

Optical hybrid package with an 8-channel 18GT/s CMOS transceiver for chip-to-chip optical interconnect

E. Mohammed^{*a}, J. Liao^a, A. Kern^a, D. Lu^c, H. Braunisch^c, T. Thomas^b, S. Hyvonen^a, S. Palermo^a
and I. A. Young^a

^aIntel Corporation 2513 NW 229th St., Portland, OR 97124

^bIntel Corporation 2111 NE 25th Av., Portland, OR 97124

^cIntel Corporation 5000 W Chandler Blvd., Chandler, AZ 85226

ABSTRACT

We describe the design and development of a high-speed 8-channel hybrid integrated optical transceiver package with Clock and Data Recovery (CDR) circuits. The package concept has been developed to be compatible with microprocessor package technology and at the same time allow the integration of low cost, high-performance optical components. A 90nm CMOS optical transceiver chip, 850nm 10Gb/s GaAs based vertical cavity surface emitting laser (VCSEL) array and PIN photodiode array are flip-chip mounted on a standard microprocessor Land Grid Array (LGA) package substrate. The CMOS drivers and receivers on the transceiver chip and the optical components (VCSEL and Photodiode arrays) are electrically coupled using a short transmission line routed on the top surface of the package. VCSEL and photodiode arrays are optically coupled to on-package integrated polymer waveguide arrays with metallized 45° mirrors. The waveguides, which are terminated with multi-terminal (MT) fiber optic connectors, couple out/in high-speed optical signals to/from the chip. The CMOS transceiver chip fully integrates all analog optical circuits such as VCSEL drivers, transimpedance amplifiers and clock and data recovery (CDR) retiming circuit with a low jitter LC-PLL. Digital circuits for pseudorandom bit-pattern sequence generators (PRBS) and bit-error rate test (BERT) are fully integrated. 20Gb/s electrical and 18Gb/s optical eye diagrams for the transmitter were measured out of the package. A fully packaged transmitter and receiver including clock data recovery at 10Gb/s have also been measured.

Keywords: Optical I/O, CMOS, CDR, VCSEL arrays, photodetector arrays, polymer waveguides, hybrid packaging, flip-chip

1. INTRODUCTION

The recent transition to multi-core or many core architectures in the computer industry combined with the steady increase of central processing unit (CPU) clock frequency brings new bandwidth demands to package-to-package (chip-to-chip) electrical I/O between CPU and chipset on computer motherboards and among multiple CPUs in blade servers¹. The era of terascale computing is quickly approaching¹. Traditional practices to meeting these demands require increasing the I/O bus width and individual I/O data transmission rate. However, unless alternative ways are found, this method of increasing I/O bandwidth eventually comes at the cost of compromising the integrity of the signal and increasing power consumption. This section focuses on defining the underlying problems both in the personal computers (PC) and server space using existing I/O architectures but being mindful of the potential future architectural changes in the CPU, I/O or package. These architectural changes may shift or delay such problems in electrical I/O but optics has the potential to provide the ultimate solution because of its good attributes of low signal propagation loss in fibers/waveguides. Figure 1. shows the system I/O architecture in a typical modern personal computer (PC). It shows the CPU and chipset, consisting of memory control hub (MCH) and I/O control hub (ICH) all connected to each other by interconnect buses. The key potential future bandwidth bottleneck points, and therefore candidates for optics, are the front side bus (FSB) that connects the CPU and MCH and the PCI Express bus that connects the graphic accelerator and the MCH. Lower bandwidth PCI express buses are also used in the system architecture below for external I/O communication such as Gigabit Ethernet.

*edris.m.mohammed@intel.com; phone (503) 613-4298; fax (971) 214-7807; Intel Corporation




Figure 1. The current Intel system I/O architecture for x86 PC.

Figure 2. is a schematic showing the corresponding physical interconnection between the CPU and the MCH through an individual FSB electrical I/O. As illustrated in the figure the transmitted signal from the CPU reaches the MCH after traversing approximately 18 centimeters on the printed circuit board (PCB) and passing through several discontinuities on both sides of the package. In this process the received signal at the MCH will suffer from frequency dependent losses, reflections and cross-talk that will close the data eye and increases its bit-error-rate unless the interconnect is compensated.




