

Optical Technology for Energy Efficient I/O in High Performance Computing

Ian A. Young, Edris M. Mohammed, Jason T. S. Liao, and Alexandra M. Kern, Intel Corporation
Samuel Palermo, Texas A&M University

Bruce A. Block, Miriam R. Reshotko, and Peter L. D. Chang, Intel Corporation

ABSTRACT

Future high-performance computing systems will require optical I/O to achieve their aggressive bandwidth requirements of multiple terabytes per second with energy efficiency better than 1 pJ/b. Near-term optical I/O solutions will integrate optical and electrical components in the package, but longer-term solutions will integrate photonic elements directly into the CMOS chip to further improve bandwidth and energy efficiency. The presented near-term optical I/O uses a customized package to assemble CMOS integrated transceiver circuits, discrete VCSEL/detector arrays, and polymer waveguides. Circuit simulations predict this architecture will achieve energy efficiency better than 1 pJ/b at the 16 nm CMOS technology node. Monolithic photonic CMOS process technology enables higher bandwidth and improved energy efficiency for chip-to-chip optical I/O through integration of electro-optical polymer based modulators, silicon nitride waveguides, and polycrystalline germanium (Ge) detectors into a CMOS logic process. Experimental results for the photonic CMOS ring resonator (RR) modulators and Ge detectors demonstrate performance at up to 40 Gb/s and analysis predicts that photonic CMOS will eventually enable energy efficiency of 0.3 pJ/b with 16 nm CMOS. Optical interconnect technologies with multilane communication or wavelength-division multiplexing will further increase bandwidth to provide the multiple-terabyte-per-second optical interconnect solution that enables scaling of high-performance computing into and beyond the tera-scale era.

INTRODUCTION

The microprocessor architecture transition from multicore to many-core will increase chip-to-chip input/output (I/O) bandwidth demands at processor/memory interfaces and in multiprocessor systems. Near-term projections, shown in Fig. 1, estimate that CPU-to-memory interconnects will require 100 Gbytes/s bandwidth in 2012–2013. Future many-core architectures will require bandwidths from 200 Gbytes/s to 1.0 Tbyte/s and begin the era of tera-scale computing.

To meet these bandwidth demands, tradition-

al chip-to-chip electrical interconnect techniques will require increased transceiver circuit complexity and costlier materials. However, due to electrical channel loss, increasing I/O bandwidth in electrical links eventually comes at the cost of reducing interconnect link length, reducing signal integrity, or increasing power consumption.

In contrast, optical interconnects have negligible frequency-dependent loss and low crosstalk. Performance is independent of link length (for lengths of interest in chip-to-chip I/O), and little or no equalization is required. This motivates chip-to-chip I/O architects to consider optical I/O as a means of scaling data rates in a power-efficient manner.

ELECTRICAL LINK ISSUES

Figure 2 shows the components of a typical high-speed electrical link, including the transmitter, receiver, timing system, and channel. A phase-locked loop (PLL) frequency synthesizer generates the transmit serialization clocks, and the receiver timing system provides the serial data sampling clocks. The design complexity of the transmitter and receiver increases to include additional equalization circuitry as data rates scale above electrical channel bandwidths.

Electrical channel frequency characteristics are dependent on channel length. An inset in Fig. 2 shows the channel response for three typical electrical channels, a 17-in server backplane channel with two connectors, a 7-in desktop channel with no connectors, and an 8-in high-performance cable channel. The frequency-dependent loss exponentially increases with channel length, as illustrated by the loss difference between a 17-in backplane channel and a 7-in desktop channel. Attenuation and dispersion in these low-pass channels introduces intersymbol interference (ISI) in high-speed data patterns. Equalization can cancel ISI and open the received data eye, but requires additional circuit complexity, which increases I/O power and area. Equalization is typically implemented with a progressive combination of transmitter (Tx) feed-forward equalization (FFE), receiver (Rx) continuous-time linear equalizers (CTLEs), and decision feedback equalization (DFE).

S. Palermo was with Intel Corp. and is now with Texas A&M University.

A detailed circuit simulation study shows that electrical link bandwidth is limited by either the channel, at the frequency beyond which loss cannot be overcome with equalization, or the complementary metal oxide semiconductor (CMOS) technology, at the frequency beyond which the required equalization techniques cannot be implemented in an energy-efficient manner [1]. In 45 nm CMOS with constant 1 V_{pp} transmit signaling, power efficiency of transmitter and receiver front-end circuits initially improves as the data rate increases. This trend reverses, and power efficiency begins to decline with data rate as more complex equalization becomes necessary. While transmit equalization can be implemented with little additional energy, a CTLE with sufficient gain-bandwidth product requires significant power, so the energy efficiency degrades rapidly once a CTLE becomes necessary. Ultimately, the maximum data rates in all three channels shown in Fig. 2 are limited by the equalization circuit speed, as the 45 nm technology cannot support efficient DFE in the 20 Gb/s range.

Circuit simulation estimates based on a predictive 16 nm CMOS technology node reveal that the faster transistors remove the CMOS technology limitations and allow efficient implementation of all equalization circuitry necessary to operate the two conventional electrical channels at their fundamental limits [1]. Channel loss, transmit peak power constraints, receiver sensitivity, and jitter eventually limit the maximum data rate at which the desired 10⁻¹² bit error rate (BER) can be achieved in backplane and desktop channels, even though significant equalization is used. For the shorter lengths, a high-performance low-loss flex cable channel is still technology-limited because it does not require DFE until the data rate exceeds 40 Gb/s, at which point a DFE cannot be implemented efficiently in the projected 16 nm node.

