

Ultralow Power 80 Gb/s Arrayed CMOS Silicon Photonic Transceivers for WDM Optical Links

Xuezhe Zheng, *Senior Member, IEEE*, Frankie Y. Liu, Jon Lexau, *Member, IEEE*, Dinesh Patil, *Member, IEEE*, Guoliang Li, Ying Luo, Hiren D. Thacker, *Member, IEEE*, Ivan Shubin, Jin Yao, Kannan Raj, *Senior Member, IEEE*, Ron Ho, *Senior Member, IEEE*, John E. Cunningham, and Ashok V. Krishnamoorthy, *Fellow, IEEE, Fellow, OSA*

Abstract—Silicon photonic interconnects offer a promising solution to meeting the ever growing demand for more efficient I/O bandwidth density. We report an ultralow power 80 Gb/s arrayed silicon photonic transceiver for dense, large bandwidth inter/intrachip interconnects. Low parasitic microsoldier-based hybrid bonding enables close integration of silicon photonic array devices optimized on a 130 nm silicon-on-insulator CMOS platform with CMOS very large scale integration circuits optimized on a 40 nm silicon CMOS platform to achieve unprecedented energy efficiency. The hybrid CMOS transceiver consists of eight 10 Gb/s channels with a total consumed power below 6 mW/channel. The eight-channel wavelength division multiplexing transmitter array using cascaded tunable ring modulators demonstrated better than 100 fJ/bit energy efficiency for 10 Gb/s operation excluding the laser power and tuning power, while the eight-channel receiver array using broadband Ge p-i-n waveguide detectors show sensitivity of better than -15 dBm for a bit error rate of 10^{-12} at a data rate of 10 Gb/s with energy efficiency of better than 500 fJ/bit.

Index Terms—Optical communications, optical interconnections, optical receivers, optical transmitters, silicon photonics, transceiver array.

I. INTRODUCTION

CONTINUING to scale the performance of computer systems within reasonable power constraints requires increasing total system parallelism. The number of processors per chip and the chips per board must all steadily increase in order to grow the aggregate computational power, thereby raising requirements for aggregate interconnect bandwidth across the entire system. At the processor level, multicore architectures




Fig. 1. I/O bandwidth scaling of general purpose processors and commercial FPGA chips.

have enabled continued performance growth, even as clock rates have stopped scaling. Consequently, as shown in Fig. 1, designers have been pushing the I/O bandwidth of general purpose multicore processors and commercial field-programmable gate arrays (FPGAs) to follow Moore’s Law scaling, doubling every two years to maintain balanced systems [1]–[8]. Given that such chips already require 2–5 Tb/s, it is extremely challenging to maintain such scaling using traditional electrical solutions. Transmission line driver technologies have achieved impressive results; state-of-the-art commercial electrical links recently demonstrated 28 Gb/s per channel, 8.82 pJ/bit energy efficiency, and 35 Gbps/mm² bandwidth per unit area [9]. The many on-going research and academic efforts further pushing these aggregate bandwidth and energy targets have found the design space to be tightly constrained; however, significant improvements are nonetheless necessary to keep up with integration scaling trends. Projecting forward using the metrics above, 20 Tb/s of I/O bandwidth would consume 90 W and 285 mm². Higher data rates will reduce the number of links, but will need more equalization and noise cancellation due to channel loss, thus increasing power and area. It is not clear that electrical links can be denser and more power efficient even with more advanced CMOS technology node. In addition, error encoding may also be needed for low bit error rate (BER), resulting in increased link latency. These, and other considerations, thus motivate a need for an interconnect technology with low power, low latency, and high bandwidth density.

Silicon photonic interconnects offer a promising solution to help meet these demands [10]–[12]. Many different inter/intrachip architectures have been proposed to exploit these interconnects [13]–[15]; in particular, by using optical proximity communication [16]–[18], photonic networking can allow a collection of discrete chips to function as a logically contiguous piece

Manuscript received July 09, 2011; revised October 17, 2011, November 29, 2011; accepted December 01, 2011. Date of publication December 14, 2011; date of current version February 01, 2012. This work was supported in part by the U.S. Defense Advanced Research Project Agency under Grant HR0011-08-09-0001.

X. Zheng, G. Li, Y. Luo, H. D. Thacker, I. Shubin, J. Yao, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy are with Oracle Labs, San Diego, CA 92121 USA (e-mail: xuezhe.zheng@oracle.com; glenn.li@oracle.com; ying.luo@oracle.com; hiren.thacker@oracle.com; ivan.shubin@oracle.com; jin.yao@oracle.com; raj.kannan@oracle.com; john.cunningham@oracle.com; ashok.krishnamoorthy@oracle.com).

F. Y. Liu, J. Lexau, and R. Ho are with Oracle Labs, Redwood City, CA 94002 USA (e-mail: Frankie.liu@oracle.com; jon.lexau@oracle.com; ron.ho@oracle.com).

D. Patil was with Oracle Labs, Menlo Park, CA 94025 USA. He is now with Rambus, Sunnyvale, CA 94089 USA (e-mail: patil7ue@gmail.com).

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/JLT.2011.2179287



Fig. 2. Conceptual view of WDM interconnects using hybrid-integrated silicon photonic transceivers.

of silicon. By seamlessly scaling silicon beyond its physical reticle limits, such a “macrochip” offers unprecedented computational density, energy efficiency, bisection bandwidth, and reduced message latency [19].

A low-power wavelength division multiplexing (WDM) photonic link is key to achieving energy-efficient and high-bandwidth interconnects for macrochip and other system architectures. Such a link must be built around arrayed low-power WDM transceivers. Using high-density microsolder bumping technology [20], we first demonstrated single-channel low-power silicon photonic transmitters and receivers at 5 Gb/s with energy costs as low as 320 fJ/bit [21] and 690 fJ/bit [22], respectively. Later, using the same hybrid integration technique to bond improved silicon photonic devices to improved CMOS circuits, we showed single-channel 10 G transmitters and receivers with a record energy efficiency of 135 fJ/bit and 380 fJ/bit, respectively [23].

In this study, we report on the development of an ultra-efficient WDM ready 80 Gb/s CMOS photonic arrayed transceiver. We show optimized silicon photonics devices arrays built in a Luxtera/Freescale 130 nm silicon-on-insulator (SOI) CMOS platform, and hybrid-integrated with high-speed very large scale integration (VLSI) circuits optimized for low power using the TSMC 40 nm Silicon CMOS process.