Figure 2. A schematic of interconnection between CPU and MCH through the front side bus (System Bus).

In the server space similar issues are encountered for I/O buses between CPU chips (packages) located on different server blades. However, due to the longer length the signal travels and the additional electrical discontinuities in its path, much more severe signal integrity issues are expected as the data transfer rate is increased. Also, optical transmission becomes an increasingly viable solution for longer interconnect lengths. Therefore, we will examine the blade server physical layer to understand the source of bandwidth limitation and the potential for optical interconnects to be deployed. Figure 3. illustrates an arrangement of two server blades on a backplane connected electrically. For chip-to-chip I/O data transmission on the backplanes, the high-speed signal driven out off-chip from the I/O bumps of CPU1 has to travel through two multi-layer package substrates, two CPU sockets, three multilayer PCBs and two backplane connectors before being received at CPU2 on a different blade 54 centimeters away. During this travel the signal suffers from frequency dependent loss and encounters many discontinuities such as various pad/bump capacitances and package via inductances as shown in Figure 3. The frequency dependant transmission losses are caused primarily by skin effect and dielectric loss. High frequency components of the signal are greatly attenuated in comparison to lower frequencies. The attenuation of signal on a transmission line causes intersymbol interference (ISI) because this attenuation prevents the signal from reaching full strength within its symbol time, causing it to spread into the next bit time slot. This signal deterioration effect will be reflected in the smearing of the eye diagram and increase in the bit-error rate at the receiver.




Figure 3. An example of CPU Package-to-package electrical interconnection between server blades. Colored traces on the board show I/O signal path from CPU1 to CPU2.

Figure 4 shows simulated insertion loss data for a situation found in the above mentioned server blade for 54 cm long electrical interconnect connecting the two CPUs. The result includes the effects of packages, pad capacitance, via inductance and two backplane connectors in addition to the linear frequency dependent losses associated with the transmission line traces on PCB made out of FR4 material. The insertion loss data shows a largely linear dependence on frequency indicating the dominant factor to be dielectric attenuation. At 10GHz the simulation for this link predicts an insertion loss of 50dB which is near the limit for electrical transmission. A dip around 10GHz is likely to be from discontinuity due to via inductance or impedance mismatch. Generally driver pre-emphasis or various receiver equalization techniques are used to compensate for the signal loss. Pre-emphasis operates by boosting the high frequency energy every time there is a transition in the data, since this is where the most signal loss occurs. Receiver equalization provides functionality in the receiver to help overcome high frequency signal losses of the transmission medium. It acts as a high pass filter and amplifier to the data as it enters into the receiver.




Figure 4. Projected channel loss for electrical in electrical link in servers.

However, both methods have some associated overhead since they use additional circuitry. The success of long, medium and short distance optical communication has proven that at high frequencies optical fibers/waveguides have very high bandwidth compared to electrical interconnects, minimal propagation loss and negligible cross-talk. However, viable implementation of optical interconnects in ultra-short distances in computing systems requires looking at the problem from a new perspective. A mere replacement of physical layers may not be sufficient unless combined with a change in some aspect of the system architecture. As most of the problems are parasitic related, the architectural change should include very close integration of electronic and optical chips together with revolutionary packaging approaches that

emphasize miniaturization. Figure 5 is an illustrative hypothetical example for a possible approach to use optics in the server architecture that emphasizes close integration of electronic and optoelectronic dies. In this arrangement the laser driver and the receiver are assumed to be integrated with the CPU. They are then interfaced with the respective optoelectronic dies (lasers or photodetectors) through short microstrip lines (3-5 mm in length) on the package as shown below. Typical insertion losses for these microstrip lines would be ~1dB at 20GHz. Multimode fiber arrays that are terminated with multi-terminal (MT) optical connectors physically connect the two chips to enable signal transmission between CPU1 and CPU2. If an efficient optical coupling is achieved at every interface, there will be a potential for the total power consumption to decrease appreciably.

