Optical I/O Technology for Tera-Scale Computing

Ian A. Young, *Fellow, IEEE*, Edris Mohammed, Jason T. S. Liao, *Member, IEEE*, Alexandra M. Kern, *Member, IEEE*, Samuel Palermo, *Member, IEEE*, Bruce A. Block, Miriam R. Reshotko, *Member, IEEE*, and Peter L.D. Chang, *Member, IEEE*

***Abstract*—This paper describes both a near term and a long term optical interconnect solution, the first based on a packaging architecture and the second based on a monolithic photonic CMOS architecture. The packaging-based optical I/O architecture implemented with 90 nm CMOS transceiver circuits, 1×12 VCSEL/detector arrays and polymer waveguides achieves 10 Gb/s/channel at 11 pJ/b. A simple TX pre-emphasis technique enables a potential 18 Gb/s at 9.6 pJ/b link efficiency. Analysis predicts this architecture to reach less than 1 pJ/b at the 16 nm CMOS technology node. A photonic CMOS process enables higher bandwidth and lower energy-per-bit 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 modulators and Ge detectors demonstrate performance above 20 Gb/s and analysis predicts that photonic CMOS will eventually enable energy efficiency better than 0.3 pJ/b with 16 nm CMOS. Optical interconnect technologies such as these using multi-lane communication or wavelength division multiplexing have the potential to achieve TB/s interconnect and enable platforms suitable for the tera-scale computing era.**

***Index Terms*—Computers, interconnect, optical, chip-to-chip, I/O, photonic CMOS, bandwidth, energy efficiency, modulator, photodetector, VCSEL.**

I. Introduction

THE microprocessor architecture transition from multi-core to many-core will drive increased chip-to-chip I/O bandwidth demands at processor/memory interfaces and in multi-processor systems. Near-term projections estimate the CPU-to-memory bandwidth to be 100 GB/s in 2012–2013 [1]. Future architectures will require bandwidths from 200 GB/s to 1.0 TB/s and begin the era of tera-scale computing.

To meet these bandwidth demands, traditional electrical interconnect techniques require increased circuit complexity and costlier materials. However, without low-loss electrical interconnects, increasing I/O bandwidth in electrical links eventually comes at the cost of reducing interconnect link length, reducing signal integrity, or increasing power consumption. Optical interconnects with the potential for terahertz bandwidth, low loss, and low crosstalk have been proposed to replace electrical interconnects between chips [2].

Chip-to-chip optical interconnects have negligible frequency-dependent loss and, unlike electrical interconnects, require little or no equalization. This motivates I/O architects to consider optical I/O as a means of scaling data rates in a power efficient manner. This paper presents analysis demonstrating the need for optical interconnect, describes experimental results for both near-term and longer-term optical link technologies, and predicts how the relative power and performance of optical and electrical links will scale with technology and data rate.

The near-term and long-term chip-to-chip optical interconnect architectures and techniques described in the paper have potential data rates suitable for tera-scale computing applications. To motivate the need for optical interconnects, Section II discusses the design challenges of increasing electrical interconnect data rate with an analysis of circuit complexity and power efficiency versus data rate in both a current 45 nm and a predictive 16 nm CMOS technology. Section III outlines a near-term optical interconnect architecture which uses a microprocessor organic land grid array (OLGA) package MCM technology that combines GaAs optical sources, detectors and waveguides with a CMOS chip to enable a hybrid optical chip-to-chip I/O. Section IV describes an architecture and process technology to integrate the photonic elements monolithically with the CMOS drivers and receivers. The reduced parasitic capacitance achieved through monolithic integration will enable improved speed and I/O power efficiency. Section V presents projections for the power efficiency and data rate of the hybrid and integrated optical transceivers, and compares this with electrical alternatives. Section VI concludes the paper and proposes a future optical interconnect for on-chip networks that maximizes the link performance by using wavelength division multiplexing (WDM) to multiplex multiple wavelengths of modulated data onto a single fiber or waveguide.

II. Electrical Link Issues

Fig. 1 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 sampling clocks. The receiver timing system is less complex for a periodically-trained, forwarded-clock architecture [3], [4] and more complex for a continuously-tracking, embedded-clock architecture [5]. The design complexity of the transmitter and receiver increases to

Manuscript received May 13, 2009; revised August 09, 2009. Current version published December 23, 2009. This paper was approved by Guest Editor Kevin Zhang.

I. A. Young, E. Mohammed, J. T. S. Liao, A. M. Kern, B. A. Block, M. R. Reshotko, and P. L. D. Chang are with Intel Corporation, Hillsboro, OR 97124 USA (e-mail: ian.young@intel.com).

S. Palermo was with Intel Corporation, Hillsboro, OR 97124 USA, and is now with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2009.2034444

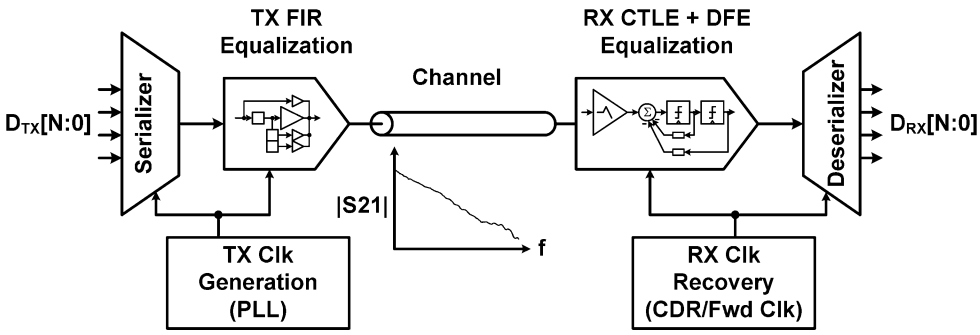

Fig. 1. High-speed electrical link block diagram.

include equalization circuitry as data rates scale past electrical channel bandwidths.

Electrical channel frequency characteristics are dependent on channel length and channel quality. Fig. 2 shows the channel response for three typical electrical channels, a 17 inch server backplane channel with two connectors, a 7 inch desktop channel with no connectors, and an 8 inch high-performance cable channel. The low-pass, frequency-dependent loss is exponential with channel length, as illustrated by the loss difference between a 17 inch backplane channel and a 7 inch 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) FIR filtering (sometimes called feed-forward equalization (FFE)) [6], [7], continuous-time linear equalizers (CTLE) [4], and receiver (RX) decision-feedback equalization (DFE) [8]–[10]. Results from a statistical BER simulator [6], shown in Fig. 3, illustrate the increased circuit complexity required for conventional non-return-to-zero signaling at higher data rates and/or over channels with higher loss. Over the illustrated channels, equalization options include one to four tap TX FIR (TX1–TX4), CTLE, and one to eight tap RX DFE (DFE1–DFE8). The refined 17 inch backplane channel requires equalization to achieve greater than 6 Gb/s and, even with significant equalization complexity, fails to achieve greater than 20 Gb/s due to the high channel loss. However, the 8 inch high-performance cable channel can achieve 20 Gb/s without equalization and potentially higher than 50 Gb/s with substantial equalization.

Even when equalization can theoretically compensate channel loss, circuit technology can limit the maximum speed at which the equalization systems operate in an energy efficient manner. Fig. 4(a) shows simulation-based power efficiency estimates of transmit and receive front-end circuits, excluding the timing systems, in a 45 nm CMOS technology. Since a constant 1 V_{pp} transmit signal is assumed, the power efficiency initially improves as the data rate increases. This trend reverses and power efficiency begins to decline with data rate as increasingly-complex equalization becomes necessary. While transmit equalization can be implemented with little additional energy, a CTLE with sufficient gain-bandwidth requires significant power, so the energy efficiency degrades rapidly once a CTLE becomes necessary. Ultimately, the maximum data rates in all three systems are limited by circuit speed, as the 45 nm technology cannot support efficient DFEs in the 20 Gb/s range due to their critical timing path [8], [10].

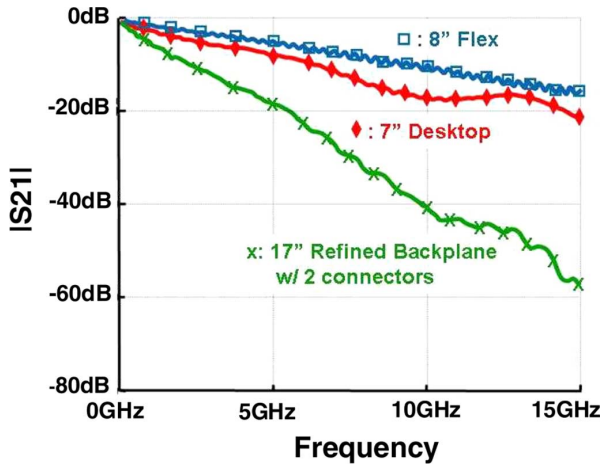

Fig. 2. Electrical channels frequency response.