II. 80 GB/S HYBRID INTEGRATED CMOS PHOTONIC ARRAY WDM TRANSCEIVER

Previous single channel demonstrations [20], [21], [23] have proven that hybrid integration is a valid approach for building efficient silicon photonic transceivers. It allows photonics and electronics to be independently optimized on different technology nodes for best-of-breed performance from each. Densely integrated multiple transmitters and receivers in an arrayed WDM transceiver can take advantage of the same technology. Using such hybrid integration, a conceptual view of WDM interconnects is shown in Fig. 2. Silicon photonic opto-electronic (OE) device chips are hybrid integrated with VLSI driver circuits as photonic bridges, which are then interconnected optically via a WDM optical routing layer using optical proximity coupling. The WDM multiplexer (Mux) and demultiplexer (DeMux) devices can be integrated together with the OE devices on the photonic bridges, or with the routing waveguides on the WDM routing chip.

To achieve an energy-efficient WDM silicon photonic transceiver array, we need high-speed OE device arrays with low parasitic loads and low-voltage operation, low-power, high-speed CMOS VLSI arrayed driver circuits, and high-yield low-parasitic integration [20].

A. Arrayed Cascaded Ring Modulators With Tuning

Silicon ring modulators offer compelling size and power characteristics. High-speed voltage modulation across a re-

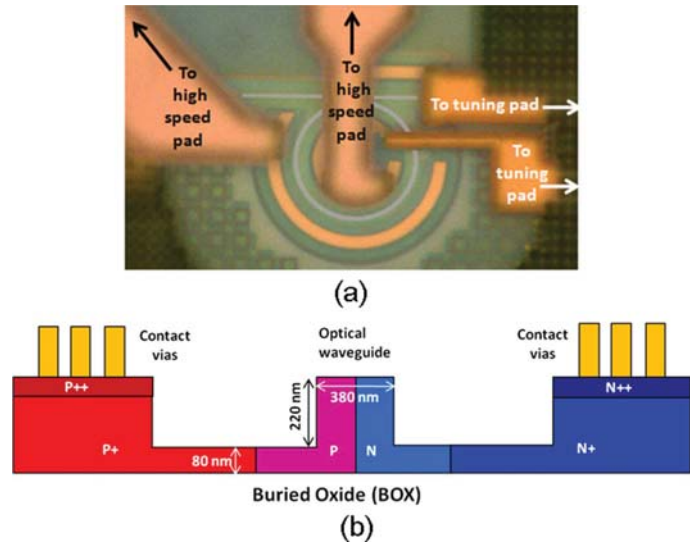


Fig. 3. (a) Ring modulator. (b) Cross section diagram of the ring waveguide high-speed section.

verse-biased p-n junction embedded in the ring waveguide can be very efficient [21]. The highly wavelength-selective nature of the ring modulators, often considered a disadvantage, can be used to make compact WDM array transmitters by cascading multiple ring modulators along the same bus waveguide without significant interference. We designed eight slightly differently sized rings, each with a radius of approximately $7.5 \mu\text{m}$, to accommodate eight wavelength channels at 1.6 nm channel spacing within one free spectral range (FSR) of 12.8 nm. To overcome the static resonance shifts due to fabrication tolerances, and dynamic drift from changes in ambient temperature, doped silicon-resistor-based thermal tuners are integrated into each modulator.

Fig. 3(a) shows a photograph of one fabricated ring modulator fabricated in Freescale’s 130 nm HIP7_SOI process. 67% of the ring is doped as a p-n junction [$1.5 \times 10^{18} \text{ cm}^{-3}$ doping, shown in Fig. 3(b)] for high-speed modulation, while the upper-right 25% is n-type doped as an Si resistor for thermal tuning. There are $2 \mu\text{m}$ wide isolation gaps, where the ring waveguide is undoped, between the p-n diode and the thermal-resistor sections. Multiple metal layers are used for the layout of electrodes. The ring waveguide is 380 nm wide, with 220 nm etch depth and 80 nm thick Si slab. We made the ring waveguide wider than the bus waveguide (300 nm wide) for smaller bending loss and better coupling phase matching. The doping densities in the p-n junction are determined by the tradeoff between the modulation efficiency and the speed.

The high-speed behavior of the ring modulator was examined using a circuit model extracted by curve-fitting the measured S11 data for both phase and amplitude. The circuit model and the extracted circuit values at 1 V are shown in Fig. 4. C_P denotes the capacitance between the electrodes through the top dielectrics, R_s and C_J model the current path through the reverse-biased p-n junction, R_{Si} and C_{OX} model the current path through the BOX and the Si handler.




Fig. 4. Small-signal circuit model for the ring modulator with circuit values at 1 V reverse bias.

The modulator EO frequency response was tested using a microwave network analyzer, and a 3 dB bandwidth of >10 GHz was obtained at 1 V. The modulation bandwidth of the device is mostly subject to the photon lifetime limit, which is 12 GHz based on the measured quality factor.

B. Ge *p-i-n* Waveguide Photo Detectors

Ge *p-i-n* waveguide photodiodes (PDs), fabricated in Freescale's HIP7_SOI 130 nm CMOS technology with Luxtera's Ge-enabled optoelectronic process [24], have demonstrated exceptional performance. Such evanescent coupled Ge detectors are fabricated by selective epi growth of Ge on a $1.5 \mu\text{m}$ wide and $15 \mu\text{m}$ long multimode silicon waveguide, tapered from a $0.3 \mu\text{m}$ wide single-mode waveguide. These devices are very compact in area and have electrical bandwidths exceeding 12 GHz, room temperature dark current less than $3 \mu\text{A}$ at 0.5 V reverse bias, measured responsivity of $> 0.8 \text{ A/W}$ for wavelength less than 1535 nm (which drops to about 0.7 A/W at 1540 nm, and 0.6 A/W at 1550 nm due to Ge band edge), and diode capacitance of only 10 fF [23]. Because a Ge waveguide PD has very broad optical bandwidth, eight identical PDs were designed into an array WDM receiver.

C. Photonic EO Device Array Layout

The photonic device arrays were optimized using a specialized process design kit developed at Luxtera. The floorplan of the chip is shown in Fig. 5(a). Eight ring modulators are cascaded along a shared bus waveguide, and each ring modulates a wavelength channel out of an eight-wavelength comb with 1.6 nm channel spacing within the 12.8 nm FSR. These rings are slightly different in size with nominal radius of approximately $7.5 \mu\text{m}$. Eight Ge waveguide detectors, however, are arranged as a 1×8 array with separate waveguide because they are broadband devices. The photonic device array is laid out in a column with $250 \mu\text{m}$ cell pitch.