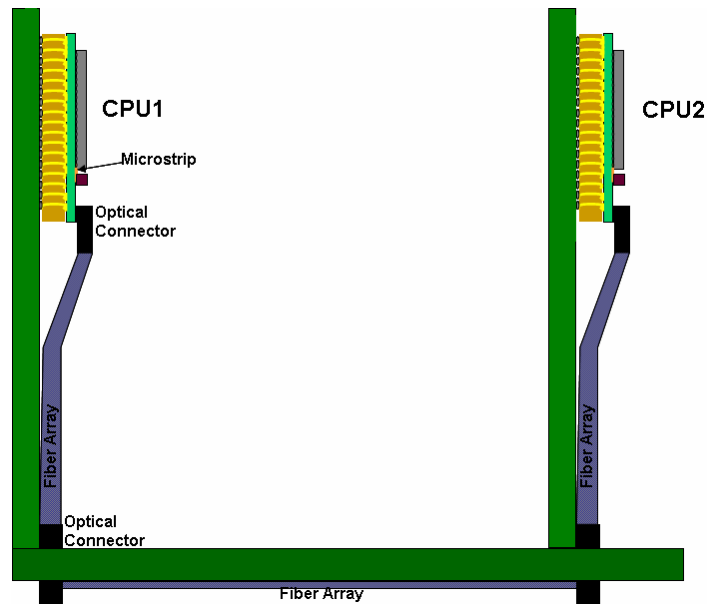


Figure 5. Optical interconnection between two microprocessor units on blade servers.

The following sections describe the development of an 8-channel optical transceiver hybrid package which attempts to demonstrate integration of CMOS and optics using standard microprocessor packaging. Section 2 discusses the architecture and design of the chip including the key circuit blocks for optical. Section 3 will describe the package architecture and design. Experimental results will be discussed in section 4 and a summary will be given in section 5.

2. ARCHITECTURE and DESIGN

2.1 Transceiver chip architecture

The dual transceiver chip shown in Figure 6 was designed and fabricated in Intel’s 90nm digital CMOS process technology with seven metal layers on high-resistivity substrate². The high resistivity reduces noise coupling through the substrate. The total dimension of the dual chip is 5 x 10 mm² but 1/3 of the chip is unoccupied. The chip has eight transmitter (Tx) and receiver (Rx) channels, six of which contain key circuit blocks such as transimpedance and limiting amplifiers (TIA, LIA), clock and data recovery (CDR), pseudorandom bit-pattern sequence (PRBS) generator and VCSEL drivers. The two remaining Tx and Rx channels only contain TIA, LIA and VCSEL drivers and they are intended to test these individual blocks separately using external inputs. All logic functions on the chip are implemented using high speed differential current mode logic (CML). The differential nature of the circuits makes them immune to power supply variations and substrate noise. The high-speed gates of CML topology have been optimized to provide bandwidth in excess of 10GHz and they do not require any inductive peaking as used in many similar recent optical interconnect reports^{3,4}. Inductors are not desired because they complicate chip fabrication and require more real estate increasing the total chip area. The center block of the chip is used for global distribution of power and bias as well as testing scan bits and reference clock signals to all channels. An on-chip bandgap reference circuit provides bias to all the CML circuits. Three separate power supplies were used for the core logic, the voltage control oscillator (VCO) and for the cascode TIA. These are 1.2V, 1.4V and 1.8V, respectively. In order to minimize switching noise decoupling and

bypass capacitors were used at the chip, package and board level.

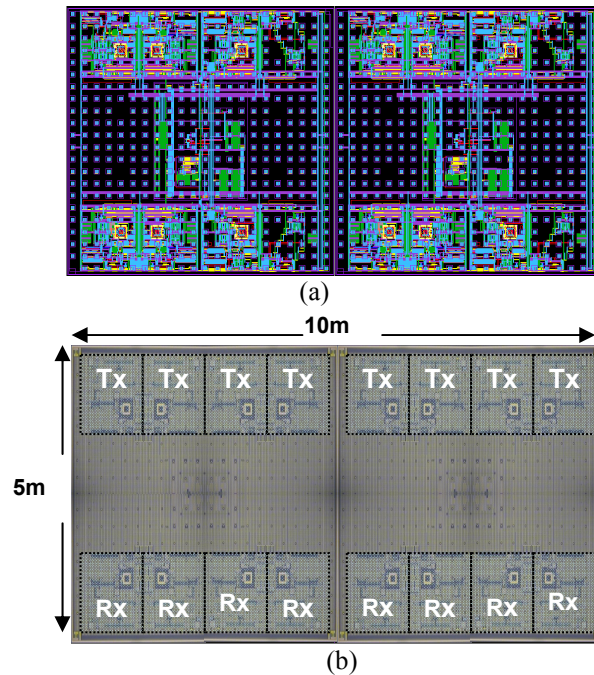




Figure 6. (a) transceiver chip layout (b) picture of actual transceiver.