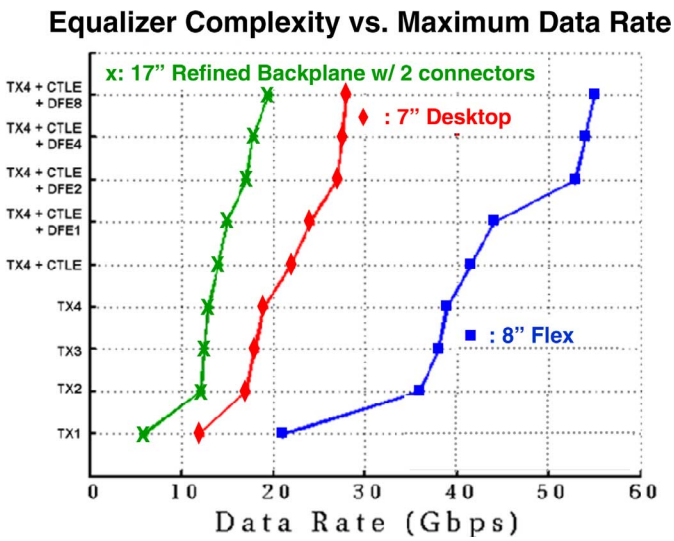

Fig. 3. Equalization complexity required versus maximum data rate ($\mathrm{BER} = 10^{-12}$) for three electrical channels.

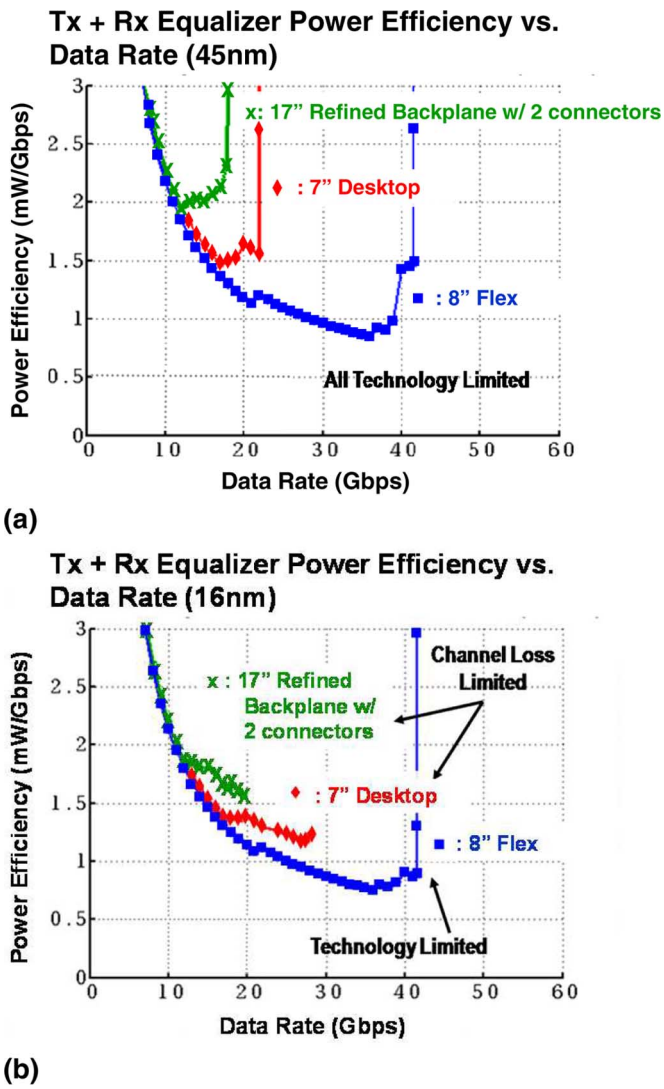

Fig. 4. Circuit simulation-based power efficiency estimates of transmit and receive front-end circuits versus data rate. (a) 45 nm CMOS technology. (b) Predictively modeled 16 nm CMOS technology.

Estimates based on a predictive 16 nm technology node, shown in Fig. 4(b), reveal that 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. Channel loss, transmit peak power constraints, receiver sensitivity, and jitter eventually limit the maximum data rate at which the desired 10^{-12} BER can be achieved in backplane and desktop channels, despite significant equalization. The high-performance 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.

III. Optical I/O Implementation Using a Hybrid MCM Package

In the near term, a 12-channel optical transceiver architecture is proposed which allows integration of low-cost, high-performance optical components with existing microprocessor

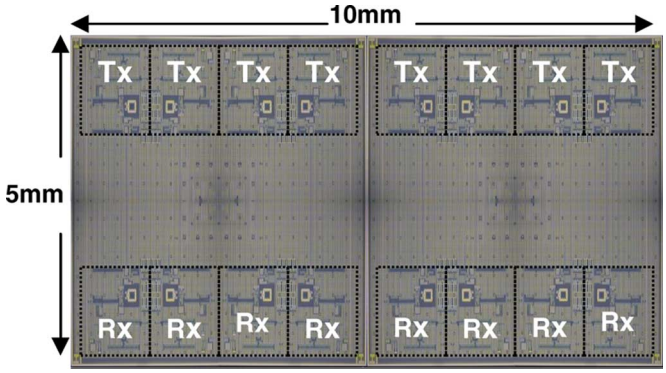

Fig. 5. Transceiver chip die photo [19].

package technology. This hybrid architecture integrates CMOS and discrete optical components in a multi-chip module (MCM) package. In this design, a 90 nm CMOS multi-channel optical transceiver chip, an 850 nm 10 Gb/s GaAs vertical-cavity surface-emitting laser (VCSEL) 1×12 linear array and a PIN photodiode 1×12 linear array are flip-chip mounted on a standard microprocessor 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 multi-terminal (MT) fiber optic connectors, which connect to 1×12 waveguide or fiber arrays to couple the light off-chip.

A. The Transceiver Chip Architecture

The transceiver chip shown in Fig. 5 was fabricated in a 90 nm digital CMOS process with seven metal layers on high-resistivity substrate [11]. The chip is 5 mm $\times$ 10 mm and has sixteen individual 10 Gb/s transceiver channels arranged in two 1×8 ports [12]. While standard optical connectors support 1×12 ports, die size considerations limited this design to 1×8, but 1×12 would be possible in a 45 nm CMOS technology. Each channel (Fig. 6) contains a VCSEL driver, a transimpedance amplifier (TIA) and limiting amplifier (LA), clock and data recovery (CDR) with a phase-frequency detector (PFD) for frequency acquisition and a bang-bang phase detector for phase acquisition for the receiver, and a pseudorandom bit-pattern sequence (PRBS) generator and BER tester (BERT) for self test. While all channels contain complete transmit and receive circuits, each one is programmed with the scan chain to be either a transmitter or receiver. A few receiver channels are implemented without a CDR in order to enable characterization of the TIA and LA path without re-timing. In contrast to other recent work, all high-speed circuits and logic functions are implemented with current-mode logic (CML) without inductive peaking which saves die area [13], [14]. The center of the chip contains circuitry for the test interface, reference clock distribution, and a bandgap reference. Three separate power supplies were used for the core logic (1.2 V), the LC-VCO (1.4 V), and

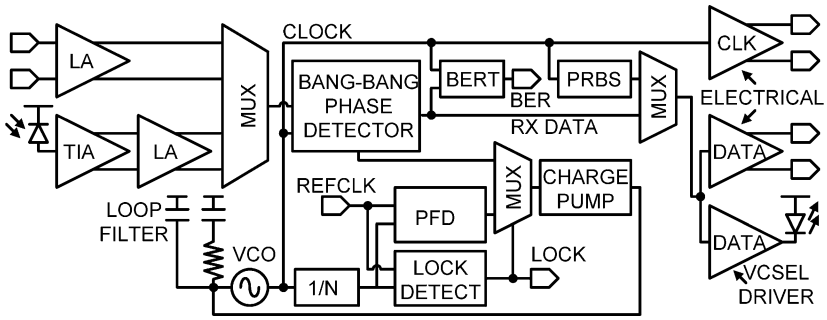

Fig. 6. Block diagram of the Tx/Rx channel architecture.

TABLE I
POWER EFFICIENCY BREAKDOWN OF THE ELEMENTS IN OPTICAL I/O

Data Rate (Gb/s)	**Tx**	Current (mA)	Voltage (V)	Power (mW)
10	Predriver	23.9	1.2	28.7
	VCSEL Driver	12.9	1.2	15.5
	VCSEL	6.2	3.0	18.7
	Total			62.9
	Power Efficiency **Tx** (mW/Gb/s)	6.3		
	Rx			
	TIA	7	1.8	12.6
	LIA	20	1.2	24.0
	Total Power			36.6
	Power Efficiency **Rx** (mW/Gb/s)	3.7		
	Power Efficiency Tx + Rx (mW/Gb/s)	**10.0**		
18	**Tx**			
	Predriver	63.6	1.2	76.3
	VCSEL Driver	32.5	1.2	39.0
	VCSEL	6.1	3.0	18.3
	Total			133.6
	Power Efficiency **Tx** (mW/Gb/s)	7.4		
	Rx			
	TIA	7.8	1.8	14.0
	LIA	20	1.2	24.0
	Total Power			38.0
	Power Efficiency Rx (mW/Gb/s)	2.1		
	Power Efficiency Tx + Rx (mW/Gb/s)	**9.54**		

the TIA and VCSEL driver (1.8 V). An off-chip supply (3.0 V) is used to bias the anode of the VCSEL. Table I summarizes the power breakdown for each of the elements of the hybrid MCM package optical I/O.