Grating couplers are used as the optical I/O interface, with a pitch of $250 \mu\text{m}$ to support optical testing using fiber arrays. To form a complete WDM link using such arrayed modulators and PDs, a WDM comb source or a wavelength Mux is needed to multiplex single wavelength sources together to drive the cascaded modulators, and a wavelength DeMux is needed to separate the wavelength channels before the PD array. We assume that the WDM Mux/DeMux devices [25], [26] can be integrated on the WDM routing chip as one possible option. Using the grating coupler interface, complete WDM links can be established between the transceiver bridges through a WDM routing

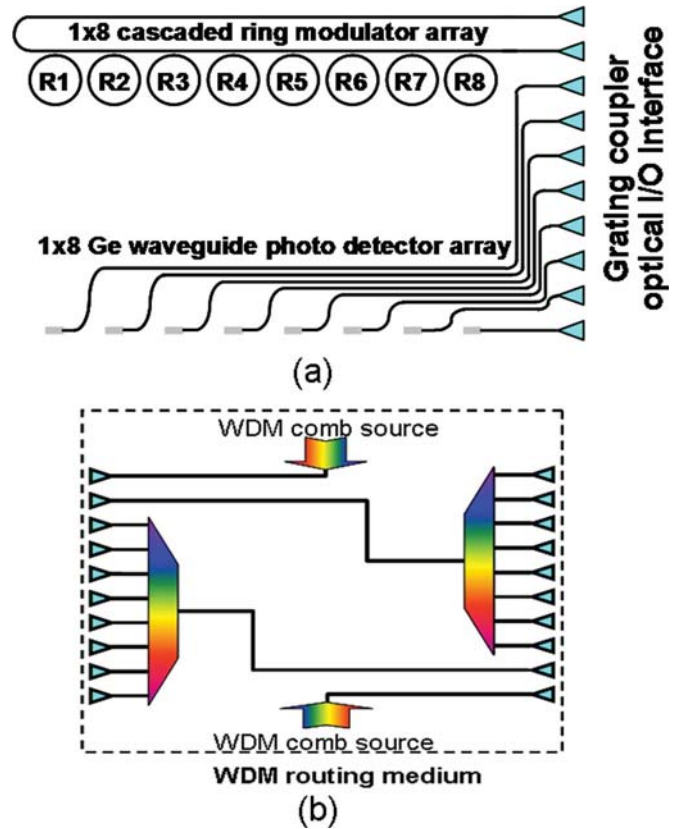


Fig. 5. (a) and (b) Arrayed photonic device layout of an 80 Gb/s silicon photonic transceivers for WDM links using cascaded ring modulators, broadband Ge waveguide photodetectors, and grating couplers for optical I/O.

chip equipped with matching optical I/O configuration, as the schematics shown in Fig. 5(b). One can also integrate the WDM Mux/DeMux devices directly with the OE devices on the bridge to simplify the routing chip as an alternative option.

D. Low-Power VLSI Circuits

Low-power, high-speed VLSI circuits were developed in TSMC's 40 nm bulk CMOS node. As the circuit model indicates, the reverse-biased depletion ring modulator appears largely as an RC load to a CMOS driver. We use cascode drivers to modulate 10 Gb/s data on the cascaded rings with up to 2 V swing, as the circuits diagram shown in Fig. 6.

Pulse generators on the cascode devices activate them within a timing window designed to minimize V_{gs} and V_{gd} overstress during output transitions [27]. Each driver consumes a silicon area of only $120 \mu\text{m}^2$. Sixteen copies of these compact modulator driver are arranged in a column with a pitch of $125 \mu\text{m}$. For testing purposes, the modulator drivers can take data from either a common on-chip pseudorandom bit sequence (PRBS) generator or from an off-chip pin. The datapath, however, can selectively invert and/or delay that data stream at each individual modulator. This allows data bits at different experiments to be independent and thus enables more realistic device array testing. In order to simplify the on-chip digital circuitry, data are moved across the chip in two parallel "single-data-rate" 5 Gb/s streams and converted at the modulator driver into one "double-data-rate" 10 Gb/s stream.




Fig. 6. Cascode modulator driver circuit.




Fig. 7. Diagram of low-power receiver circuit with retiming and power supply regulator.

Tuning helps to align the ring modulators with their corresponding wavelength channels. To drive the silicon resistor heater built in each ring modulator, current-source-based tuning circuits with up to 6 mW output power to a 400 Ω resistor, and with 12-bit resolution, are also included in every modulator driver cell.

The low-power receiver [28], [29] is based on a transimpedance amplifier (TIA) followed by a strobed sense-amplifier (SA). The SA, more efficient than traditional limiting amplifiers, restores the received signal to full CMOS voltage levels. As the circuit diagram depicts in Fig. 7, the receiver includes a clock-phase recovery circuit, based on a delay-locked loop (DLL), to properly clock the SA. The receiver also includes a facility to adjust the input current so that the TIA output swings symmetrically around the SA switching threshold, so that SA can correctly determine whether the data were a “0” or a “1.” Because SAs suffer from intrinsic mismatch and comparison inaccuracy, the receiver implements offset compensation.

Setting the DLL for appropriate clocking of the SA, setting the input level to center the SA input around the voltage threshold, and setting two distinct offset compensation circuits requires a total of four digital-to-analog converters (DACs) in the receiver. Digital finite-state machines drive and set these four DACs during periodic recalibration of the optical link, once every millisecond or so.

The three-stage TIA is designed for a gain of 4 k Ω at a bandwidth of 7 GHz, assuming an input capacitive load of 60 fF that include both PD and hybrid bonding pads. This enables operation at -15 dBm average optical input power with a worst case extinction ratio (ER) of 3 dB and a PD responsivity of




Fig. 8. Arrayed transceiver VLSI chip.

0.7 A/W, and targeting a signal-to-noise ratio sufficient for a BER of 10^{-12} .

Each low-power receiver consumes silicon area of less than 8000 μm^2 excluding the bonding pads. Bonding pads are smaller than the circuits and may be placed directly above the CMOS circuits. Sixteen copies of such receiver circuits are arranged in a column with a pitch of 125 μm . To minimize the supply noise and crosstalk from neighboring channels in an array, a linear regulator on the power supply of each receiver offers low-frequency noise rejection better than 30 dB [29].

The VLSI test chip was fabricated in TSMC’s 40 nm bulk CMOS technology. Its organization, in Fig. 8, shows four transmit (TX) columns of 16 modulator drivers each, as well as three columns arrayed receivers instantiated [29].