Single channel architecture

Figure 7 shows the block diagram of a fully integrated single transceiver channel with all the necessary circuit blocks for standard CDR based optical I/O transmission. This is a universal channel that is used as a transmitter (Tx) or a receiver (Rx) depending on the configuration set by the scan chain through a computer interface. When the unit is in a transmission mode, the key functional blocks used are the phase locked loop (PLL) with input reference clock, the pseudorandom bit sequence (PRBS) generator and the VCSEL driver. The PLL uses an LC-VCO⁵ topology for its superior low jitter performance. The VCO has four digitally controlled switch capacitors which extend its tuning range from 7.5-11.2GHz. In a transmit mode the PLL is locked using a standard PFD loop. The clock from the PLL is then fed to a pattern generator circuit block which produces $2^{15}-1$ PRBS non-return to zero (NRZ) data signals to drive the VCSEL driver as shown in Figure 7. The transmit block also has an option where it receives a high-speed external clock using the input LIA/MUX circuits to drive the PRBS directly. The PRBS has the capability for half-rate sampling in order to double the output frequency of the PRBS data stream. The receive mode allows both electrical and optical inputs to be selected using the input MUX circuit. In optical receive mode a photocurrent from an on-package photodetector is fed to a transimpedance amplifier (TIA) for amplification and then through a limiting amplifier (LIA) to CML level. It then passes to a PLL based CDR unit as shown in Figure 7. The CDR has a dual loop PLL based on phase frequency detector (PFD) and a bang-bang phase detector (BBPD). During startup it first uses the external reference clock and PFD loop for frequency locking, as in the transmit mode described above. When the frequency difference between VCO and reference clock is less than 500ppm, the lock signal goes high, and the PLL starts to lock to the received serial data using a bang-bang phase detector (PD). The CDR retimes the incoming high-speed data to remove jitter. The output MUX selects the received and retimed signal from the CDR and passes it to a 50Ω output driver for off-chip detection.



(a)



(b)

Figure 7. Single universal channel architecture (a) block diagram (b) channel layout design.

VCSEL drivers

The VCSEL devices used in this development are rated for 10Gb/s. Beyond 10Gb/s they are bandwidth limited with a slow tail due to intrinsic and extrinsic parasitic effects such as carrier diffusion⁶ and device capacitance of 700fF. In order to compensate for these effects and increase the frequency, drivers with pre-emphasis were used. The pre-emphasis VCSEL drivers⁷ integrated on the transceiver chip have a dual edge pre-emphasis and generate a sub-bit-period timing directly from the input data. A pre-emphasized current waveform is generated by summing the main modulation current with a delayed and weighted peaking current in order to produce pre-emphasis pulses at each data transition as shown in Figure 8. The VCSEL supply voltage is independent from the core supply to allow test flexibility. A 5-bit digital-to-analog converter (DAC) provides digital control of the output currents from 0–46mA and 0–15mA for the main and pre-emphasis drivers, respectively. Typical average currents provided to VCSELs range from 6–10mA which corresponds to an average optical power of 1.5–2mW. The VCSEL driver is output terminated and connected to the VCSEL through a 50Ω microstrip transmission line that was routed on the top surface of the package. The design of the VCSEL drivers comprehended this microstrip transmission line and the VCSEL optical and electrical models. Extracted optical and electrical models of the VCSEL have been included to accurately simulate the effects of carrier dynamics. The detail description of the design of this pre-emphasis VCSEL driver is available⁷.