B. VCSEL Driver

The VCSELs in this technology development vehicle are rated for 10 Gb/s. Beyond 10 Gb/s they are bandwidth-limited with a slow transient tail due to intrinsic and extrinsic parasitic effects such as carrier diffusion [15] and device capacitance of 700 fF. Pre-emphasis can compensate for these effects and increase the achievable data rate. The VCSEL driver [16], shown in Fig. 7(a), directly generates dual-edge pre-emphasis with sub-bit-period pre-emphasis waveform timing precision. The pre-emphasized current waveform is generated, as shown in Fig. 7(b), by summing the main modulation current with a delayed and weighted peaking current in order to produce pre-emphasis pulses at each data transition. The VCSEL supply voltage (3.0 V) is independent from the transceiver core or driver supply to allow testing flexibility. Two 5-bit digital-to-analog converters (DACs) provide independent digital control of the output currents for the main and pre-emphasis drivers. Typical average currents provided to the VCSELs range from 6 mA to 10 mA which corresponds to an average optical power of 1.5 mW to 2 mW. The VCSEL driver is output terminated and connected to the VCSEL through a 50 Ω microstrip transmission line routed on the top surface of the package.

C. Transimpedance Amplifier

The transimpedance amplifier has the differential symmetric-feedback topology shown in Fig. 8. This differential topology 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 in excess of 12.5 Gb/s when the input parasitic capacitance C_p is less than 250 fF. It has a feedback resistance $(\mathrm{R_F})$ of 314 Ω and open loop gain of 3.9. It receives a single-ended photocurrent of 200 μA from the photodiode and generates a differential $2 \times 50\ \mathrm{mV_{pp}}$ output that is fed to the

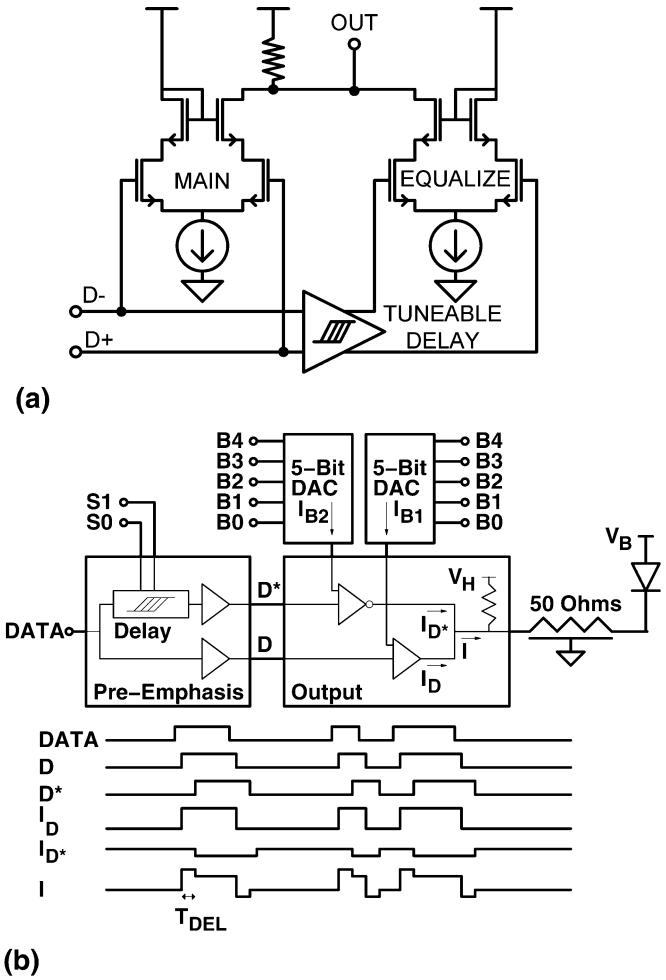

Fig. 7. (a) Schematic of the pre-emphasis VCSEL driver. (b) Pre-emphasis generation circuit.

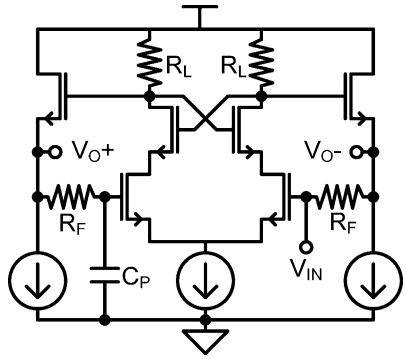

Fig. 8. Schematic of a symmetric transimpedance amplifier.

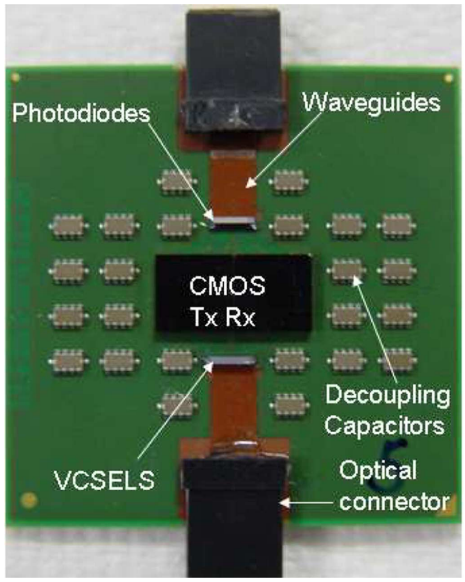

Fig. 9. A fully assembled optical transceiver unit [19].

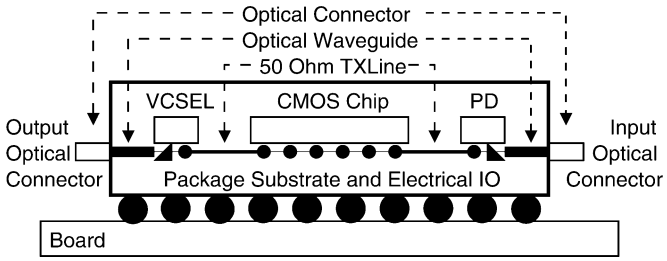

Fig. 10. A side view of optical coupling scheme of VCSELs/photo-detectors to waveguides through a 45 degree mirror.

LA which converts it to a CML level output. The LA consists of a cascade of CML buffers. A current DAC, not shown in the figure, can be used to cancel the DC signal current at the TIA input caused by the zero-state optical power level of the VCSEL. In the packaged transceiver, the combined capacitance of the photodiode, metal pad, C4 bump and ESD could be as high as 500 fF. This limited the maximum data rate that could 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 [16].

D. VCSEL Arrays, Photodetector Arrays, and Polymer Waveguide

Oxide-confined, 850 nm, 10 Gb/s VCSEL arrays were used in the transmitter. The VCSEL array is 3200 μm $\times$ 485 μm with a die thickness of 625 μm and VCSEL aperture spacing of 250 μm. The metal pads are 85 μm $\times$ 85 μm, and the pad spacing is 125 μm. The VCSEL and OLGA substrate are designed with matched layout so they can be directly flip-chip bonded. The VCSEL peak optical output power is more than 3 mW ($\sim$5 dBm).

A 10 Gb/s, 1 $\times$ 12 GaAs PIN detector array receives the optical signal. The detector array is 450 μm $\times$ 3200 μm with a pad size of 75 μm and pad pitch of 125 μm. The detector has 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.

Figs. 9 and 10 show the fully assembled optical transceiver package and a side view of the electrical and optical path between optoelectronics chips and waveguides. The 12-channel polymer waveguide is multi-mode and is 10 mm long and 3 mm wide. It is free-standing and consists of a 45° mirror on one end and a standard MT optical connector on the other. The 45°

TABLE II
HYBRID OPTICAL I/O LINK BUDGET

Avg VCSEL TX Power	**3.0dBm**
VCSEL to MMF Coupling	-1.1dB
MMF to Photodetector Coupling	-1.1dB
Extinction Ratio (7.3dB) Penalty	-1.6dB
Margin	-3.0dB
Link Budget	**-6.8dB**
Required RX Sensitivity	**-3.8dBm**

mirror turns the light 90° and the MT connector couples the input/output optical signals on/off the package. The 45° mirror cut is formed either by microtome or laser ablation and its reverse side is metalized. The waveguides have square apertures with a total height of 100 μm, core dimension of 35 μm $\times$ 35 μm and pitch of 250 μm. The waveguides were fabricated from acrylate using photo-bleach processing [17].