For our arrayed transmitters and receivers, capacitive and inductive crosstalk between array sites is small. Traditional electrical links, which employ large soldered structures to carry signal currents, can suffer from significant crosstalk, particularly from high-current transmitters placed in close proximity to sensitive low-swing input receivers. In such circumstances, capacitive coupling between soldered pins, plus capacitive and inductive coupling between transmission lines, can generate noise that overwhelms the sensitivity limits of the receiver circuits. In our case, however, the transmitters drive small capacitive loads of only a few tens of fF, over small (25 μm) hybrid bonding pads, so the circuits are small and do not couple strongly to each other. Also, the waveguides carry no net electrical current and generate neither inductive loops nor injected capacitive noise. We saw no effects of running many other transceivers while measuring one selected channel.

However, noise may still couple between receivers, or from transmitters to receivers, by means of the power supply. For a TIA aiming at -15 dBm sensitivity, a power supply transient can easily swamp out the received signal: in fact, simulations show the TIA has a nearly 18 dB gain from power supply to output. To compensate, our test chip employs power supply regulators on the TIA, on re-timing DLLs, and on DACs used for receiver calibration. The regulators, based on a differential amplifier, show in excess of 18 dB of broadband noise rejection (simulated) across all process corners, and hence just cancel the TIA’s supply noise gain. In order to further avoid coupling noise between the TIA and the timing circuits, we use separate power supply regulators for these different blocks.




Fig. 9. Hybrid integrated WDM transceiver on a test PCB.

E. Hybrid Integration

Hybrid integration [30] allows photonics and electronics to be independently optimized on different technology nodes for best-of-breed performance from each. Employing a special under-bump-metallization solution and microsolder bumps fabricated through lithography, metal deposition, and liftoff, we developed a low-parasitic microsolder bump technology compatible with 40 nm bulk CMOS using extreme low- κ interlayer dielectrics [31]. 25 μm bonding pads are used for achieving the targeted less than 20 fF hybrid integration capacitance.

OE device chips and VLSI chips are flip-chip bonded using thermal compression to collapse the microsolder bumps. The hybrid chip assembly is then die-attached and wire bonded to a carrier substrate to provide power, clocking and digital data I/Os needed for a complete transceiver. Fig. 9 shows a photo of the complete arrayed silicon photonic transceiver, with the photonic device and VLSI hybrid chip assembly wire-bonded to a test PCB.

III. WDM TRANSCEIVER PERFORMANCE

A. Arrayed WDM Transmitter Performance

A fiber array enables characterization of the transceiver performance through the grating coupler I/O interface. Since the on-chip silicon waveguide only supports TE polarization, the source laser light was polarization controlled and polarization maintained fibers were used to match the waveguide polarization. As discussed earlier, deep optical trench etching was used to enable ultracompact ring design, which unfortunately made the fiber grating coupler performance worse. A ~ 5 dB coupling loss was measured through the fiber/grating coupler interface. Using optimized grating trench depth [32], and introducing backside mirror [33], such coupling loss is expected to be less than 1 dB. For thermal stability, a room temperature passive copper heat sink was attached to the photonic device chip. Fig. 10(a) shows the “through” spectra of the cascaded ring modulators with -5 dBm optical power in the input fiber. As expected, due to manufacturing tolerances, the ring resonances are not evenly distributed on the designed 1.6 nm channel spacing. However, it is encouraging to see that the resonances of the cascaded rings (designed using slightly different radius for each ring) in close physical proximity are separated in order. Using

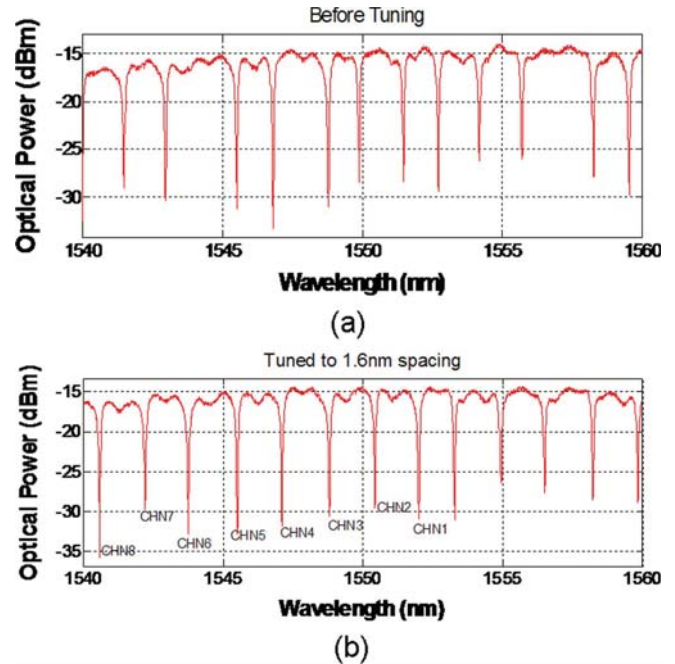


Fig. 10. Through spectra of the cascaded ring modulator array (a) before and (b) after tuning.

the built-in thermal tuners, we were able to position the ring modulator resonances to a 1.6 nm channel spacing as intended [see Fig. 10(b)], overcoming any resonance shifts due to fabrication variances [see Fig. 10(a)]. The resonance shift needed were 0.543, 0.577, 0.043, 0.346, 0.005, 0.8, 0.77, and 0.588 nm for CHN1 to CHN8 as labeled in Fig. 10(b), corresponding to individual ring turning power of 2.74, 2.92, 0.22, 1.75, 0.03, 4.09, 3.89, and 2.97 mW, respectively.

The built-in silicon resistor thermal tuner was driven by the 12 bit digital current source. The tuning capability of such a tuner is shown in Fig. 11. The resistance of the doped resistor was measured about 600 Ω , slightly larger than the design target of 400 Ω , which limits the maximum power that can be delivered to the heater to be about 4 mW. As indicated by Fig. 11, maximum tuning of only about 0.8 nm can be achieved with 4 mW of tuning power. Using thermal engineering by trench isolation and undercut [34], or substrate removal by back side etch pit [35], up to 20 times tuning improvement have been demonstrated. With improved tuning efficiency, we will be able to position a ring resonance anywhere within one FSR with less than 4 mW of tuning power.

With the modulators tuned to the correct channel spacing, high-speed transmission performance was tested by injecting WDM continuous wave wavelengths into the input bus waveguide with about 1 mW of in-waveguide optical power, and driving the modulators with a 10 Gb/s PRBS stream generated on-chip. Limited by WDM laser source availability, a single tunable laser source was used to test different wavelength channels sequentially with electrical modulation turned on for all rings. An erbium-doped fiber amplifier (EDFA) was used to amplify the modulated optical signal before an Agilent reference receiver. The 10 Gb/s modulated optical “eyes” of the eight WDM transmitter channels are presented in Fig. 12 and show dynamic




Fig. 11. Ring modulator resonance shift in wavelength driven by the digital current source.