OPTICAL I/O IMPLEMENTATION USING A HYBRID MCM PACKAGE

For the near term, the proposed 12-channel optical transceiver architecture allows package integration of low-cost high-performance optical components in existing microprocessor package technology. This hybrid architecture integrates CMOS and discrete optical components in a multichip module (MCM) package. In this architecture a multichannel optical transceiver chip, an 850 nm 10 Gb/s GaAs vertical-cavity surface-emitting laser (VCSEL) 1 × 12 linear array (or *n* × 12 array), and a PIN photodiode 1 × 12 linear array (or *n* × 12 array) are flip-chip mounted on a standard microprocessor organic land grid array (OLGA) package substrate. The CMOS drivers and receivers on the transceiver chip are electrically coupled to the VCSELs and photodiodes with very short transmission lines routed on the top surface of the package. The VCSEL and photodiode arrays are optically coupled to on-package integrated polymer waveguide arrays with metalized 45° mirrors. The waveguides couple the optical signals from the VCSELs and photodetectors to standard multiterminal (MT) fiber optic connectors, which connect to 1 × 12 (or *n* × 12 array) waveguide or fiber arrays to couple the light off-chip.




Figure 1. Historical CPU trend — I/O bandwidth vs. year.

THE TRANSEIVER CHIP ARCHITECTURE

The transceiver architecture for optical I/O shares many common features with the typical electrical transceiver shown in Fig. 2, including the serializer, deserializer, clock generation, and clock recovery. Transmitter clocks are generated with a PLL, and receiver clocks are either recovered from the data with a CDR (as in the presented link) or forwarded from the transmitter. Equalization complexity is significantly reduced compared to electrical links, but new circuits are required to perform the electrical-optical-electrical conversion. New package technologies are required to integrate the optical and electrical components of the link. The following sections describe the package and circuit innovations that enable a 10 Gb/s hybrid optical link in 90 nm CMOS.

PACKAGE ARCHITECTURE

The package architecture allows the integration of low-cost high-performance optical components with standard microprocessor flip-chip OLGA package technology [2]. Figure 3 shows a photograph of the fully assembled optical transceiver package and a drawing illustrating the subcomponents. The package substrate is a stack of laminated copper layers separated by a dielectric. A trench is fabricated in the substrate to accommodate the multimode polymer waveguides, which have square apertures with a total height of 100 μm, core dimension of 35 μm × 35 μm, and pitch of 250 μm. The 12-channel polymer waveguide array is 10 mm long and 3 mm wide. A standard 12-channel MT optical connector on one end of the waveguide array connects to a fiber optical cable to couple light in and out of the package. An array of 45° mirrors on the other end of the waveguide bend the optical signal 90° in order to couple into and out of the VCSEL and photodiode arrays, which are flip-chip bonded face down onto the package. The 45° mirror cut is formed either by microtome or laser ablation and its reverse side is metalized. The loss from this 45° mirror is 0.3 dB. The high-speed electrical lines used to connect the CMOS chip to the optical

Current VCSEL technology is rated for 10 Gb/s, beyond which the VCSELs are bandwidth-limited with a slow transient tail due to intrinsic and extrinsic parasitic effects such as carrier diffusion and device parasitic capacitance. Pre-emphasis can compensate for these effects and increase the achievable data rate.

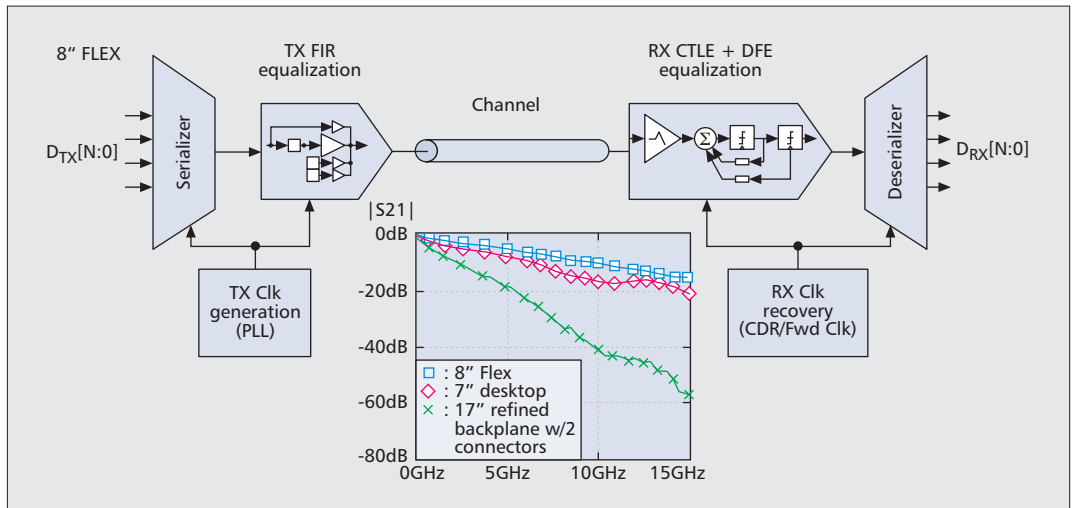


Figure 2. High-speed electrical link block diagram showing serializer, TX PLL, TX finite impulse response (FIR) equalizer, RX continuous-time linear equalizer (CTLE), RX decision feedback equalizer (DFE), CDR, and deserialize.