E. Package Architecture

The package architecture allows the integration of low-cost, high-performance optical components with standard microprocessor C4 package technology [18]. The package substrate is 31 mm $\times$ 31 mm with a stack of laminated copper layers separated by a dielectric. A trench was fabricated in the substrate to accommodate the polymer waveguide. All the high-speed electrical lines on the substrate are routed as controlled impedance (50 Ω, single-ended or 90 Ω, differential) microstrip traces. The single-ended microstrip lines are routed on the top surface of the substrate and connect the VCSEL driver and TIA bumps to the VCSEL and photodiode bumps on the package. The optical components are flip-chip bonded with their apertures face down and polymer waveguides couple from below. The total optical loss budget for the end-to-end link includes VCSEL and photodiode coupling loss through the 45° mirrors at either end of the optical link, 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 [1], [19]. Potential improvements to the hybrid package architecture are in development to reduce the optical loss to the 6.8 dB budget described in Table II.

F. Experimental Results

Fig. 11 shows the 10 Gb/s optical measurement results for a channel including a transmitter and a receiver, but no CDR. For the transmitter measurement in Fig. 11(a), external differential electrical PRBS 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. Fig. 11(b) shows the receiver 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 measurement method is described in [19].

Fig. 12 shows the 20 Gb/s electrical and 18 Gb/s optical output eye diagrams of the transmitter with a $2^{15}-1$ PRBS data

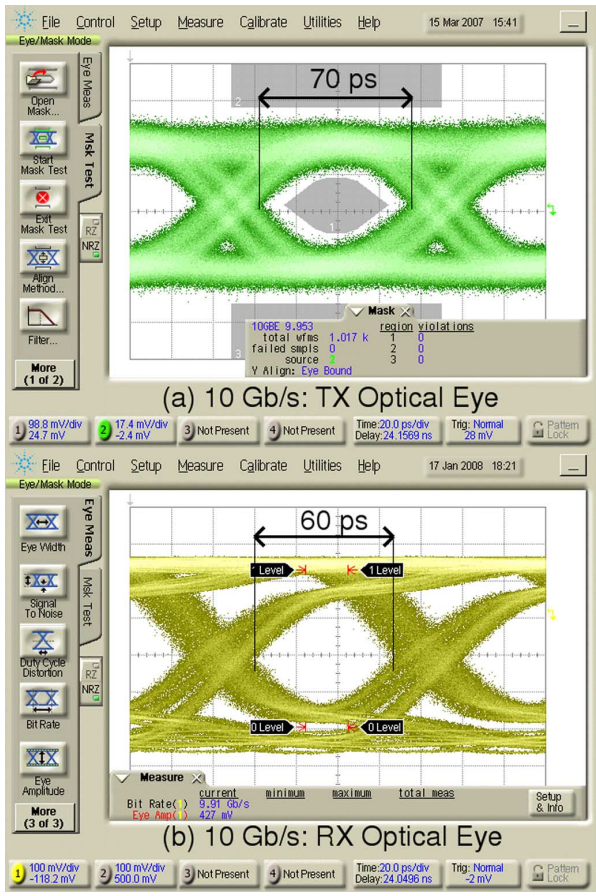

Fig. 11. Optical eye diagrams for 10 Gb/s tested with fully packaged (a) transmitter optical output and (b) receiver optical input.

pattern. The electrical transmit data of Fig. 12(a) was measured by directly probing one of the package contact output pads of the pre-emphasis driver, where a VCSEL array would normally be flip-chip bonded, with a high-speed ground-signal-ground (GSG) coplanar probe. Despite some ISI, a 20 Gb/s electrical eye is observed with an eye opening of 175 mV and 36 ps with peak-to-peak jitter of 16 ps. The optical output data shown in Fig. 12(b) was measured by feeding an 18 Gb/s electrical signal similar to Fig. 12(a) directly to the VCSEL using high-speed coplanar probes and measuring the optical output with a 12 GHz Newfocus photoreceiver through a multimode fiber. Since the signal path now includes combined parasitics from the package and the VCSEL, the driver bias and pre-emphasis currents were readjusted to optimize the 18 Gb/s optical eye. The VCSEL was biased with a 9 mA average drive current and a 2.8 V external bias voltage to provide an average optical power of 2 mW measured at the 12 GHz photoreceiver. The 18 Gb/s optical eye was observed with a vertical and horizontal eye opening of 70 mV and 30 ps, which represents 60% of the total eye. TIA expanded testing used a binary capacitor array integrated in the circuit at the TIA input to test the TIA over a range of input capacitance. Different capacitor combinations were selected by using a focused ion beam to cut the metal and remove capacitance at the input node. As shown in Fig. 13, the

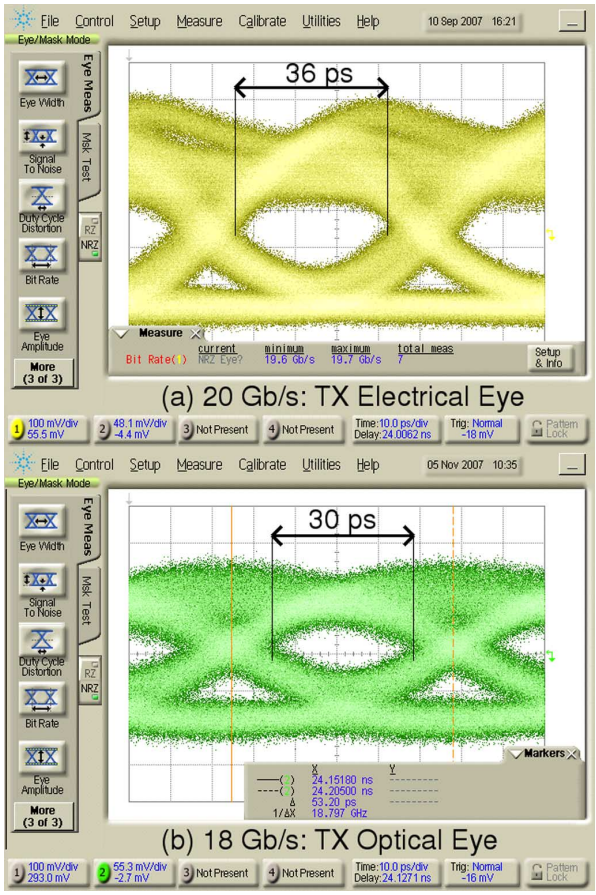

Fig. 12. Transmitter eye diagrams: (a) 20 Gb/s electrical, (b) 18 Gb/s optical.

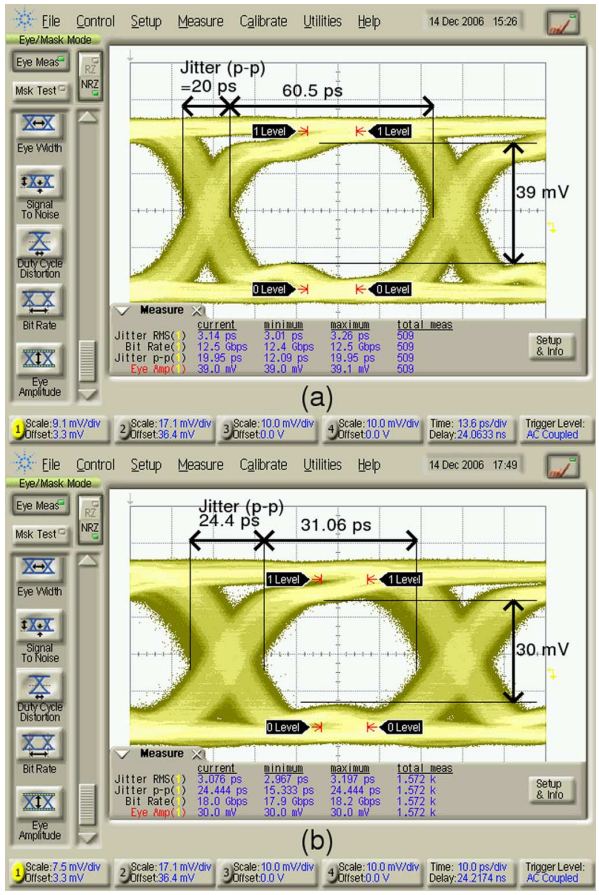

Fig. 13. (a) 12.5 Gb/s with 260 fF input capacitance and 200 μA input current. (b) 18.0 Gb/s with 90 fF input capacitance and 200 μA input current.

TIA operates at 12.5 Gb/s with 260 fF input capacitance and 18 Gb/s with 90 fF input capacitance.

IV. Photonic CMOS Optical I/O

A. Architecture

In the long term, monolithic integration of photonic elements in the CMOS process can enable significant improvements in I/O performance, energy efficiency and cost. The proposed monolithic photonic CMOS process, illustrated in Fig. 14, 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 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 LA.

Monolithic integration of photonics onto the microprocessor promises to reduce power and cost. 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, while certain larger discrete alternatives such as Mach-Zender interferometers 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 processing. Furthermore, the process must allow fabrication of all required optical components on the same die. The demonstrated 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 [20]–[22], and waveguide-embedded metal-semiconductor-metal (MSM) detectors fabricated from poly crystalline germanium on the CMOS process back-end silicon-dioxide dielectric.

