides with this many discontinuities would exhibit prohibitive losses, making them unusable in this configuration.

The 4X magnification of the bulk optics resulted in an overall coupling efficiency of only 50% into the HMWG. This low coupling efficieny is due to the imperfect collimation of the input beam (NA ~0.065) resulting in the excitation of lossy higher-order modes in the waveguide. A higher magnification would make the bulk optic large and the waveguide spacing unwieldy.

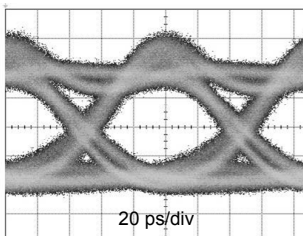

Figure 9 *10Gb/s eye thru 12.5cm HMWG with 9 slots of 0.8 mm gaps and no taps.*

A turning mirror comprising of Au coated mirror was inserted at the Tx side, while 4 pellicle beamsplitters each comprising of a single 250nm thick layer of SiN were inserted into the optical waveguide starting from the last tap. The reflectivity of the single-layer film is very polarization sensitive, with the R_s approximately 23% and R_p close to 5%. This large difference in reflectance is due to the proximity of 45° to the Brewster, or polarizing, angle of SiN. At the Brewster angle, the *p* polarized light passes without attenuation while the *s* polarized light is totally reflected. For SiN, the brewster angle is around 63.4°. This large discrepancy in reflectivity together with the poor coupling efficiency of the bulk optic resulted in a very uneven power distribution among the taps. As a result only two taps had enough power to create an open eye.

In order to increase the coupled power into the waveguide, we replaced the Tx board with a single VCSEL which was collimated by a 500µm ball lens. The ball lens produced a collimated input beam with an NA of around < 0.02. The overall coupling efficiency into the waveguide was increased to 67%. The results are shown in Figure 10. An input power of 1.3mW was coupled into the waveguide. The outcoupled power from the 4 SiN beamsplitters were 164, 110, 84 and 60µW respectively. 450 µW of optical power was left-over at the output of the waveguide. This is enough power to yield 3 to 4 more taps, assuming we have the correct beamsplitter ratios. The total loss through the system was around 33%. In order to improve the efficiency of the system, it is crucial that nonpolarizing beamsplitters be used. Translating the receiver evaluation board under each of the taps resulted in the measured eye-diagrams @ 5Gb/s as shown in Figure 10(b). The speed of the demo is limited by the speed of the evaluation boards.

These results demonstrate for the first time a vaible optical multi-drop bus. Further optimization of the beamsplitter and waveguide technology is underway. In the the next generation demo we are integrating both *fan-out* (northbound) and and *fan-in* (southbound) buses onto a single motherboard. The data rate will be increased to 10Gb/s.

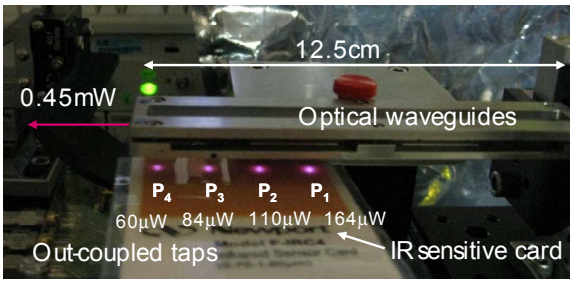

(a)

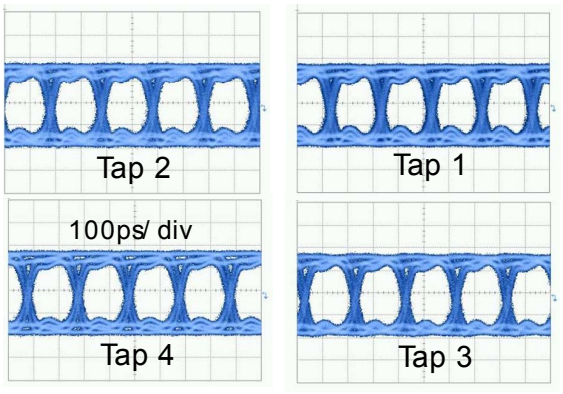

(b)

Figure 10 *(a) Photograph showing the out-coupled light from the 4 SiN pellicle beam splitters onto an IR sensitive card. The measured average power from each tap is also indicated. (b) Measured 5 Gb/s eye diagrams from each tap.*

6. Further applications

A possible application of the optical bus is to enable large disaggregated memory using a 1x8 fan-out. This is accomplished by interconnecting several memory bridge chips together via the optical bus. Command, address and write signals are sent on the northbound bus while read data is sent back on the southbound bus. The memory bridge chips serve as the interface to standard DDR3 DRAM interface (see Figure 11).

The basic optical bus structure demonstrated provides a master/slave bus where one node is responsible for scheduling all transfers. By adding appropriate bus sequencing and arbitration protocols the optical bus can be used in multi master configurations. In the current configuration only the master node can broadcast; slave nodes have only point-to-point communication back to the master.