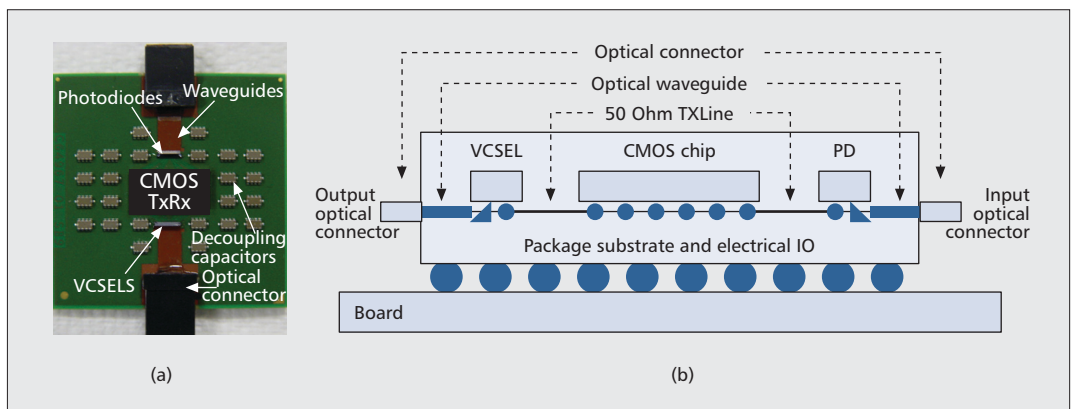


Figure 3. a) A fully assembled optical transceiver unit; b) a schematic side view of the same unit, showing the optical coupling scheme of VCSELS/photo-detectors to waveguides through a 45° mirror.

components are routed as controlled impedance (50Ω single-ended or 90Ω differential) microstrip traces on the top surface of the substrate where they have the best high-frequency characteristics. The optical signal is transmitted by an oxide-confined 850-nm 10 Gb/s 1 × 12 VCSEL array with peak optical output power greater than 3 mW (~5 dBm) and received by a 10 Gb/s, 1 × 12 GaAs PIN detector array with a diameter of 75 μm, a capacitance of 330 fF, a 3 dB bandwidth of 8 GHz, and a responsivity of 0.6 A/W. The total optical loss budget for the end-to-end link includes VCSEL and photodiode coupling loss through the 45° mirrors, propagation loss through the waveguide, MT connector loss, and Fresnel losses at the interfaces in the connectors. The total optical loss budget calculated for the complete link is 10 dB [2]. Improvements in optical coupling for this hybrid package architecture are in development to reduce the optical loss budget to as low as 6.8 dB.

CIRCUITS: VCSEL DRIVER AND TRANSIMPEDANCE AMPLIFIER

Current VCSEL technology is rated for 10 Gb/s, beyond which the VCSELs are bandwidth-limit-

ed with a slow transient tail due to intrinsic and extrinsic parasitic effects such as carrier diffusion and device parasitic capacitance. Pre-emphasis can compensate for these effects and increase the achievable data rate. The VCSEL driver described in [1] directly generates dual-edge pre-emphasis with sub-bit-period pre-emphasis waveform timing precision. The 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. Typical average currents provided to the VCSELs range from 6 to 10 mA, which corresponds to an average optical power of 1.5 to 2 mW. The VCSEL driver is output terminated and connected to the VCSEL through a 50Ω microstrip transmission line routed on the top surface of the package. As the VCSEL technology develops for higher modulation speed (using quantum dots rather than quantum well technology), high-data-rate VCSELs at 20 Gb/s and higher will still benefit from these pre-emphasis techniques to further extend data rates.

The transimpedance amplifier (TIA) uses the differential symmetric-feedback topology [1],

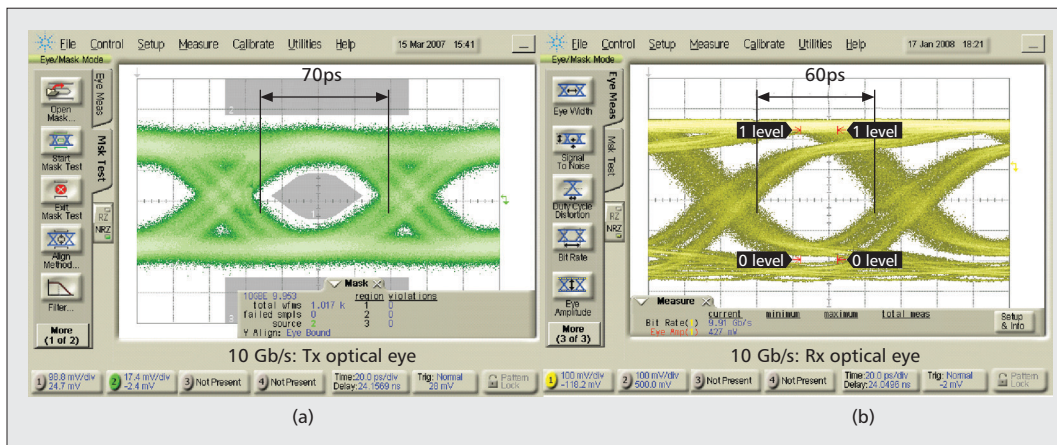


Figure 4. Optical eye diagrams for 10 Gb/s tested with fully packaged: a) transmitter optical output; b) receiver optical input.