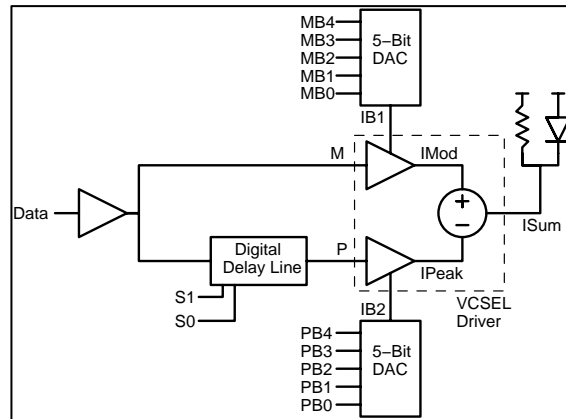


Figure 8. Block diagram of the pre-emphasis VCSEL driver.

The transimpedance amplifier

The transimpedance amplifier integrated on each transceiver channel has symmetric feedback and has a differential topology as shown in Figure 9. This differential topology converts the single-ended current input to a differential output voltage to help mitigate supply noise at subsequent gain states and improve stability by allowing negative feedback from a two-stage amplifier. The TIA was designed to have a dominant pole of 10GHz which provides a data rate in excess of 12.5Gb/s for an input parasitic capacitance of $C_p=250\text{fF}$. It has a feedback resistance of 314Ω and open loop gain of 3.9. It receives a single-ended photocurrent of $200\mu\text{A}$ from the photodiode and generates a differential $2 \times 50\text{mVp-p}$ output that is fed to the limiting amplifier (LIA) which converts it to a CML level output. The LIA consists of a cascade of buffers based on current mode logic (CML). In the first stage of the LIA, sizing the differential pair devices for optimum gain bandwidth product minimizes input referred noise. Typically this LIA achieves an input sensitivity of $2 \times 22.5\text{mVp-p}$. The LIA design does not require any inductive peaking since the 90nm CMOS technology² has NMOS transistors with an f_t of at least 130GHz. A digital-to-analog (DAC) DC offset cancellation circuit is implemented at the TIA input to remove the DC current component caused by the off-state DC bias of the VCSEL transmitter. In the packaged transceiver the combined photodiode and pad/bump/ESD capacitances could go as high as 500fF. This increase in the parasitic capacitance decreases the targeted data rate that could be measured in the receiver channel. At the chip level for the same TIA an electrical eye diagram of 18Gb/s was measured previously for an input capacitance of 88fF indicating a strong dependence on the parasitic capacitance⁷.




Figure 9. Block diagram of a symmetric transimpedance amplifier connected to a photodiode

3. TRANSCEIVER PACKAGE

3.1 Package architecture and design

The package architecture has been developed to be compatible with the current microprocessor C4 organic package technology and at the same time allow the integration of low cost, high-performance optical components. Since this packaging approach has been discussed in detail previously, only a brief relevant description is given here^{8,9}. Figure 10 shows the package layout design. The package substrate has a dimension of 31x31mm² with a standard build up consisting of laminated copper layers separated by dielectric. The electrical design of the package involves the routing of signal and power lines as well the use of decoupling capacitors. All the high-speed electrical lines on the substrate are routed as controlled impedance (50Ω, single-ended or 90Ω, differential) 5GHz microstrip or stripline traces. The single-ended microstrip lines are routed on the top surface of the substrate and connect the VCSEL driver (TIA) output bumps to VCSEL (photodiode) bumps on the package. The close proximity between them minimizes frequency dependent loss. The length of these microstrip lines varies because of the relative size difference between the Si die and optoelectronic chips as illustrated in Figure 10. The different layers of the substrate provide signal, power and biasing. Bias planes were used to externally bias VCSELs and photodiodes. The decoupling capacitors on the top surface are used to minimize switching noise on power and bias planes/lines.