The characteristics of the optical bus are well suited to

parallel links with source synchronous clocking. The transmission medium of HMWGs has essentially no crosstalk, with the electronic components being the only

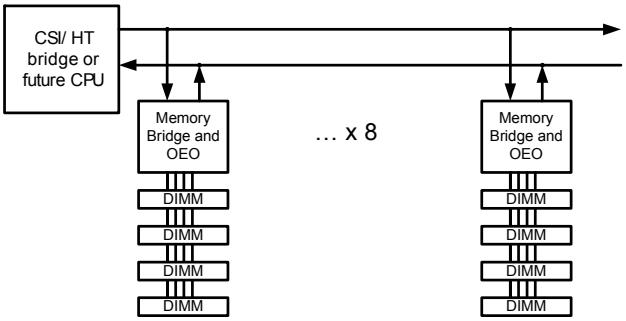

Figure 11 *Application of an optical bus to a large disaggregated memory system.*

source of crosstalk in the system. The propagation delays down the parallel paths are closely matched and stable. Source synchronous clocking is highly desirable from the perspective of the fan-in bus, as this permits rapid bus turnaround between different transmitting modules. A source synchronous optical communication link would have significant advantages in terms of both latency and power. Other possible area of application for the optical bus include expansion buses in modular switch architectures, midplanes for blade systems, and I/O systems.

The two main impediments to bringing the multi-drop bus into wide use are the cost and reliability of the optoelectronic engine and the need to support for broadcast protocols in high speed interfaces which in most case assume only point-to-point connectivity. In order to address this, we need to work closely with external optoelectronic vendors to drive down the cost and reliability of the optoelectronic engine.

7. Conclusions

The basic feasibility of an optical multi-drop bus using hollow metal waveguides and pellicle beamsplitters was demonstrated. The bus allows for power to be distributed equally among multiple nodes attached to the bus providing for an efficient utilization of the available optical power and also reduced latency. It also provides for a 2x improvement in bandwidth and power over a networked point-to-point optical link using the same number of optoelectronic transceivers. The fundamental component technologies, the hollow metal waveguide and the optical beamsplitters have been investigated. A 1x4 optical bus has been demonstrated using off-the-shelf optoelectronic components and ICs at a data rate of 5Gb/s. Future demonstrators will include both northbound and southbound buses running at 10Gb/s. With further optimization, the optical bus will provide for both improved bandwidth and connectivity with reduced power and lower latency over its electrical counterpart.

Acknowledgements

The authors would like to acknowledge the contributions from their colleagues Sriram Ramamoorthi and Len Seals at AMS/TDO for their valuable technical contributions. We would also like to acknowledge the contributions from our colleagues Lennie Kiyama, Eric Montgomery, Plary Mendoza and Gil Perusa in HPLabs for their technical support. We would also like to thank Laura King and Stan Williams for their continued support in this project.

References

[1] Dally, William J.; Poulton, John W; "Digital Systems Engineering", Cambridge University Press, 1998

[2] Graham, Martin; Johnson, Howard; "High-Speed Signal Propagation – Advanced Black Magic" Prentice Hall, 2003

[3] Jacob, Bruce; Wang, David; "Memory Systems Architecture and Performance Analysis" University of Maryland ECE Department Lecture http://www.ece.umd.edu/class/enee759h.S2005/lectures/Lecture15.pdf

[4] Goodin, Dan; "IT Confronts the Datacenter Power Crisis" Infoworld, October 2006

[5] Schares *et al*, "Terabus: Terabit/Second-Class Card-Level Optical Interconnect Technologies," *IEEE JSTQE,* Vol. **12**, No. 5, pp 1032-1044, September/October 2006.

[6] Alexei Glebov, Michael G. Lee, Kishio Yokouchi, "Optical Interconnect Modules With Fully Integrated Reflector Mirrors," *IEEE Photonics Technology Letters*, Vol. **17,** No. 7, pp. 1540-1542, July 2005

[7] D. M. Chiarulli, S. P. Levitan, R. G. Melhem, M. Bidnurkar, R. Ditmore, G. Gravenstreter, Z. Guo, C. Qao, J. Teza,"Optoelectronic Buses for High Performance Computing," *Proceedings of the IEEE*, Vol. **92**, No. 11, pp.1701-1709, November 1994

[8] Keiichiro Itoh, Riyo Konno, Yoshitada Katagiri, and Testuo Mikazuki, "Data Transmission Performance of an Optical Backboard Bus," Proc. Electronic Manufacturing Symposium, 1995

[9] Seiki Hiramatsu, Kiyotaka Miura, and Kazuyuki Hirao, "Optical Backplane Connectors Using Three-Dimensional Waveguide Arrays," Journal of Lightwave Technology, Vol. 25, Issue 9, pp. 2776-2782

[10] R.T.Chen, Lei Lin, Chulchae Choi,Y.J. Liu, B. Bihari, L. Wu, S. Tang, R. Wickman, B. Picor, M.K. Hibb-Brenner, J. Bristow, Y.S. Liu, **"**Fully embedded board-level guided-wave optoelectronic interconnects," Proceedings of the IEEE Volume 88,

Issue 6, June 2000 pp.780 - 793
[11] Xuliang Han, G. Kim, G.J. Lipovski, R.T. Chen, "An optical centralized shared-bus architecture demonstrator for microprocessor-to-memory interconnects," *IEEE JSTQE,* Vol. 9 (2), pp 512-517, March/April 2003
[12] Alexei Glebov, Michael Lee and Kishio Yokouchi, "Integration Technologies for optical backplane interconnects," *Optical Eng* 46(1), 015403 January 2007
[13] Marcatili *et al*, " Hollow Metallic and Dielectric Waveguides for Long Distance Optical Transmission and Lasers," *Bell Syst. Tech. J.* **43**, 1783 (1964).
[14] Pavel Kornilovich, "Optical Modes of Rectangular Hollow Metal Waveguides," Hewlett Packard/AMS Internal Memo, August 2007
[15] TFCalc, *Software Spectra Inc.*