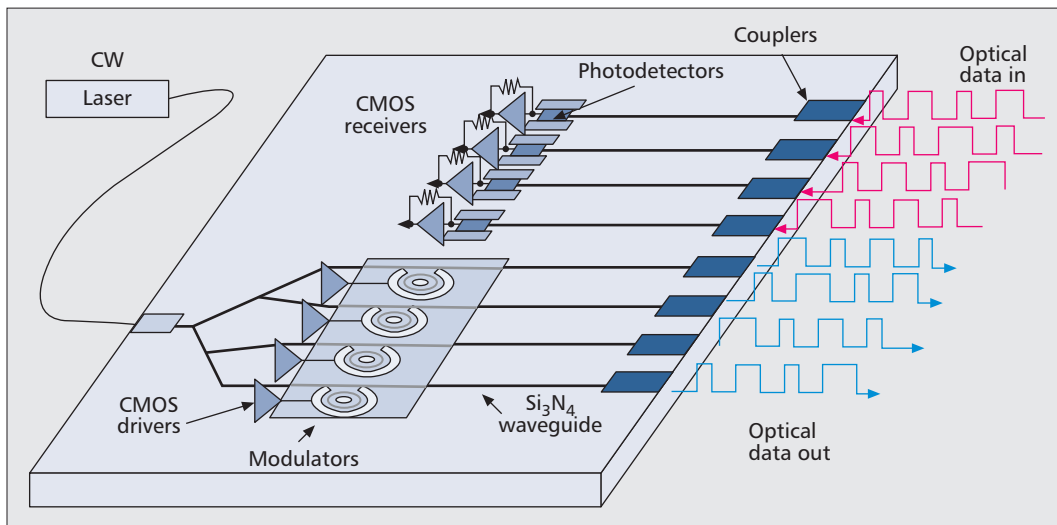


Figure 5. Photonics optical interconnect architecture.

which converts the single-ended input current to a differential output voltage to help mitigate supply noise at subsequent gain stages and provides a data rate above 12.5 Gb/s when the total input parasitic capacitance C_p is less than 250 fF. The TIA receives a single-ended photocurrent of 200 μ A from the photodiode and generates a differential 2×50 mV_{pp} output that is fed to a following limiting amplifier (LA), which converts it to a CML level output. The LA consists of a cascade of CML buffers. In the packaged transceiver the combined capacitance of the photodiode, metal pad, bump, and ESD could be as high as 500 fF. This capacitance limits the maximum data rate that can be measured for the packaged receiver channel. The same TIA tested electrically with wafer probing had an open electrical eye diagram at 18 Gb/s for an input capacitance of 90 fF. This indicates there is a strong dependence of bandwidth on the input parasitic capacitance.

EXPERIMENTAL RESULTS

10 Gb/s optical measurement results are shown in Fig. 4 for a fully assembled transmitter and receiver. For the transmitter measurement, external differential electrical pseudorandom data was

sourced into the chip to drive the CMOS pre-emphasis VCSEL driver, and the VCSELs were biased with an average current of 7 mA. The measured transmitter optical eye opening was 70 ps. The receiver demonstrated an open electrical eye for optical 10 Gb/s input data. The electrical received signal eye opening was 60 ps with a peak-to-peak jitter of 30 ps. The individual transmitter and TIA receiver circuits are capable of operation at up to 18 Gb/s [2].

PHOTONIC CMOS OPTICAL I/O

ARCHITECTURE

In the longer term, monolithic integration of photonic elements in a CMOS process will enable significant improvements in I/O performance, energy efficiency, and cost. The proposed monolithic photonic CMOS process, illustrated in Fig. 5, integrates modulators, waveguides, and detectors on top of the metal interconnect layers in the far back-end of a standard CMOS process. Light from a continuous-wave (CW) source is coupled onto the die and modulated using integrated waveguide-based modulators driven by on-chip circuits, such that the electrical signals do not

In a photonic CMOS process for integrated optical links, the additional process steps required for photonics must not degrade or interfere with the front-end CMOS transistor performance. Furthermore, the process must allow fabrication of all required optical components on the same die.




Figure 6. a) Schematic of top view of full on-die optical link showing bus waveguide connecting modulator to photo detector; b) cross section SEM image (along the dotted line in 6a) showing optical components in one piece of silicon.

leave the die. The modulated light is coupled off the die through a fiber or waveguide to a receiving chip, where it is coupled through an integrated waveguide into a compact photodetector. The photodetector output current is converted to a full-swing electrical signal by a TIA and an LA.

Monolithic integration of photonics onto the microprocessor will reduce the power and the cost of I/O. Integration reduces the capacitive load on the driver and receiver circuits and leads to higher bandwidth and lower power. Parasitic capacitance is reduced because integration of the circuits and optical devices on the same die removes the bump, package, and ESD capacitance from the signal path. The intrinsic device capacitance of integrated optical components is smaller than the capacitance of discrete alternatives. Static power consumption is reduced because small integrated optical devices do not require termination, in comparison to larger discrete alternatives such as Mach-Zender interferometers which require 50 Ohm termination for high-speed operation. Cost is reduced by decreasing the required number of discrete optical components.

In a photonic CMOS process for integrated optical links, the additional process steps required for photonics must not degrade or interfere with the front-end CMOS transistor performance. Furthermore, the process must allow fabrication of all required optical components on the same die. The optical components in previously demonstrated integrated optical links were fabricated in the front-end of a semiconductor-on-insulator (SOI) CMOS process [3], which constrains the transistor processing. The presented experimental photonic process is based on a silicon nitride single-mode waveguide with silicon dioxide cladding and provides waveguides, electro-optic (EO) polymer ring resonator (RR) modulators [2, 4], and waveguide-embedded metal-semiconductor-metal (MSM) detectors fabricated from polycrystalline germanium on the

bulk CMOS (not SOI) process back-end compatible silicon-dioxide dielectric [4, 5].