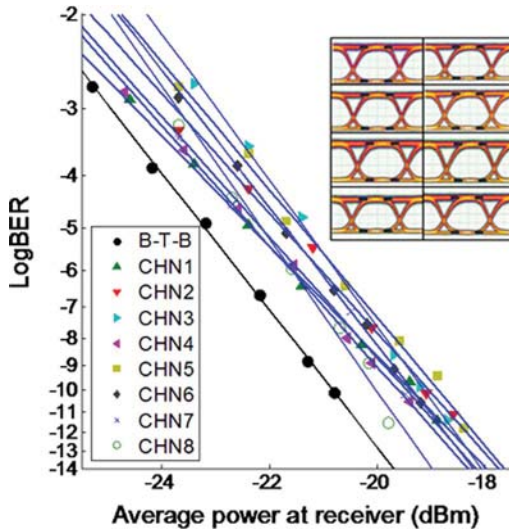


Fig. 12. Measured received power penalty using a reference InGaAsP-based receiver at 10 Gb/s for the cascaded ring transmitters with the inset picture showing the optical eye diagrams of eight WDM channels.

ER ranging from 5.6 to 8 dB for the same “ON” state loss of -4 dB (excluding the fiber to waveguide coupling loss). The plots in Fig. 12 indicate a received power penalty of 1–2 dB for the WDM array transmitter when compared to an Agilent lightwave transmitter with 12 dB ER, mostly attributable to the ring modulator array’s imperfect dynamic ERs.

With an array spacing of $250 \mu\text{m}$, we expect little optical crosstalk. Electrical crosstalk impairment was measured by measuring the “eye” closure penalty on a channel with and without all the neighboring channels turned on as aggressors. The channels in the middle of the array suffer from the most electrical crosstalk. The measurement result of one such channel, CHN5, was shown in Fig. 13. The measured received power penalty with crosstalk is about 0.5 dB better than without crosstalk at BER of 10^{-12} . This is mostly due to the error in average power measurements. The EDFA output power fluctuates about ± 0.2 dB. In addition, we had to switch the fiber connection between an optical receiver and a power meter to measure the average power. The fiber connection brings in additional uncertainty about ± 0.1 dB. We, therefore, believe that the crosstalk is negligible.

The power consumption of the modulators was measured directly from the driver supply voltage and current. The modulation power was measured to be 0.94, 0.71, 0.70, 0.70, 0.92, 0.68,

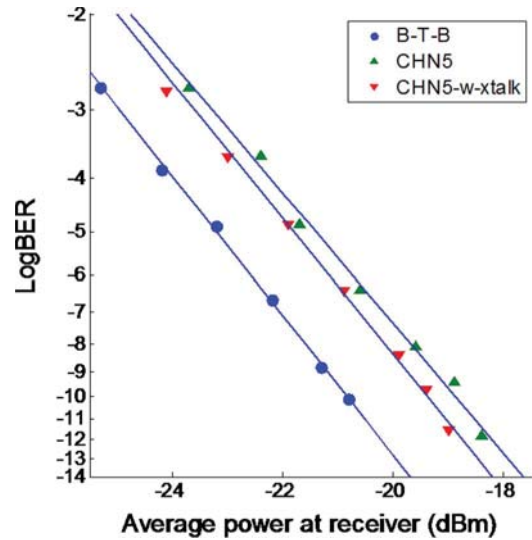


Fig. 13. Worst case channel crosstalk induced “eye” closure penalty of the arrayed WDM transmitters.

0.66, and 0.91 mW, respectively, for the 8 transmitters. A record high efficiency of less than 100 fJ/bit per channel is achieved.

B. Arrayed Receiver Performance

The arrayed receiver used multiple individual clocks: each receiver employed its own DLL to phase-align its SA strobe to the incoming data stream. The array was tested using a modulated 10 Gb/s PRBS ($2^{31} - 1$) input optical signal from an Agilent lightwave transmitter. An on-chip PRBS error checker measured the receiver performance. The receiver circuit allows a test interface to externally set both the SA’s decision threshold voltage and the DLL’s phase delay. By sweeping this voltage and time, and evaluating the BER for each (voltage, time) pair, we obtain the receiver eye diagram. The receiver also offers a test mode with direct measurement of the average input photocurrent. Fig. 14 shows the eight measured receiver eyes with $24 \mu\text{A}$ of average input photocurrent at wavelength of 1535 nm. The vertical axis of each “eye” represents the sense amplifier threshold voltage range of 160 mV, while the horizontal axis represents the timing of one UI or 100 ps. The white opening in the middle of the eyes represents a BER under 10^{-12} . Because the error counter used for the on-chip error checker has only 8 bits, and a 20 min error counting time period was used to make sure credible measurement of BER under 10^{-12} , BER above 2×10^{-11} would overflow the error counter and processed as the black region in each “eye.” The abrupt changes around the edge of the “eye” opening are also due to such effect. Assuming a PD responsivity of 0.8 A/W, all the receivers achieved a sensitivity of better than -15 dBm for 10 Gb/s at BER of 10^{-12} without optical crosstalk from WDM channels.

With no WDM DeMux integrated, channels spacing of $250 \mu\text{m}$, and waveguides laid out with big enough spacing to minimize optical crosstalk, we believe that electrical crosstalk dominates. We measure the crosstalk impairment from neighboring channels for this study. We reduce the average input optical power of one channel in the middle of the array to about -17 dBm at the detector such that the receiver “eye” has




Fig. 14. 10 Gb/s eye diagrams of eight receiver channels measured one channel at a time using the on-chip PRBS error checker.




Fig. 15. Measured receiver “eye” at 10 Gb/s with minimum opening at BER of 10^{-12} (a) before and (b) after crosstalk aggressor channel was turned ON.

minimum opening at BER of 10^{-12} , as shown in Fig. 15(a). We then turn ON the crosstalk aggressor channel with about -15 dBm average on the detector to see its impact on the measured channel “eye” opening. As the measured result shown in Fig. 15(b), the “eye” opening is actually slightly better than the “eye” with the neighboring channel turned OFF. We believe that such small difference is most likely due to the power fluctuation during the hours long measurement, and the crosstalk in the arrayed receivers is negligible.