B. Fabrication

The first process step in the photonic CMOS process is the deposition of a thick 2 μm layer of SiO_2 on the 300 mm silicon wafers. This layer isolates the optical signal from the substrate

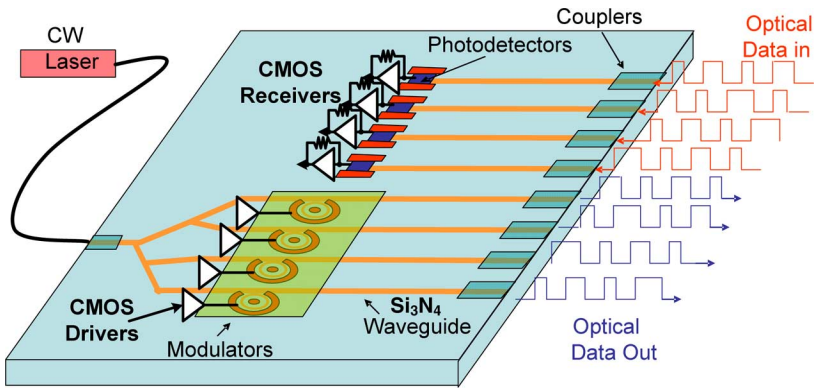

Fig. 14. Photonic optical interconnect architecture.

and simulates the interlayer dielectric (ILD) of the upper metal layers upon which the optical layer would sit in the monolithic photonic CMOS process. Next, a 450 nm layer of silicon nitride is deposited by PECVD and patterned with photolithography and plasma dry etch to form waveguides. This shared waveguide layer is used to build all the waveguides, ring resonators, 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 the modulator regions. The photodetector regions are filled with polycrystalline germanium in a damascene process, the electrodes are formed in a standard damascene process and the modulators are formed by depositing EO polymer cladding over the ring resonators in the modulator regions. The photonic devices in the CMOS process require just four additional photolithography steps, which keeps the cost low.

Fig. 15(a) illustrates a top view of the modulator, waveguide, and detector comprising a complete optical link and Fig. 15(b) shows a cross section SEM image of the same components. It can be seen that 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 the photodetector. In contrast to a monolithic integrated optical transceiver that was previously reported [23], this work presents the first complete link in a back-end compatible process flow. 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. While the poly-crystalline germanium for the detectors used in the full optical link was deposited at higher temperatures, alternative deposition methods at temperatures lower than 450°C have yielded successful results in discrete waveguide-coupled photodetectors.

C. Experimental Results

Waveguide: The waveguide is the foundation for the proposed photonic CMOS technology and is described in [20]. To

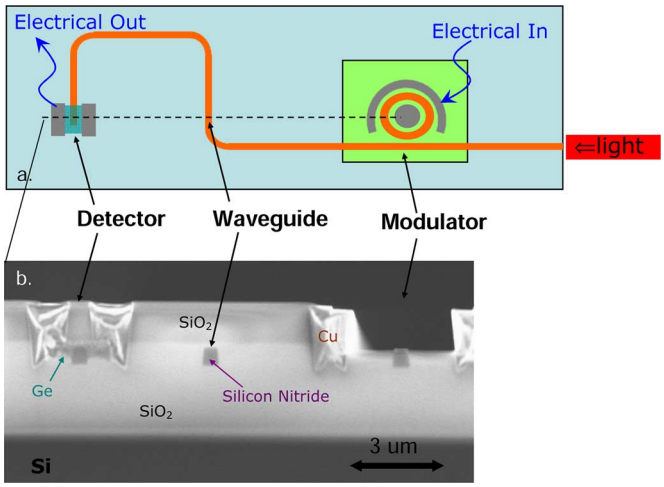

Fig. 15. (a) Schematic of top view of full on-die optical link showing waveguide bus connecting modulator to photo-detector. (b) Cross section SEM image, along the dotted line in (a), showing optical components in one piece of silicon.

form the waveguide, a 450 nm PECVD silicon nitride film is deposited on a 2 μm silicon dioxide under-cladding 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 ring-resonator modulator and the 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 EO polymer from a commercial supplier. The modulator design is optimized for high Q so that a small resonance shift can result in a large modulation depth. The EO polymer is poled before wafer processing completion using an electric field of 100 V/cm at 143°C. The electrodes have a 4.5 μm gap centered on the waveguide ring and the ring has a radius of 28 μm. A SEM image of the modulator is shown in the right hand portion of

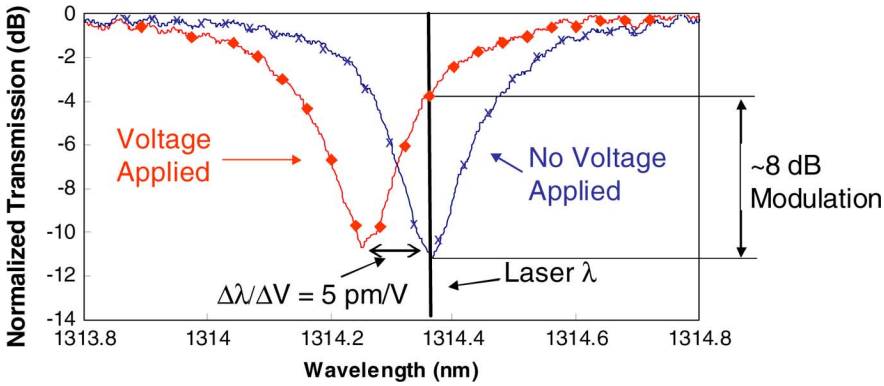

Fig. 16. Resonance spectra obtained with −20 V (♦) and 20 V (x) bias on the EO modulator.

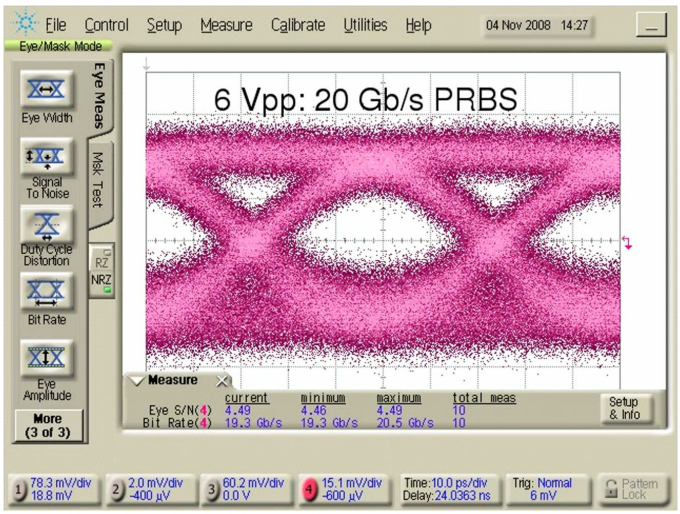

Fig. 17. 20 Gb/s PRBS eye diagram of EO polymer modulator.

Fig. 15(b). Several devices were characterized. The resonance spectrum of a typical modulator under +20 V and −20 V bias is shown in Fig. 16. The resonance shift calculated with a linear fit to the resonance frequencies measured at +20 V and −20 V bias is 5 pm/V. The Q is ∼7000 and the resonance depth is ∼11 dB. The highest measured modulation depth for a 10 GHz clock input at 6 V was 8 dB. A 20 Gb/s PRBS eye diagram for a typical device is shown in Fig. 17.

Photodetector: Unlike a PIN detector, the lateral metal-semiconductor-metal (MSM) detector requires only one lithography step to form the contacts. An evanescently coupled waveguide, shown in the left hand portion of Fig. 15(b), efficiently couples the light into the absorbing active material of the photodectector. The polycrystalline germanium in the detector was deposited by CVD processing at 600°C and further work is in progress to reduce the deposition temperature below ∼450°C. However, fabrication of a photodiode from polycrystalline germanium deposited on ILD is already an important step toward compatibility with a standard CMOS process.

Photodetectors with various electrode designs were fabricated and characterized by measuring current-voltage characteristics, DC responsivity, impulse response, and response to PRBS data. Key results including a PRBS eye diagram at 20 Gb/s were reported in [24] and the following results are from improved devices. The signal-to-noise ratio (SNR) at 20 Gb/s was improved as shown in Fig. 18(a) and a reasonably open eye at 40 Gb/s was demonstrated in Fig. 18(b).

Although the performance is not yet sufficient for a robust 40 Gb/s system, 40 Gb/s measurements are presented here to show the present device capability because the PRBS generator did not allow measurements at rates between 20 Gb/s and 40 Gb/s. The photodiode metallization had an ohmic contact with the polycrystalline germanium film, resulting in a high dark current on the order of 1 mA. To improve the noise performance, an experimental device was fabricated using bandgap engineering to create a Schottky barrier at the metal/germanium contact in order to reduce the dark current. While the increased Schottky barrier also reduces the photo-current for a given bias voltage, the dark current is reduced by ∼90% while the photocurrent is reduced by only ∼50%. This improves the ratio of photo-current to dark-current by ∼80%. Fig. 18(c) shows the 40 Gb/s eye diagram for a photodetector design similar to Fig. 18(b) except for the alternative metallization and dark current of ∼70 μA. Finally, in order to demonstrate that a back-end compatible germanium deposition process can produce high-quality photodetectors at lower temperatures, photodetectors were fabricated using sputter deposition of germanium at 350°C. The 40 Gb/s eye diagram in Fig. 18(d), measured under the same conditions as that in Fig. 18(b), confirms that these photodetectors support similar speeds as those fabricated with 600°C CVD germanium. These results demonstrate that with further attention to device design, bandgap engineering, and materials improvement, a back-end compatible photodetector suitable for high-speed optical links is within reach.