FABRICATION

The photonic elements are added to the CMOS process in the metal interconnect fabrication back-end section of the process after all the high temperature front-end processing of transistors is completed. The waveguides are formed with a 450 nm silicon nitride layer deposited by plasma-enhanced chemical vapor deposition (PECVD) on the SiO₂ interlayer dielectric (ILD) and patterned with photolithography and plasma dry etch. This shared waveguide layer is used to build all the waveguides, RRs, and coupling waveguides for the active electro-optic devices. After patterning the waveguides, silicon dioxide cladding is deposited, and three subsequent lithography steps define the detector regions, the electrodes for all active devices, and modulator regions. The photodetector regions are filled with polycrystalline germanium in a damascene process, the detector electrodes are formed in a standard copper damascene process, and then the modulators are formed by depositing EO polymer cladding over the ring resonators in the regions defined for the modulators. The additional cost to add photonic devices to the CMOS process is low, since only four additional photolithography steps are required.

Figure 6 shows both a top view and an SEM cross-section of the modulator, waveguide, and detector constituting a complete optical link. A single patterned silicon nitride layer forms all of the waveguides in the active and passive components. Similarly, one metal layer forms all the electrodes for both the modulator and photodetector. Furthermore, this optical layer is compatible with standard microprocessor CMOS as it is created on an amorphous ILD and can therefore be fabricated in the back-end metal interconnect section of the CMOS process. In order to stay within the thermal budget for standard back-end processing, all steps in the process flow must occur below ~450°C.

EXPERIMENTAL RESULTS

Waveguide — The waveguide is the foundation for the proposed photonic CMOS technology [6]. The waveguide is processed as a 450 nm PECVD silicon nitride film deposited on a 2 μm silicon dioxide undercladding layer at 400°C. The waveguide is patterned using conventional 248 nm lithography and plasma etching. Loss measurements at 1310 nm using the cut-back method show that the silicon nitride waveguide loss is ~1 dB/cm for waveguides with a width of 0.5 μm. This loss is sufficiently low for on-die applications where the total waveguide length is on the order of 1 cm.

Modulator — The electro-optic cladding RR modulator and photodetector share the high index contrast waveguide fabrication process. The modulator design is based on a high-performance ring resonator built with a silicon nitride waveguide and ring. Copper damascene electrodes are fabricated around the ring, and the top cladding is removed and replaced with the EO polymer. This work uses a proprietary chromophore-doped EO polymer [6]. The modulator design is optimized for a quality factor (Q) between 5000 and 10,000: high enough that a small resonance shift

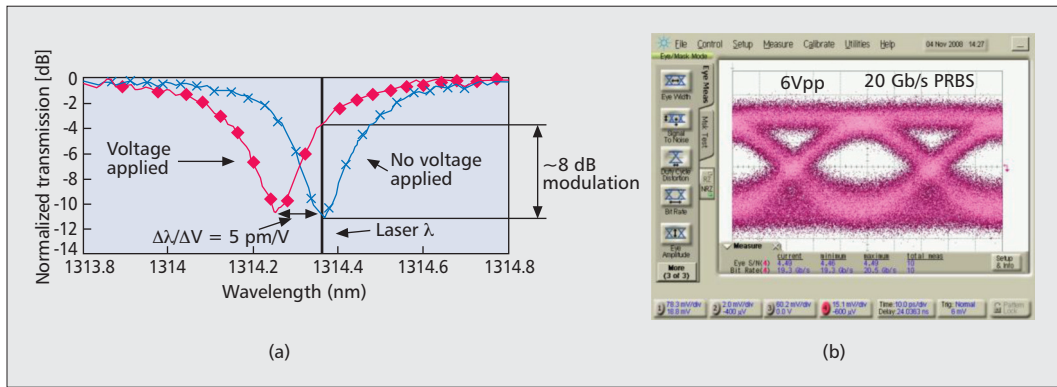


Figure 7. a) Resonance spectra obtained with -20 V (\blacklozenge) and 20 V (\times) bias on the EO modulator; b) 20 Gb/s PRBS eye diagram of EO polymer modulator.

results in a large modulation depth, but not so high that the modulator is unable to switch at high data rates. The EO polymer is poled before wafer processing is completed using an electric field of 100 V/cm around the glass temperature of 143°C . The electrodes have a $4.5\text{ }\mu\text{m}$ gap centered around the waveguide ring, and the ring has a radius of $28\text{ }\mu\text{m}$. An SEM image of the modulator is shown on the right of Fig. 6b. The resonance spectrum of a typical modulator under $+20\text{ V}$ and -20 V bias is shown in Fig. 7a. The resonance shift calculated with a linear fit to the resonance frequencies measured at $+20\text{ V}$ and -20 V bias is 5 pm/V . The measured Q was ~ 7000 , and the resonance depth was $\sim 11\text{ dB}$. The highest measured modulation depth for a 10 GHz clock input with a 6 V swing was 8 dB . A 20 Gb/s pseudo-random binary sequence (PRBS) eye diagram for a typical device was measured [4] (Fig. 7b).

Photodetector — Unlike a PIN detector, the lateral MSM detector requires only one lithography step to form the contacts. An evanescently coupled waveguide, shown on the left of Fig. 6b, efficiently couples the light into the absorbing active material of the photodetector. The polycrystalline germanium in the detector was deposited by CVD processing at 600°C . Fabrication of a photodiode from polycrystalline germanium deposited on ILD is an important step toward compatibility with a standard back-end (BE) CMOS process. Measurements at a higher frequency showed that 40 Gb/s operation is within reach. To improve the noise performance, band-gap engineering can be used to create a Schottky barrier at the metal/germanium contact in order to reduce the dark current further. Another important step toward BE compatibility is lowering the process temperature. Devices were fabricated using PVD Ge at 350°C , and the best measured devices had a dark current of $77\text{ }\mu\text{A}$ with open PRBS eyes at both 20 and 40 Gb/s [4].