Similar to the transmitter measurement, the receiver power was also measured directly from its supply voltages and currents. The measured power consumptions are 4.52, 2.72, 2.6,

2.92, 3.2, 2.84, 2.72, and 4.04 mW, respectively, for eight channels. At less than 500 fJ/bit per channel, it is the most energy efficient receiver array with retiming ever reported [22], [23].

IV. CONCLUSION

Energy efficiency and bandwidth density of interconnects are becoming the most critical issues for computing system scaling. While pure electrical solutions are running out of tricks in the tool box, and struggling to meet the need of future system scaling, WDM optical link solutions based on silicon photonics offer a promising alternative with potentially orders of magnitude improvement in both bandwidth density and energy efficiency.

Low-parasitic hybrid integration again proves to be a viable approach to enable high energy efficiency for dense silicon photonic interconnects, with which an arrayed silicon photonic transceiver was successfully developed for use in 80 Gb/s WDM silicon photonic links by employing a cascaded ring modulator array with thermal tuning and a Ge waveguide PD array, optimized and cointegrated using the Luxtera/Freescale 130 nm SOI CMOS process. A full accounting of the power of a silicon photonic link, particularly when being considered as a replacement for a digital electronic link, must include transmitter and receiver functionality, link retiming requirements, additional supply regulation needs, as well as tuning requirements and the required laser power. In the transceivers described in this paper, we have incorporated the first three functions. Low-power, high-speed VLSI circuits were designed in the TSMC 40 nm bulk CMOS process. Excluding tuning and laser power, the 80 Gb/s arrayed transceiver consumes a record-low power of less than 6 mW per channel. In other work, we have shown that the tuning power for cascaded ring devices similar to those described here, with appropriate back-end CMOS processing [35] can be reduced to less than 0.5mW per microring in fully processed CMOS wafers with 193nm lithography [36]. To complete the power accounting, a final consideration is to develop suitable tunable optical sources with high electrical-to-optical conversion efficiency.

The transceiver circuits reported here, together with input–output grating couplers for power and I/O, provide an I/O density of 700 Gbps/sq.mm. Through the optical interface using grating couplers, the highly efficient hybrid WDM transceiver serves as a promising key element enabling various intra/interchip networking applications with compelling bandwidth, density, and energy efficiency advantages.

ACKNOWLEDGMENT

The authors would like to thank Dr. J. Shah of Defense Advanced Research Project Agency’s Microsystems Technology Office for his inspiration and support of this program. The authors would also like to thank Dr. T. Pinguet and the Luxtera team for contribution to this work.

REFERENCES

- [1] J. L. Shin, D. Huang, B. Petrick, C. Hwang, K. W. Tam, A. Smith, H. Pham, H. Li, T. Johnson, F. Schumacher, A. S. Leon, and A. Strong, “A 40 nm 16-core 128-thread SPARC SoC processor,” *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 131–144, Jan. 2011.