Full Link (Modulator-Waveguide-Photodetector): To demonstrate the monolithic integration of the modulator and photodetector in the same process flow, a modulator and detector were fabricated on the same die and connected with an integrated waveguide. The modulator was driven with a 5 GHz electrical clock signal having an amplitude of 8 V_{pp} and the wavelength tuned to the maximum modulation depth. The modulated signal was detected with a 2 V DC bias on the photodetector. Fig. 19 shows the electrical signal received at

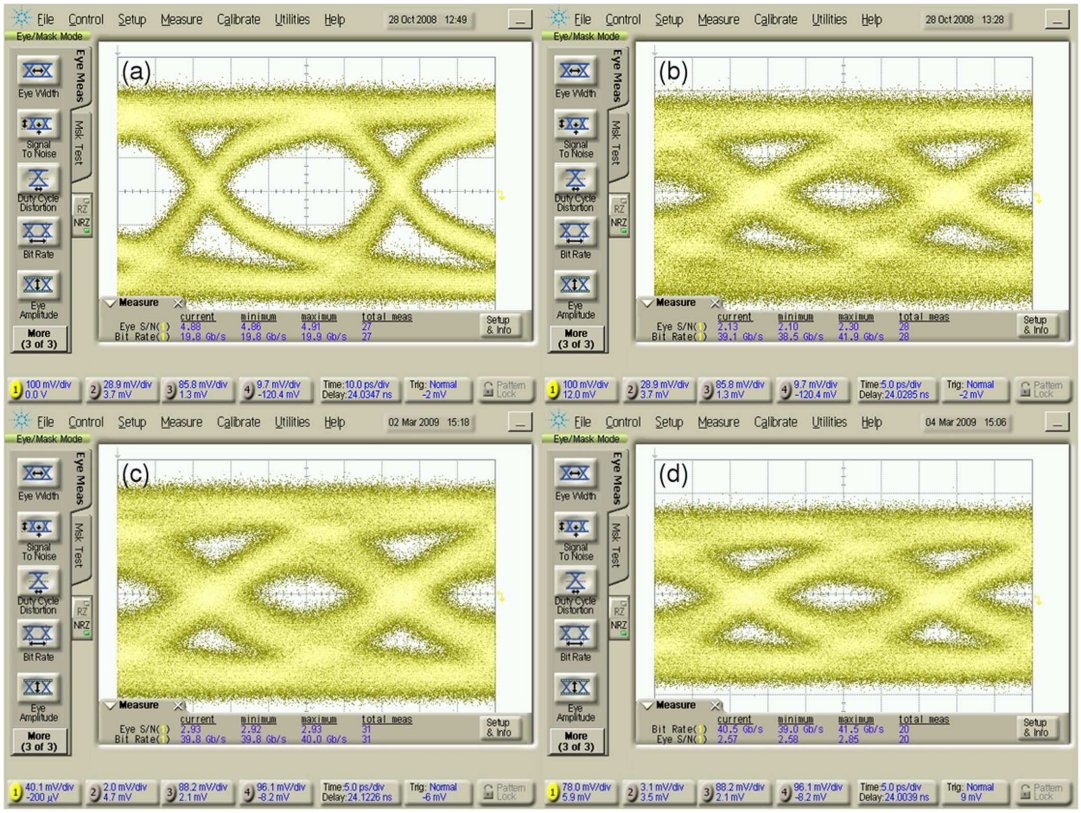

Fig. 18. (a) 20 Gb/s PRBS eye diagram of photodetector with ohmic contacts. (b) Same detector's response to 40 Gb/s PRBS. (c) 40 Gb/s eye diagram of PD with alternate contacts and lower dark current. (d) 40 Gb/s eye diagram of PD with PVD Ge deposited at 350°C.

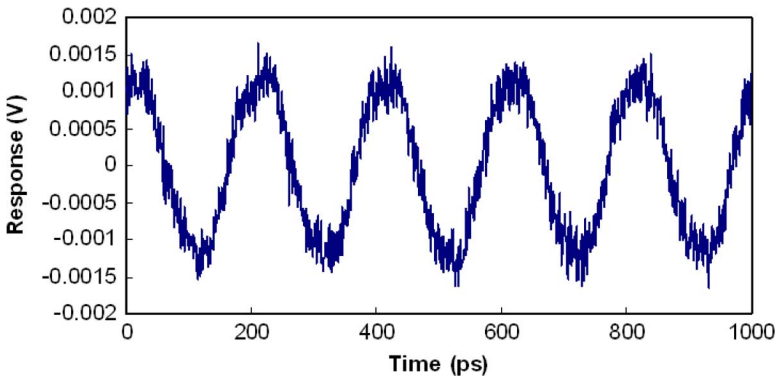

Fig. 19. A 5 GHz signal detected at the Ge photodetector after the CW light is encoded at the modulator.

the photodetector after subtracting the electrical noise due to crosstalk signal coupling across the electrical probes. Although the individual devices support the high data rates described above, strong electrical crosstalk between the probes prevented eye diagram measurements of the full link integrated on one die. In addition, large optical coupling losses prevented testing in a chip-to-chip configuration where the transmitter and receiver are on separate chips which eliminates the probe crosstalk problem. In the future, optical couplers will be fabricated onto the platform to enable this measurement.

V. 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. The link budgets for the hybrid and integrated optical I/O links are shown in Tables II and III, respectively. A feasible best case value for the hybrid optical link budget is -6.8 dB with some packaging improvements [1]. This is dominated by coupling losses from the VCSEL and photodetector to the multi-mode fiber (MMF) and the finite extinction ratio penalty. The integrated optical link budget is nearly 9 dB worse than the hybrid optical link budget due to the off-chip single-mode fiber coupling directly to 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 roughly 13 dB higher sensitivity at the same bandwidth, which results in significant system power savings.

TABLE III
INTEGRATED OPTICAL I/O LINK BUDGET

Source CW Laser Power	**3.0dBm**
Source Laser to SMF Coupling	-2.0dB
SMF to Modulator Coupling	-2.0dB
Modulator Loss	-2.0dB
Modulator to SMF Coupling	-2.0dB
SMF to Photodetector Coupling	-3.0dB
Extinction Ratio (8.0dB) Penalty	-1.4dB
Margin	-3.0dB
Link Budget	**-15.4dB**
Required RX Sensitivity	**-12.4dBm**

Fig. 20 shows circuit simulation-based power efficiency estimates of transmit and receive front-end circuits, excluding the clock timing systems, for these two optical I/O architectures in CMOS technologies spanning from 45 nm to predictive 16 nm. A current-mode VCSEL driver similar to the main driver branch in Fig. 7 and a CMOS inverter-based voltage-mode modulator driver are modeled for the hybrid and integrated optical systems, respectively. In both systems, a TIA similar to Fig. 8 followed by simple differential-pair LA stages make up the optical receiver. The models are constructed with the circuits optimized to provide the minimum bandwidth necessary for a particular data rate, and thus approximate a power optimal solution. The hybrid optical link power efficiency, shown in Fig. 20(a), 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 1 pJ/b will be realized in the future. Assuming a 1310 nm CW laser source with 3 dBm optical power, the integrated optical link power efficiency, shown in Fig. 20(b), 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 allows 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, as seen by the discrete jumps in the power efficiency curves. These projections indicate that photonic CMOS will enable integrated optical interconnect to reach 0.3 pJ/b.

The power-performance analysis of the hybrid optical link is compared with electrical link systems that employ the three electrical channels discussed in Section II. 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. This data rate is dependent on the channel loss characteristics and is 13 Gb/s and 19 Gb/s for the 17 inch backplane and 7 inch 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 in research [25]. The reduced parasitics offered by the integrated photonics with CMOS optical architecture allow this architecture to achieve superior power efficiency compared to the three electrical channels and the hybrid optical architecture over the majority of data rates. 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 [20].