OPTICAL LINK MODELING AND COMPARISONS

The optical I/O link power efficiency is a strong function of the received optical power, which is determined by the transmit power and the link optical loss budget. A feasible best case value for

the hybrid optical link budget is -6.8 dB with some packaging improvements. This is dominated by coupling losses from the VCSEL and photodetector to the multimode fiber (MMF) and the finite extinction ratio penalty. The hybrid optical I/O link budget is calculated using the following assumptions:

Average TX power	3.0 dBm
VCSEL to MMF coupling	-1.1 dB
MMF to photodetector coupling	-1.1 dB
Extinction ratio (7.3 dB) penalty	-1.6 dB
Margin	-3.0 dB
Link budget	-6.8 dB
Required RX sensitivity	-3.8 dBm

The integrated optical link budget is nearly 9 dB worse than the hybrid optical link budget due to the coupling loss between the off-chip single-mode fiber and the on-chip single-mode waveguide, and the extra coupling loss from the off-chip CW laser. However, the integrated photodetector's ultra-low capacitance allows the integrated optical receiver to achieve approximately 13 dB of sensitivity improvement at the same bandwidth, which results in significant system power savings. The integrated optical I/O link budget is calculated using the following assumptions:

Average VCSEL TX power	3.0 dBm
Source laser to SMF coupling	-2.0 dB
SMF to modulator coupling	-2.0 dB
Modulator loss	-2.0 dB
Modulator to SMF coupling	-2.0 dB
SMF to photodetector coupling	-3.0 dB
Extinction ratio (8.0 dB) penalty	-1.4 dB
Margin	-3.0 dB
Link budget	-15.4 dB
Required RX sensitivity	-12.4 dBm

Circuit simulation-based power efficiency estimates of both transmit and receive front-end circuits for these two optical I/O architectures was performed for CMOS technologies starting from a 45 nm node and ending with a predictive 16 nm CMOS node [1]. A current-mode VCSEL driver and a simple CMOS inverter-based voltage-mode modulator driver are modeled for the hybrid and integrated optical systems, respectively. In both systems a TIA is followed by simple differential-pair LA stages to realize the optical receiver. The models are constructed with the circuits optimized to provide the minimum bandwidth necessary for a particular data rate, and

Measurements at a higher frequency showed that 40 Gb/s operation is within reach. To improve the noise performance, band-gap engineering can be used to create a Schottky barrier at the metal/germanium contact in order to reduce the dark current further.

The comparison reveals that the hybrid optical architecture is equal to or better in power efficiency than both the electrical backplane channel and the desktop channel at data rates near where RX equalization becomes necessary.

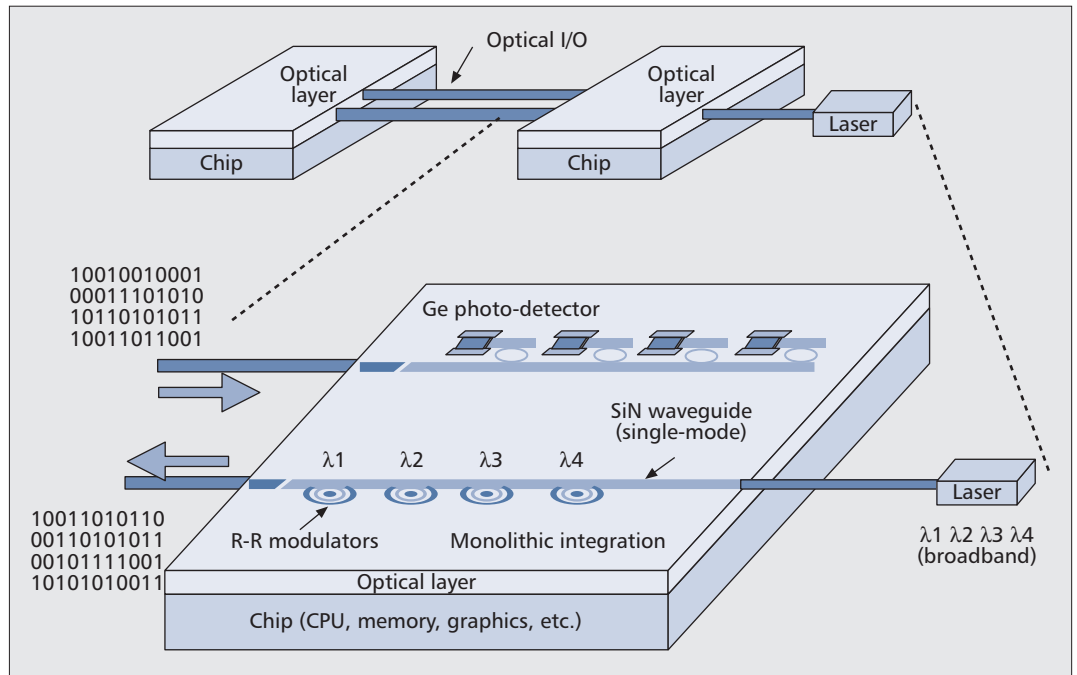


Figure 8. Photonic CMOS enabled wavelength-division multiplexing architecture.