- [2] M. Shah, J. Barreh, J. Brooks, R. Golla, G. Grohoski, H. Gura, R. Hetherington, P. Jordan, M. Luttrell, C. Olson, B. Saha, D. Sheahan, L. Spracklen, and A. Wynn, "UltraSPARC T2: A highly-threaded, power efficient, SPARC SoC," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2007, pp. 22–25.
- [3] U. G. Nawathe, M. Hassan, K. C. Yen, A. Kumar, A. Ramachandran, and D. Greenhill, "Implementation of an 8-core, 64-thread, power efficient SPARC server on a chip," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 6–20, Jan. 2007.
- [4] A. S. Leon and D. Sheahan, "The ultraSPARC T1: A power-efficient high-throughput 32-thread SPARC processor," in *Proc. IEEE Asian Solid-State Circuits Conf.*, 2006, pp. 27–30.
- [5] M. Gschwind, "Optimizing data sharing and address translation for the cell BE heterogeneous chip multiprocessor," in *Proc. Int. Conf. Comput. Design*, 2008, pp. 478–485.
- [6] J. Stuecheli, "System performance scaling of IBM POWER6™ based servers," *Hotchips* 19, 2007 [Online]. Available: http://www.hotchips.org/archives/hc19/2_Mon/HC19.01/HC19.01.02.pdf
- [7] J. M. Tendler, "POWER7 processors: The beat goes on," 2011 [Online]. Available: <http://www.ibm.com/developerworks/wikis/download/attachments/104533501/POWER7+-+The+Beat+Goes+On.pdf>
- [8] "Xilinx FPGA device families support documentation." [Online]. Available: <http://www.xilinx.com/support/documentation/index.htm>
- [9] D. Mansur, "Stratix V with 28 Gbps transceivers in 28 nm," in *Hotchips* 22, 2010 [Online]. Available: <http://www.hotchips.org/conference-archives/hot-chips-22>, Available:
- [10] R. Soref, "The past, present, and future of silicon photonics," *IEEE J. Sel. Topics Quantum Electron*, vol. 12, no. 6, pp. 1678–1687, Nov./Dec. 2006.
- [11] M. Lipson, "Guiding, modulating and emitting light on silicon—Challenges and opportunities," *J. Lightw. Technol.*, vol. 23, no. 12, pp. 4222–4239, Dec. 2005.
- [12] S. Assefa, W. M. Green, A. Rylyakov, C. Schow, F. Horst, and Y. Vlasov, "CMOS integrated nanophotonics: Enabling technology for exascale computing systems," presented at the Opt. Fiber Commun. Conf., Los Angeles, CA, 2011, Paper OMM6.
- [13] A. Shacham, K. Bergman, and L. P. Carloni, "On the design of a photonic network-on-chip," in *Proc. First Int. Symp. Netw.-on-Chip*, 2007, pp. 53–64.
- [14] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2008, pp. 153–164.
- [15] 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 network with monolithic silicon photonics," in *Proc. IEEE Symp. High Performance Interconnects*, Aug. 2008, pp. 21–30.
- [16] X. Zheng, J. E. Cunningham, I. Shubin, J. Simons, T. Sze, D. Cohen, A. V. Krishnamoorthy, M. Asghari, D. Feng, H. Lei, D. Zheng, H. Liang, C. C. Kong, and J. Luff, "Optical proximity communication using reflective mirrors," *Opt. Exp.*, vol. 16, no. 19, pp. 15052–15058, 2008.
- [17] A. V. Krishnamoorthy, J. E. Cunningham, X. Zheng, I. Shubin, J. Simons, D. Feng, H. Liang, C. Kung, and M. Asghari, "Optical proximity communication with passively aligned silicon photonic chips," *IEEE J. Quantum Electron*, vol. 45, no. 4, pp. 409–414, 2009.
- [18] J. Yao, X. Zheng, G. Li, Y. Luo, H. Thacker, A. Mekis, T. Pinguet, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "Grating coupler based optical proximity coupling for scalable computing systems," in *Proc. Optoelectron. Interconnects Compon. Integration XI*, 2011, pp. 794405-1–794405-7.
- [19] A. V. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J. E. Cunningham, "Computer systems based on silicon photonic interconnects," *Proc. IEEE*, vol. 97, no. 7, pp. 1337–1361, Jul. 2009.
- [20] J. E. Cunningham, A. V. Krishnamoorthy, I. Shubin, J. G. Mitchell, and X. Zheng, R. Ho and R. Drost, Eds., "Aligning chips face to face for dense capacitive communication," in *Coupled Data Techniques*. Berlin, Germany: Springer, 2010, pp. 163–168.
- [21] G. Li, X. Zheng, J. Lexau, Y. Luo, H. Thacker, P. Dong, S. Liao, D. Feng, M. Asghari, J. Yao, Y. Shi, P. Amberg, N. Pinckney, K. Raj, R. Ho, J. E. Cunningham, and A. V. Krishnamoorthy, "Ultralow-power, high-performance Si photonic transmitter," presented at the Opt. Fiber Commun. Conf., San Diego, CA, 2010, Paper OMI2.
- [22] X. Zheng, F. Liu, D. Patil, H. Thacker, Y. Luo, T. Pinguet, A. Mekis, J. Yao, G. Li, J. Shi, K. Raj, J. Lexau, R. Ho, J. E. Cunningham, and A. V. Krishnamoorthy, "An sub-picojoule-per-bit CMOS photonic receiver for densely integrated systems," *Opt. Exp.*, vol. 18, no. 1, pp. 204–211, 2010.
- [23] X. Zheng, D. Patil, J. Lexau, F. Liu, G. Li, H. Thacker, Y. Luo, I. Shubin, J. Li, J. Yao, P. Dong, D. Feng, M. Asghari, T. Pinguet, A. Mekis, P. Amberg, M. Dayringer, J. Gainsley, H. F. Moghadam, E. Alon, K. Raj, R. Ho, J. Cunningham, and A. V. Krishnamoorthy, "Ultra-efficient 10 Gb/s hybrid integrated silicon photonic transmitter and receiver," *Opt. Exp.*, vol. 19, no. 6, pp. 5172–5186, 2011.
- [24] G. Masini, S. Sahni, G. Capellini, J. Witzensand, and C. Gunn, "High-speed, near infrared optical receivers based on Ge waveguide photodetectors integrated in a CMOS process," in *Advances in Optical Technologies*. Chicago, IL: Hindawi, 2008.
- [25] B. E. Little, J. S. Foresi, G. Steinmeyer, E. R. Thoen, S. T. Chu, H. A. Haus, E. P. Ippen, L. C. Kimerling, and W. Greene, "Ultra-compact Si-SiO₂ microring resonator optical channel dropping filters," *IEEE Photon. Technol. Lett.*, vol. 10, no. 4, pp. 549–551, Apr. 1998.
- [26] X. Zheng, I. Shubin, G. Li, T. Pinguet, A. Mekis, J. Yao, J. Costa, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "A tunable 1 × 4 silicon CMOS photonic wavelength multiplexer/demultiplexer for dense optical interconnects," *Opt. Exp.*, vol. 18, no. 5, pp. 5151–5160, 2010.
- [27] S. Palermo and M. Horowitz, "High-speed transmitters in 90 nm CMOS for high-density optical interconnects," in *Proc. 32nd Eur. Solid-State Circuits Conf.*, Sep. 2006, pp. 508–511.
- [28] R. Ho, J. Lexau, F. Liu, D. Patil, R. Hopkins, E. Alon, N. Pinckney, P. Amberg, X. Zheng, J. E. Cunningham, and A. V. Krishnamoorthy, "Circuits for silicon photonics on a 'macrochip,'" in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2009, pp. 17–20.
- [29] F. Liu, D. Patil, J. Lexau, P. Amberg, M. Dayringer, J. Gainsley, H. F. Moghadam, X. Zheng, J. E. Cunningham, A. V. Krishnamoorthy, E. Alon, and R. Ho, "10 Gbps, 535 fJ/b optical transceiver circuits in 40 nm CMOS," in *Proc. VLSI Symp.*, 2011, vol. C27-3, pp. 290–291.
- [30] P. Castelein, J. M. Debono, M. Fendler, C. Louis, F. Marion, L. Mathieu, and M. Volpert, "Ultra fine pitch hybridization of large imaging detectors," in *Proc. IEEE Nuclear Sci. Symp. Conf. Rec.*, 2003, vol. 5, pp. 3518–3521.
- [31] H. D. Thacker, Y. Luo, J. Shi, I. Shubin, J. Lexau, X. Zheng, G. Li, J. Yao, J. Costa, T. Pinguet, A. Mekis, P. Dong, S. Liao, D. Feng, M. Asghari, R. Ho, K. Raj, J. G. Mitchell, A. V. Krishnamoorthy, and J. E. Cunningham, "Hybrid-integrated silicon photonic bridge chips for ultralow energy inter-chip communications flip-chip integrated silicon photonic bridge chips for sub-picojoule per bit optical links," in *Proc. IEEE Electron. Compon. Technol. Conf.*, 2010, pp. 240–246.
- [32] D. Taillaert, P. Bienstman, and R. Baets, "Compact efficient broadband grating coupler for silicon-on-insulator waveguides," *Opt. Lett.*, vol. 29, no. 23, pp. 2749–2751, 2004.
- [33] F. Van Laere, G. Roelkens, M. Ayre, J. Schrauwen, D. Taillaert, D. Van Thourhout, T. F. Krauss, and R. Baets, "Compact and highly efficient grating couplers between optical fiber and nanophotonic waveguides," *J. Lightw. Technol.*, vol. 25, no. 1, pp. 151–156, Jan. 2007.
- [34] P. Dong, W. Qian, H. Liang, R. Shafiha, X. Wang, D. Feng, G. Li, J. E. Cunningham, A. V. Krishnamoorthy, and M. Asghari, "1 × 4 reconfigurable demultiplexing filter based on free-standing silicon racetrack resonators," *Opt. Exp.*, vol. 18, pp. 24504–24509, 2010.
- [35] J. E. Cunningham, I. Shubin, X. Zheng, T. Pinguet, A. Mekis, Y. Luo, H. Thacker, G. Li, J. Yao, K. Raj, and A. V. Krishnamoorthy, "Highly-efficient thermally-tuned resonant optical filters," *Opt. Exp.*, vol. 18, no. 18, pp. 19055–19063, 2010.
- [36] A. V. Krishnamoorthy, X. Zheng, G. Li, J. Yao, T. Pinguet, A. Mekis, H. Thacker, I. Shubin, Y. Luo, K. Raj, and J. E. Cunningham, "Exploiting CMOS manufacturing to reduce tuning requirements for resonant optical devices," *IEEE Photon. J.*, vol. 3, no. 3, pp. 567–579, Jun. 2011.