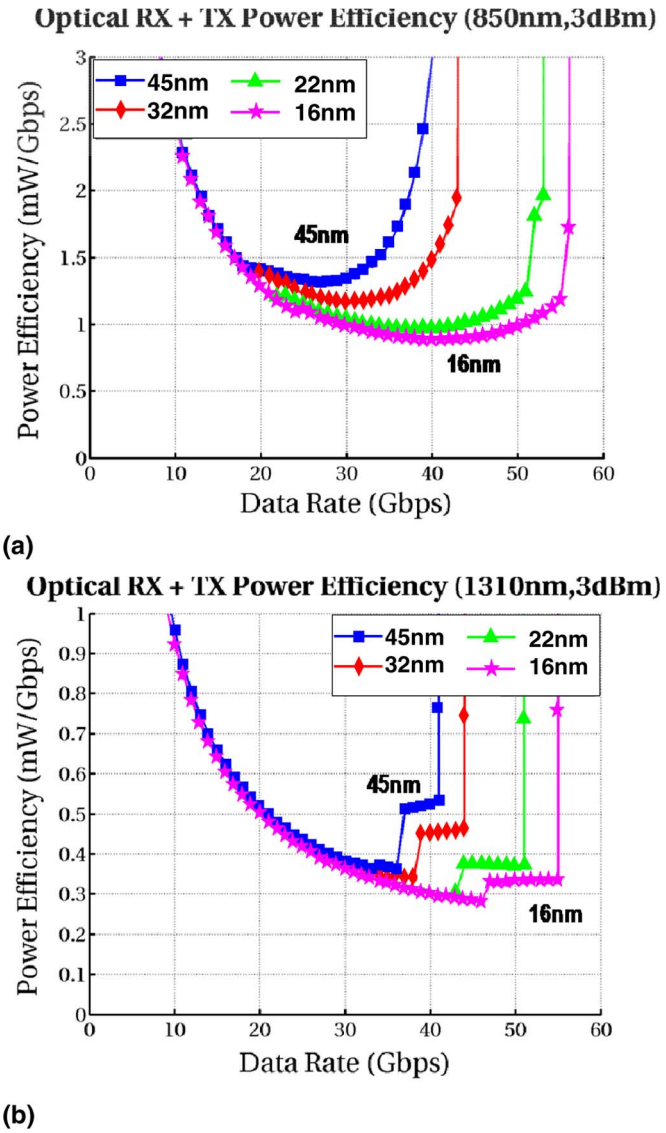

Fig. 20. Circuit simulation-based optical transceiver power efficiency estimates versus data rate for 45 nm to predictively-modeled 16 nm CMOS technology. (a) Hybrid optical I/O architecture. (b) Integrated optical I/O architecture.

VI. FUTURE DIRECTIONS AND CONCLUSIONS

As CMOS scaling continues in the future, larger numbers of CPU cores will be integrated on the microprocessor chip and it

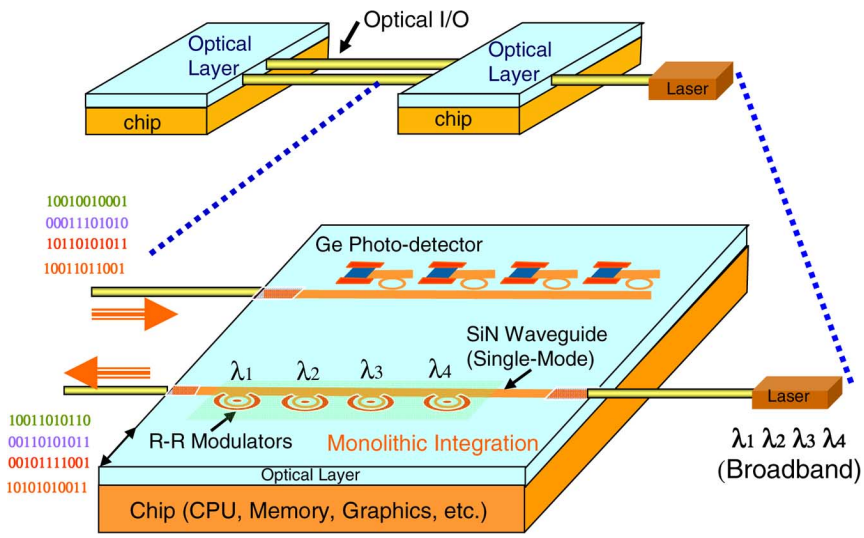

Fig. 21. Photonic CMOS enabled wavelength division multiplexing (WDM) architecture.

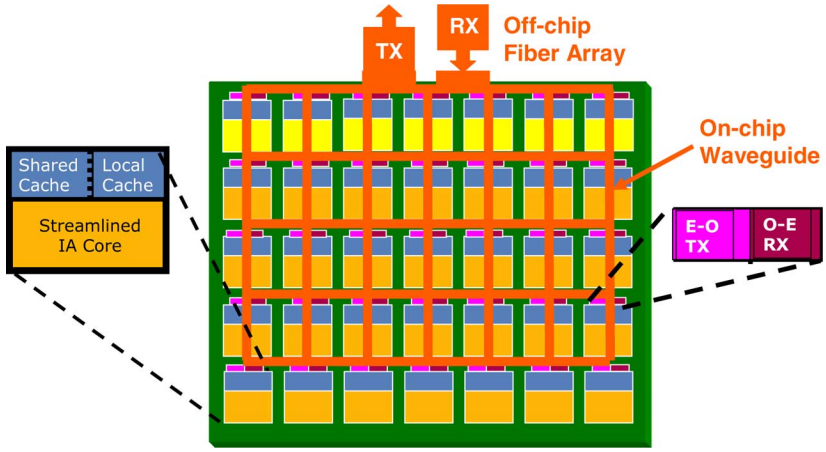

Fig. 22. Optical network on-chip and off-chip using dense WDM.

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. 21. The RR modulator selectively modulates a single wavelength from a multi-wavelength source and eliminates the need for separate optical de-mulitiplexers and multiplexers. At the receiver, passive ring resonator optical filters can de-multiplex 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. 16) the WDM wavelengths can be spaced at less than 1 nm (100 GHz in optical frequency with a reference of 230 THz). Thus, the ring resonator technology provides the means for bandwidth to scale by adding more wavelengths to each waveguide channel. In addition to bandwidth scaling through WDM, optical signals also enable a switchable high-bandwidth optical network for both on-chip and off-chip. Data can be routed to a determined core or memory node by selectively modulating and demodulating data onto a given wavelength. The WDM optical networks can be switched as fast as allowed by the electrical circuits driving the RR modulators. Network on-chip architectures that exploit the low energy-per-bit transmission and high data bandwidth of optical WDM networks will eventually be required to meet the aggregate bandwidth demands of future microprocessors.

The work described in this paper 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. A package-to-package optical I/O prototype achieved 10 Gb/s error free data transmission over a full link. This prototype demonstrated an optical output transmit driver and VCSEL operating at data rates up to 18 Gb/s. The receiver TIA was measured electrically operating up to 18 Gb/s. In the

long term, monolithic integration of optical components will provide TB/s interconnect data rates with the required energy efficiency of less than 1 pJ/bit.

REFERENCES

[1] E. Mohammed, A. Alduino, T. Thomas, H. Braunisch, D. Lu, J. Heck, A. Liu, I. Young, B. Barnett, G. Vandentop, and R. Mooney, "Optical interconnect system integration for ultra-short-reach applications," *Intel Technol. J.*, vol. 8, no. 2, pp. 115–127, May 2004.

[2] M. J. Kobrinsky, B. A. Block, J.-F. Zheng, B. Barnett, E. Mohammed, M. Reshotko, F. Robertson, S. List, I. Young, and K. Cadien, "On-chip optical interconnects," *Intel Technol. J.*, vol. 8, no. 2, pp. 129–142, May 2004.

[3] N. Kurd, J. Douglas, P. Mosalikanti, and R. Kumar, "Next generation Intel® micro-architecture (Nehalem) clocking architecture," in *Symp. VLSI Circuits Dig.*, Jun. 18–20, 2008, pp. 62–63.

[4] B. Casper, J. Jaussi, F. O'Mahony, M. Mansuri, K. Canagasaby, J. Kennedy, E. Yeung, and R. Mooney, "A 20 Gb/s forwarded clock transceiver in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig.*, Feb. 6–9, 2006, pp. 263–272.

[5] B. Casper, J. Jaussi, F. O'Mahony, M. Mansuri, K. Canagasaby, J. Kennedy, E. Yeung, and R. Mooney, "A 20 Gb/s embedded clock transceiver in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig.*, Feb. 6–9, 2006, pp. 1334–1343.

[6] G. Balamurugan, B. Casper, J. Jaussi, M. Mansuri, F. O'Mahony, and J. Kennedy, "Modeling and analysis of high-speed I/O links," *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 237–247, May 2009.

[7] A. Rylyakov and S. Rylov, "A low power 10 Gb/s serial link transmitter in 90-nm CMOS," in *Proc. IEEE Compound Semiconductor Integrated Circuits Symp. (CSICS)*, 30 Oct.–2 Nov. 2005, p. 4.

[8] R. Payne, P. Landman, B. Bhakta, S. Ramaswamy, S. Wu, J. D. Powers, M. U. Erdogan, A.-L. Yee, R. Gu, L. Wu, Y. Xie, B. Parthasarathy, K. Brouse, W. Mohammed, K. Heragu, V. Gupta, L. Dyson, and W. Lee, "A 6.25-Gb/s binary transceiver in 0.13- μm CMOS for serial data transmission across high loss legacy backplane channels," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2646–2657, Dec. 2005.

[9] J. F. Bulzacchelli, M. Meghelli, S. V. Rylov, W. Rhee, A. V. Rylyakov, H. A. Ainspan, B. D. Parker, M. P. Beakes, A. Chung, T. J. Beukema, P. K. Pepeljugoski, L. Shan, Y. H. Kwark, S. Gowda, and D. J. Friedman, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2885–2900, Dec. 2006.