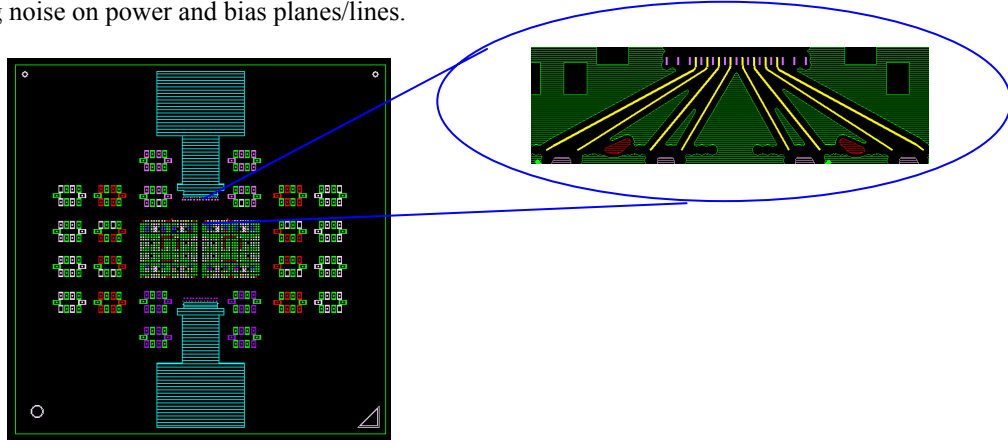


Figure 10. Package substrate layout design. Microstrip traces are magnified to show their relative size differences.

Figure 11a illustrates the fully assembled optical transceiver package. This includes the CMOS transceiver chip, VCSEL and photodiode array chips, MT connectorized polymer waveguide arrays and decoupling capacitors, all attached to a common microprocessor land grid array (LGA) substrate. The polymer waveguide device consists of a 45° mirror on one end and a modified industry standard multi-terminal (MT) connector on the other. The waveguide array has a core dimension of 35μm x 35μm.




Figure 11. (a) A fully assembled optical transceiver unit (b) a side view of optical coupling of lasers/photodetectors to waveguides through a 45 degree mirror.

The transmitter includes a VCSEL optical source array and a polymer waveguide array. VCSEL arrays are flip-chip bonded on the substrate and coupled to polymer waveguides with 45° metal mirrors to direct light at right angle for transmission through the waveguide. Similarly, the receiver consists of identical optical waveguide assembly to the transmitter section with a high-speed GaAs photodiode array replacing the VCSEL. Figure 11b shows a side view of the interfacing between optoelectronics chips and waveguides. The optoelectronic chips are flip-chip bonded with their apertures face down to allow polymer waveguides to slide under and couple. The total optical loss budget for our system includes VCSEL/photodiode coupling loss to the 45° mirror on either side of the optical link, propagation loss through the waveguide, MT connectorization loss and Fresnel losses at the interfaces of connectors. Based on various parameters of individual optical components, the total optical loss budget for the complete link was calculated to be between 7dB (best case) and 12dB (worst case)⁹.

4. MEASUREMENT RESULTS

4.1 Optical and electrical characterization setup

Figure 12 shows a photograph of the electrical/optical measurement setup. Optical transceiver packages are mounted on high-speed testing sockets on a printed circuit board (PCB). Bias and power were supplied through connectors on the top surface of the PCB and high speed differential output signals were collected from edge mount connectors. Using a computer with Labview interface software, a scan chain on chip sent signals to configure parameters in the transmitter and the receiver. The VCSEL bumps on the LGA substrate were directly probed to acquire electrical signals from the chip.




Figure 12. A photograph of the experimental set-up for optical transmission tests.

4.2 High-speed electrical and optical characterization 10Gb/s transmitter and receiver

Figure 13 below shows the 10Gb/s optical measurement results of the transmitter and receiver for the two channels mentioned in section 2.1. These channels are integrated on the main transceiver chip but do not contain circuit blocks other than the pre-emphasis VCSEL driver and the TIA/LIA pair. For the transmitter measurement in Figure 13a high-speed external electrical PRBS data is received by the on-chip integrated limiting amplifier to drive the pre-emphasis VCSEL driver. VCSELs were biased with an average current of 7mA.