thus approximate a power optimal solution. The hybrid optical link power efficiency initially improves as the data rate increases due to the assumed-constant 3 dBm optical power from the 850 nm VCSEL. Power efficiency degrades from the optimum at higher data rates due to the optical RX amplifier gain-bandwidth requirements. As technology scales, this optimum occurs at a higher data rate due to the increased transistor f_T . This analysis predicts that hybrid optical data transmission at less than 1 pJ/b will be realized in the future. Assuming a 1310 nm CW laser source with 3 dBm optical TX power, the integrated optical link power efficiency displays similar behavior at a much lower power level due to low capacitance of the modulator and photodetector allowing for very efficient optical drivers and receivers. Ultra-low receiver input capacitance enables a TIA-based receiver without any LA stages to provide sufficient sensitivity at data rates exceeding 30 Gb/s. The data rate at which extra LA stages become necessary scales with the improved CMOS technology f_T . These projections indicate that photonic CMOS could enable integrated optical interconnect to reach 0.3 pJ/b.

The power-performance analysis of the hybrid optical link was compared with electrical link systems that employ all three electrical channels discussed early in the electrical link analysis. The comparison reveals that the hybrid optical architecture is equal to or better in power efficiency than both the electrical backplane and desktop channels at data rates near where RX equalization becomes necessary. This data rate is dependent on the channel loss characteristics and is 13 Gb/s and 19 Gb/s for the 17 in backplane and 7 in desktop channels, respectively. While the hybrid optical link cannot outperform the high-performance electrical cable channel at the 45 nm node, the increased gain-bandwidth offered by the 16 nm node allows the hybrid optical link to become comparable near 40

Gb/s. Note that this assumes the availability of 40 Gb/s-class VCSELs, which are currently emerging from research [7]. The reduced parasitics offered by the integrated photonics with CMOS optical architecture allows it to achieve superior power efficiency over the majority of data rates compared to the three electrical channels as well as the hybrid optical architecture. This assumes further improvements in modulator EO polymer performance to enable sufficient optical modulation depth at voltage modulation levels compatible with CMOS inverter-based drivers [4, 6].

FUTURE DIRECTIONS

As CMOS scaling continues in the future, larger numbers of CPU cores will be integrated on the microprocessor chip, and it will become necessary to provide interconnect scaling to higher bandwidth between cores on chip, and between these cores and the off-chip DRAM. Wavelength-division multiplexed (WDM) links transmit multiple wavelengths through the same waveguide in order to increase the aggregate optical data transmission. A photonic CMOS architecture for optical WDM of signals monolithically integrated on chip is shown in Fig. 8. The RR modulator selectively modulates a single wavelength from a multiwavelength source, and eliminates the need for separate optical demultiplexers and multiplexers. At the receiver, passive RR optical filters can demultiplex the optical data by selecting a single unique wavelength for detection at each photodetector. Since the photonic CMOS RR modulators have such a narrow tuning range (Fig. 7), the WDM wavelengths can be spaced at less than 1 nm (100 GHz in optical frequency with a reference of 230 THz). Thus, the RR technology provides the means for bandwidth to scale by adding more wavelengths to each waveguide channel [1].

SUMMARY

This work provides a comparison of electrical I/O to optical I/O for chip-to-chip interconnect. While electrical interconnect will continue to use more sophisticated equalization techniques to overcome the loss of the interconnect channel, the high data rate and long interconnect lengths required by future many-core processors will require the introduction of optical interconnect. Optical interconnect for CPUs will first be introduced with optical package-to-package I/O using hybrid MCM single-package technology. In the long term, monolithic integration of optical components will provide terabyte-per-second interconnect data rates with the required energy efficiency at less than 1 pJ/bit.

REFERENCES

- [1] I. A. Young *et al.*, "Optical I/O Technology for Tera-Scale Computing," *IEEE JSSC*, Jan. 2010, pp. 235–48.
- [2] E. Mohammed *et al.*, "Optical Hybrid Package with an 8-Channel 18GT/s CMOS Transceiver for Chip-to-Chip Optical Interconnect," *Proc. SPIE*, vol. 6899, Feb. 2008.
- [3] B. Analui *et al.*, "A Fully Integrated 20-Gb/s Optoelectronic Transceiver Implemented in a Standard 0.13m CMOS SOI Technology," *IEEE JSSC*, Dec. 2006, pp. 2945–55.
- [4] I. A. Young *et al.*, "Integration of Nano-Photonic Devices for CMOS Chip-to-Chip Optical I/O," *CLEO/QELS Conf. Digest Tech. Papers*, May 16–21, 2010.
- [5] M. R. Reshotko *et al.*, "Waveguide Coupled Ge-on-Oxide Photodetectors for Integrated Optical Links," *IEEE/LEOS Int'l. Conf. Group IV Photonics*, Sept. 17–19, 2008, pp. 182–84.
- [6] B. A. Block *et al.*, "Electro-Optic Polymer Cladding Ring Resonator Modulators," *Optics Express*, vol. 16, no. 22, Oct. 2008, pp. 18326–33.
- [7] T. Anan *et al.*, "High-Speed 1.1-mm-Range InGaAs VCSELs," *Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Feb. 24–28, 2008, pp. 1–3.