[10] L. Chen, X. Zhang, and F. Spagna, "A scalable 3.6-to-5.2 mW 5-to-10 Gb/s 4-tap DFE in 32 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig.*, Feb. 6–9, 2006.

[11] K. Kuhn, M. Agostinelli, S. Ahmed, S. Chambers, S. Cea, S. Christensen, P. Fischer, J. Gong, C. Kardas, T. Letson, L. Henning, A. Murthy, H. Muthali, B. Obradovic, P. Packan, S. W. Pae, I. Post, S. Putna, K. Raol, A. Roskowski, R. Soman, T. Thomas, P. Vandervoorn, M. Weiss, and I. Young, "A 90 nm communication technology featuring SiGe HBT transistors, RF CMOS, precision R-L-C RF elements and 1 μm^2 6-T SRAM cell," in *Int. Electron Devices Meeting (IEDM) Dig.*, Dec. 2002, pp. 73–76.

[12] H. Muthali, T. Thomas, and I. Young, "A CMOS 10-Gb/s SONET transceiver," *IEEE J. Solid-State Circuits*, vol. 39, no. 7, pp. 1026–1033, Jul. 2004.

[13] D. Kucharski, Y. Kwark, D. Kuchta, D. Guckenberger, K. Kornegay, M. Tan, C.-K. Lin, and A. Tandon, "A 20 Gb/s VCSEL driver with pre-emphasis and regulated output impedance in 0.13 μm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig.*, Feb. 2005, pp. 222–223.

[14] C. Kromer, C. Sialm, C. Berger, T. Morf, M. L. Schmatz, F. Ellinger, D. Erni, G. L. Bona, and H. Jackel, "A 100 mW 4 × 10 Gb/s transceiver in 80 nm CMOS for high-density optical interconnects," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2667–2679, Dec. 2005.

[15] J. J. Morikuni, P. V. Mena, 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.

[16] A. Kern, A. Chandrakasan, and I. Young, "18 Gb/s Optical IO: VCSEL driver and TIA in 90 nm CMOS," in *Symp. VLSI Circuits Dig.*, Jun. 14–16, 2007, pp. 276–277.

[17] B. L. Both, J. E. Marchegiano, C. T. Chang, R. J. Furmanak, D. M. Graham, and R. G. Wagner, "Polyguide\texttrademark polymeric technology for optical interconnect circuits and components," in *Proc. SPIE (Photonics West: Optoelectronic Interconnects and Packaging IV)*, Feb. 1997, vol. 3005, pp. 238–251.

[18] 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 (Photonics West: Photonics Packaging and Integration IV)*, Jun. 2004, vol. 5358, pp. 60–70.

[19] E. Mohammed, J. Liao, A. Kern, D. Lu, H. Braunisch, T. Thomas, S. Hyvonen, S. Palermo, and I. Young, "Optical hybrid package with an 8-channel 18 GT/s CMOS transceiver for chip-to-chip optical interconnect," in *Proc. SPIE (Photonics West: Parallel Optical Links)*, Feb. 2008, vol. 6899.

[20] B. A. Block, T. R. Younkin, P. S. Davids, M. R. Reshotko, P. Chang, B. M. Polishak, S. Huang, J. Luo, and A. K. Y. Jen, "Electro-optic polymer cladding ring resonator modulators," *Optics Express*, vol. 16, no. 22, pp. 18326–18333, Oct. 2008.

[21] A. K. Y. Jen, J. Luo, T.-D. Kim, B. Chen, S.-H. Jang, J.-W. Kang, N. M. Tucker, S. Hau, Y. Tian, J.-W. Ka, M. Haller, Y. Liao, B. Robinson, and L. Dalton, "Exceptional electro-optic properties through molecular design and controlled self-assembly," in *Proc. SPIE (Linear and Nonlinear Optics of Organic Materials V)*, Aug. 2005, vol. 5935.

[22] D. Chen, H. R. Fetterman, A. Chen, W. H. Steier, L. R. Dalton, W. Wang, and Y. Shi, "Demonstration of 110 GHz electro-optic polymer modulation," *Appl. Phys. Lett.*, vol. 70, no. 25, pp. 3335–3337, Jun. 1997.

[23] A. Narasimha, B. Analui, E. Balmater, A. Clark, T. Gal, D. Guckenberger, S. Gutierrez, M. Harrison, R. Ingram, R. Koumans, D. Kucharski, K. Leap, Y. Liang, A. Mekis, S. Mirsaidi, M. Peterson, T. Pham, T. Pinguet, D. Rines, V. Sadagopan, T. J. Sleboda, D. Song;, Y. Wang, B. Welch, J. Witzens, S. Abdalla, S. Gloeckner, and P. De Dobbelaere, "A 40-Gb/s QSFP optoelectronic transceiver in a 0.13 μm CMOS silicon-on-insulator technology," in *Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Feb. 24–28, 2008, pp. 1–3.

[24] M. R. Reshotko, B. A. Block, B. Jin, and P. Chang, "Waveguide coupled Ge-on-oxide photodetectors for integrated optical links," in *IEEE/LEOS Intl. Conf. Group IV Photonics*, Sep. 17–19, 2008, pp. 182–184.

[25] T. Anan, N. Suzuki, K. Yashiki, K. Fukatsu, H. Hatakeyama, T. Akagawa, K. Tokutome, and M. Tsuji, "High-speed 1.1- μm-range InGaAs VCSELs," in *Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Feb. 24–28, 2008, pp. 1–3.

[26] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Holzwarth, M. Popovic, H. Li, H. Smith, J. Hoyt, F. Kartner, R. Ram, V. Stojanovic, and K. Asanovic, "Building manycore processor-to-DRAM networks with monolithic silicon photonics.," in *Proc. IEEE Symp. High Performance Interconnects*, Aug. 26–28, 2008.

Ian. A. Young (M'78–SM'96–F'99) is a Senior Fellow and Director of Advanced Circuits and Technology Integration in the Technology and Manufacturing Group at Intel Corporation. He is responsible for research and development of mixed-signal circuits for microprocessor and SOC products along with process technology development.

Dr. Young joined Intel in 1983. Starting with the development of circuits for a 1 Mb DRAM, and the world's first 1 μm 64 K SRAM, he then led the design of three generations of SRAM products and manufacturing test vehicles, and developed the original Phase Locked Loop (PLL) based clocking circuit in a microprocessor while working on the 50 MHz Intel 486™ 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. Dr. Young has developed a number of optimization metrics for process technology development, including the transistor performance metric that provided a link between processor performance and basic transistor parameters, as well as back-end metal interconnect architecture. He is currently directing the research and development of analog mixed-signal and RF circuits for high-speed serial I/O in advanced logic processes and wireless transceivers in advanced SOC processes along with research for chip-to-chip optical I/O technology.

Born in Melbourne, Australia, 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. Prior to Intel, Dr. Young worked on analog/digital integrated circuits for Telecommunications products at Mostek Corporation (United Technologies), as well as an independent design consultant.

Dr. Young was a member of the Symposium on VLSI Circuits Technical Program Committee from 1991 to 1996, serving as the Program Committee Chairman in 1995/1996, and the Symposium Chairman in 1997/1998. He was a member of the ISSCC Technical Program Committee from 1992 to 2005, serving as the Digital Subcommittee Chair from 1997 to 2003, the Technical Program Committee Vice-chair in 2004 and Chair in 2005. He was Guest Editor for the April 1997, April 1996 and December 1994 issues of the JSSC. He has served as an elected member of the SSCS Adcom from 2006 to 2009.

Dr. Young is a Fellow of the IEEE. He holds 45 patents in integrated circuits and has authored or co-authored over 40 technical papers.

Edris Mohammed is a research engineer at Intel Corporation and is 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 integration and passive optical and electrical devices. Dr. Mohammed received a B.Sc. in Physics from Addis Ababa University in 1987, M.Sc. in Physics from Florida A&M University in 1995, M.Sc. in Electrical Engineering (Optoelectronics) and a Ph.D. in applied physics both from Georgia Institute of Technology in 2000. Since 2005 Dr. Mohammed has been serving as an associate editor of optical interconnects for SPIE Optical Engineering Journal.

Jason T. S. Liao (M'09) received the M.S. degree in communication engineering from National Chiao Tung University, Taiwan, in 1996 and the Ph.D. degree in electrical engineering from University of California at San Diego in 2004.

In 2005, he joined Intel Corporation, Hillsboro, OR, 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 the A.B. and B.E. degrees from Dartmouth College in 2002 and the 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 in Hillsboro, OR, 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'06) received the B.S. and M.S. degree in electrical engineering from Texas A&M University, College Station, TX in 1997 and 1999, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA in 2007.

From 1999 to 2000, he was with Texas Instruments, Dallas, TX, 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, Hillsboro, OR, 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. Dr. Palermo is a member of IEEE and 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 1989–1992 he worked for IBM Corporation in E. Fishkill 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 photodetector, 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 also from the Hebrew 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 both discrete optical devices, including waveguides, modulators, and photodetectors, and optical links.

Peter L. D. Chang (M'92) received the B.S. degree in physics from the National Taiwan University in 1977 and the 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 Department and the ECE Department of UC 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. From 1996 to present, 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 CPU.