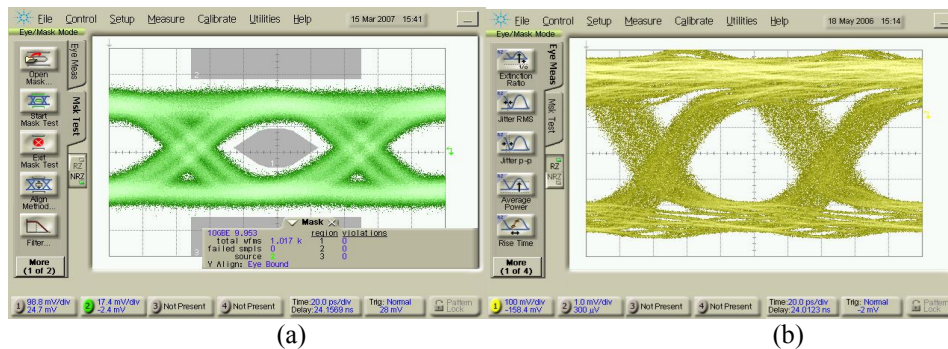


Figure 13. Optical eye diagrams from fully packaged transmitter (a) and receiver (b).


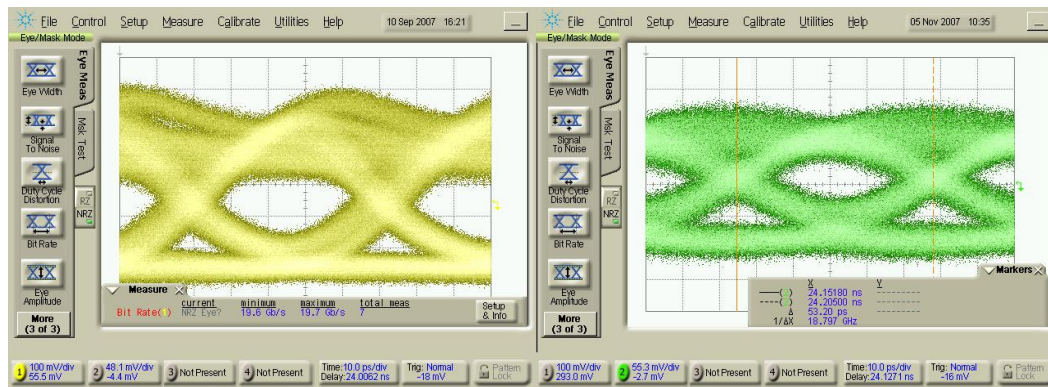


Figure 14. Received and retimed data from CDR at 10Gb/s driven electrically off-chip.

The resulting optical data was measured by coupling a 1x12 multimode fiber with MT connectors and feeding individual optical channel outputs to a 12GHz Newfocus optical photoreceiver and 30GHz Agilent DCA oscilloscope. The measurement has an eye opening of 70ps. The rise and fall times were 30ps and 35ps, respectively with a peak-to-peak jitter of 30ps. Figure 13b shows the receiver eye opening at 10Gb/s. An eye opening of 60ps and rise and fall times of 25ps and 30ps with a peak-to-peak jitter of 30ps were measured. Figure 14 shows a received and retimed optical data at 10Gb/s from the universal channel containing the CDR in the receive signal path. The eye diagram has a horizontal eye opening of 60ps and rise and fall times of 35ps and 40ps, respectively. The received eye patterns have some ISI that may have been caused by losses from the long transmission line for the electrical output signal on the PCB.

20Gb/s transmitter

Figure 15 shows the electrical and optical eye diagrams of the transmitter on the universal channel at 20Gb/s and 18Gb/s data rate using a $2^{15}-1$ PRBS data pattern which is generated by the on-chip PRBS generator after the transceiver chip was packaged. The electrical data was measured by directly probing one of the on-package VCSEL array contact pads of the pre-emphasis driver with a high-speed GSG coplanar probe and feeding the output to a 30GHz Agilent 86100 series digital component analyzer (DCA). Before its detection the high-speed signal passes through two package solder bumps for the CMOS and the VCSEL chips. A microstrip transmission line length of ~4mm connects the two bumps. Impedance discontinuities were expected at all these points that compromise the signal integrity. Some intersymbol interference (ISI) and reflection resulting from the signal path on chip and off-chip is apparent in the electrical eye diagram. However, a good 20Gb/s open electrical eye is observed with a vertical and horizontal eye opening of 175mV and 35ps, respectively. Also, the measurement showed a rise and fall time of 25ps and 30ps with a peak to peak jitter of 16ps. The optical data was measured by feeding the 20Gb/s electrical signal in Figure 15a directly to the VCSEL using high-speed coplanar probes and feeding the optical output to a 12GHz Newfocus photoreceiver using a multimode fiber. Since the signal path now encounters combined parasitics from the package and from the VCSEL, the driver currents with the pre-emphasis and the associated DC biases were re-optimized to provide the best possible eye at 18Gb/s.