BIOGRAPHIES

IAN A. YOUNG [M'78, SM'96, F'99] (ian.young@intel.com) is a senior fellow and director of Advanced Circuits and Technology Integration in the Technology and Manufacturing Group at Intel Corporation. He does research and development of mixed-signal circuits for microprocessor and SOC products along with process technology development. He joined Intel in 1983. Starting with the development of circuits for a 1 Mb DRAM, he then led the design of three generations of SRAM products and manufacturing test vehicles, and developed the original PLL-based clocking circuit in a microprocessor while working on the 50 MHz Intel486 processor design. He subsequently developed the core PLL clocking circuit building blocks used in each generation of Intel microprocessors through the 0.13 μm 3.2 GHz Pentium 4. He is currently directing the research and development of analog mixed-signal and RF circuits for high-speed serial I/O in advanced logic process, RF wireless circuits in advanced SOC process, and researching chip-to-chip optical I/O technology. He received his Bachelor's and Master's degrees in electrical engineering from the University of Melbourne, Australia, in 1972 and 1975. He received his Ph.D. in electrical engineering from the University of California, Berkeley in 1978. He was a member of the Symposium on VLSI Circuits Technical Program Committee from 1991 to 1996, serving as Program Committee Chairman in 1995–1996 and Symposium Chairman in 1997–1998. He was a member of the ISSCC Technical Program Committee from 1992 to 2005, serving as Digital Subcommittee Chair from 1997 to 2003, Technical Program Committee Vice-Chair in 2004, and Chair in 2005.

EDRIS MOHAMMED is a research engineer at Intel Corporation responsible for system design, integration, and characterization of hybrid optical interconnect at Intel's components research division. He joined Intel Corporation in 2001. Before joining Intel, he was a member of technical staff at Lucent Technologies where he worked as a test and reliability engineer for 980 nm pump laser diodes. His research interests include high-speed optoelectronic devices, package integra-

tion, and passive optical and electrical devices. He received a B.Sc. in physics from Addis Ababa University in 1987, an M.Sc. in physics from Florida A&M University in 1995, and an M.Sc. in electrical engineering (optoelectronics) and a Ph.D. in applied physics from Georgia Institute of Technology in 2000. Since 2005 he has been serving as an associate editor of optical interconnects for *SPIE Optical Engineering Journal*.

JASON T. S. LIAO [M'09] received his M.S. degree in communication engineering from National Chiao Tung University, Taiwan, in 1996 and his Ph.D. degree in electrical engineering from the University of California, San Diego in 2004. In 2005 he joined Intel Corporation, Hillsboro, Oregon, as a senior electrical engineer. His research focuses on high-speed optical I/O and receiver clocking techniques for high-speed electrical I/O links.

ALEXANDRA M. KERN [S'01, M'07] received A.B. and B.E. degrees from Dartmouth College in 2002, and S.M. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology in 2004 and 2007. In June 2007 she joined the Logic Technology Development group at Intel Corporation, where she is involved in the development of high-speed electrical serial links in advanced CMOS technology. Her research interests include electrical and optical interconnect for on-chip and chip-to-chip applications, and analog/mixed-signal circuit design techniques for deep-submicron CMOS.

SAMUEL PALERMO [S'98, M'07] received B.S. and M.S. degrees in electrical engineering from Texas A&M University, College Station in 1997 and 1999, respectively, and his Ph.D. degree in electrical engineering from Stanford University, California, in 2007. From 1999 to 2000, he was with Texas Instruments, Dallas, where he worked on the design of mixed-signal integrated circuits for high-speed serial data communication. From 2006 to 2008 he was with Intel Corporation, where he worked on high-speed optical and electrical I/O architectures. In 2009 he joined the Electrical and Computer Engineering Department of Texas A&M University, where he is currently an assistant professor. His research interests include high-speed electrical and optical links, clock recovery systems, and techniques for device variability compensation. He is a member of Eta Kappa Nu.

BRUCE A. BLOCK joined Intel in 1998 after receiving his Ph.D. degree in materials science and engineering from Northwestern University in 1997 and a B.S. degree in materials science and engineering from Cornell University in 1989. From 1989 to 1992 he worked for IBM Corporation, East Fishkill, New York, in the Advanced Packaging Laboratory. He has been working since 2000 on the process integration and characterization of CMOS-compatible optical devices in the Components Research Organization. He has developed many passive and active devices based on silicon nitride waveguides, including waveguide coupled germanium photodetectors, electro-optic polymer modulators, ring resonator filters, and surface plasmon polariton-based polarizers and detectors.

MIRIAM R. RESHOTKO [M'07] joined Intel's Components Research department in 2001 after completing a Ph.D. degree in physics at the Hebrew University of Jerusalem, Israel. She received her M.Sc. degree in physics from the same university, and her B.A. in physics from Cornell University. Her current research centers on integrated CMOS compatible optoelectronic devices for optical interconnects. Specific areas of focus are development of monolithically integrated high-speed waveguide coupled photodetectors for various interconnect applications, and process integration for discrete optical devices, including waveguides, modulators, photodetectors, and optical links.

PETER L. D. CHANG [M'92] received his B.S. degree in physics from National Taiwan University in 1977 and his Ph.D. in theoretical solid state physics from the State University of New York at Stony Brook in 1985. From 1985 to 1991 he conducted research in the Physics and ECE Departments of the University of California, Santa Barbara. From 1991 to 1994 he was with Lockheed Sanders working on MMIC. He joined Intel in 1994, developing 0.25 μm flash memory technology. Since 1996 he has been involved in R&D and high volume manufacturing of logic and SRAM process. He is currently with the Components Research of Technology, Manufacturing and Enterprise Services Group. His research interests are in high-density memory and optical interconnects for high memory bandwidth required for future generations of CPUs.

Optical interconnect for CPUs will first be introduced with optical package-to-package I/O using hybrid MCM single-package technology. In the long term, monolithic integration of optical components will provide TB/s interconnect data rates with the required energy efficiency at less than 1 pJ/bit.