(a) (b)
Figure 15. Transmitter eye diagrams (a) 20Gb/s Electrical (b) 18Gb/s Optical.

An average drive current of 9mA with 2.8V VCSEL external bias was used to provide an average optical power of 2mW to the photoreceiver. An 18Gb/s open optical eye was observed with a vertical and horizontal eye opening of 70mV and 30ps (60% of total eye), respectively. The rise and fall time were 25ps and 30ps with an increased peak-to-peak jitter of 23ps. The increase in the jitter and noise around the eye might be associated with the optical device relative intensity noise. Reflections and bandwidth limitations of the 12GHz photoreceiver also contributed in reducing the eye opening.

SUMMARY

This paper introduced the interconnect bottleneck problems in the PC and server space using existing I/O architectures and described the role optics could play in overcoming these because of its good attributes for signal propagation. As a demonstration of the viability of optics for these computing applications a high-speed optical CMOS based transceiver package has been developed. With this hybrid package 20Gb/s electrical and 18Gb/s optical eye diagrams for LC-PLL based VCSEL transmitters has been measured. 10Gb/s Fully assembled transmitters and receivers with clock data recovery were also demonstrated.

ACKNOWLEDGMENTS

The authors would like to thank Douglas Stunkard and Julie Mckenney for chip layout design and Bent Danielson for test board layout design.

REFERENCES

1. J. Held, J. Bautista, and S. Koehl, "From a few core to Many: A Tera-scale Computing Research Overview," *Research at Intel white paper*, download.intel.com/research/platform/terascale/terascale_overview_paper.pdf, 2006.
2. K. Kuhn, et al., "A 90nm Communication technology featuring SiGe HBT transistors, RF CMOS, Precision R-L-C RF elements and 1 μ m² 6-T SRAM cell", *IEDM 2002, Dig. Tech. Papers*, Dec 2002, pp. 73-76.
3. D. Kucharski, Y. Kwark, D. Kuchta, D. Guckenberger, K. Kornegay, M. Tan, C-K. Lin and A. Tandon, "A 20Gb/s VCSEL driver with pre-emphasis and regulated output impedance in 0.13 μ m CMOS" in *Proc. ISSCC*, Feb. 2005, pp. 222-223.
4. C. Kromer, C. Sialm, C. Berger, T. Morf, M. L. Schmatz, F. Ellinger, D. Erni, G. L. Bona and H. Jackel, "A 100mW 4x10Gb/s Transceiver in 80nm CMOS for High-Density Optical Interconnects" *IEEE JSSC*, vol. 40, no 12, pp. 2667-2679.
5. H. Muthali, T. Thomas and I. Young, "A CMOS 10Gb/s SONET Transceiver," *IEEE JSSC*, vol. 39, No. 7, pp. 1026-1033, July 2004.
6. J. J Morkuni, P. V. Vena, A. V. Harton, K. W. Wyatt and S. M. Kang, "Spatially Independent VCSEL Models for the Simulation of Diffusive Turn-Off Transients," *J. Lightwave. Technol.*, vol. 17, no. 1, pp. 95-102, Jan. 1999.
7. A. Kern, A. Chandrakasan and I. Young, "18Gb/s Optical IO: VCSEL Driver and TIA in 90nm CMOS," *VLSI Circuits, 2007 IEEE Symposium*, pp. 276-277, 14-16 June 2007.
8. E. Mohammed, T. P. Thomas, D. Lu, H. Braunisch, S. Towle, B. C. Barnett, I. A. Young and G. Vandentop "Optical I/O Technology for Digital VLSI" in *Proc. SPIE Photonic West Conference*, vol. 5358, June 10, 2004.
9. E. Mohammed, A. Alduino, T. Thomas, H. Braunisch, D. Lu, J. Heck, A. Liu, I. Young, B. Barnett, G. Vandentop, R. Mooney, "Optical Interconnect System Integration for Ultra-Short-Reach Applications," *Intel Technology Journal*, Vol. 8, no. 2, pp. 115-127, May 2004.