Xuezhe Zheng (M'03–SM'03) received the B.S., M.S., and Ph.D. degrees in optical instruments from Tsinghua University, Beijing, China, in 1993 and 1997, respectively.

He is currently a Consulting Hardware Engineer with Oracle Labs, San Diego, CA. He has extensive experiences in photonic switching and optical cross-connect, fiber optic components, dense wavelength division multiplexing optical networks, and optical interconnections. He is the author or coauthor of more than 100 papers in technical journals and conferences. He holds 14 U.S. patents. His current research interests include wavelength division multiplexing silicon photonics for advanced inter/intrachip interconnects.

Dr. Zheng is a recipient of the Science and Technology Development Award from the National Education Committee of China.

Frankie Y. Liu was born in Sao Paulo, Brazil. He received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, with a thesis in the area of semiconductor device physics.

He is currently with the VLSI group in Oracle Labs, Redwood City, CA, where he is involved in mixed circuit design for high-speed electrical and optical links.

Jon Lexau (M'06) received the B.S. degree in computer and systems engineering from Rensselaer Polytechnic Institute, Troy, NY, in 1989, and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 1994.

From 1989 to 1993, he was with the Amdahl Corporation designing high-performance mainframe CPUs. Since 1994, he has been with Oracle Labs, Redwood City, CA, working in the VLSI Research Group. Recently, he has been working on capacitively coupled proximity communication and electrical transceivers for optical interconnect.

Dinesh Patil (S'03–M'07) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Mumbai, India, and the Ph.D. degree from Stanford University, Stanford, CA, in 2008.

From 2007 to 2010, he was with the VLSI Research Group, Sun Laboratories, Oracle, Menlo Park, CA, where he was involved in the receiver design for interchip proximity and silicon photonics communication links. He is currently a Senior Circuit Design Engineer at Rambus, Inc., Sunnyvale, CA. His research interests include transceivers for chip-to-chip and on-chip electrical and optical interconnects, circuit optimization for energy-efficient digital circuits, and graphics/main memory subsystem design.

Guoliang Li received the B.S. and M.S. degrees in physics from Peking University, Beijing, China, and the Ph.D. degree in electrical and computer engineering from the University of California, San Diego.

He is currently a Principal Hardware Engineer at Oracle Labs, San Diego, CA. Since 2001, he has worked in multiple companies, including Bell Labs, Luxtera, Intel, and Sun Microsystems, where he has designed and developed various high-speed optical devices (in III–V and Si) and transceiver modules for applications in telecommunication, Ethernet, and optical interconnects.

Ying Luo received the Ph.D. degree in electrical engineering from the University of New Mexico, Albuquerque, in 2006.

She is currently a Senior Hardware Engineer at Oracle Labs, San Diego, CA, working on the DARPA Ultraperformance Nanophotonic Intrachip Communication project. Prior to this, she was a device engineer with Santur Corp, a leading tunable distributed feedback (DFB) laser company in Fremont, CA, where she was involved in the research on the development of 100 Gb/s directly modulated DFB laser arrays. She has extensive experience in semiconductor device fabrication, material characterization, and multichip integration and packaging. Her current research interests include novel wavelength division multiplexing silicon photonics devices, and SiP-based inter/intrachip communication.

Hiren D. Thacker (S'97–M'02) received the B.S. degree in electrical engineering and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology in 2000, 2002, and 2006, respectively.

He is currently a Principal Hardware Engineer in the Photonics, Interconnects, and Packaging Group at Oracle Labs, San Diego, CA, where he is involved in 2-D and 3-D integration and packaging of ICs deploying advanced electrical and silicon-photonics interconnects. From 2006 to 2008, he was an R&D Microelectromechanical Systems (MEMS) Design Engineer at NanoNexus Inc., where he was involved in research on stressed-metal wafer-level MEMS I/O interconnects for wafer probing. He has co-authored more than 50 articles in refereed journals and conference proceedings and has several patents issued and pending.

Dr. Thacker received the S.C. Sun Best Student Paper Award at the 2008 and 2006 IEEE International Interconnect Technology Conferences, and the Best Paper Honorable Mention Award at the 2006 IEEE International Test Conference.

Ivan Shubin received the M.S. degree from the Moscow Institute of Physics and Technology, Moscow, Russia, in 1988, and the Ph.D. degree in electrical engineering from Center for Research and Education in Optics and Lasers, School of Optics, University of Central Florida, Orlando, in 2001.

From 2001 to 2007, he was a Senior Engineer with Optium/Finisar, and then a Project Scientist at the University of California, San Diego. Since 2007, he has been with Oracle Labs, San Diego, CA, where he is a Principal Hardware Engineer working on advanced packaging solutions and platforms for electronic, optoelectronic and microelectromechanical systems applications, wafer scale packaging, 3-D integration, and novel photonic components. He is the author or coauthor of more than 50 technical papers and conference presentations.

Jin Yao received the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley, in 2007.

He is currently in the Photonics, Interconnects, and Packaging Group, Oracle Labs, San Diego, CA. He has co-authored more than 40 papers in refereed journals and conference proceedings. His research interests include photonic components, links, networking systems, and interconnects, with a focus on innovations of device and system performance. His research interests also include photonic integrated circuits, microelectromechanical systems/nanoelectromechanical systems and their system applications.

Dr. Yao received a California NanoSystems Institute Fellowship Award in 2002–2003.

Kannan Raj (S'87–M'90–SM'04) received the M.E. degree from the Indian Institute of Science, Bangalore, India, the M.S.E.E. degree from Virginia Polytechnic Institute and State University, Blacksburg, and the Ph.D. degree from George Mason University, Fairfax, VA.

He has led many successful advanced engineering, product development, and R&D teams. He is currently a Hardware Research and Operations Director in Oracle Labs, San Diego, CA, where he manages research teams in photonics, interconnects & packaging and operations of the UNIC silicon photonics program. He holds 40 issued patents and has co-authored more than 70 conference and journal publications.

Ron Ho (S'92–M'93–SM'08) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA.

His background includes ten years at Intel Corporation. He is currently an Architect at Oracle Labs, Redwood City, CA, where he is involved in a series of research programs centered around very large scale integration design, communications circuits, and advanced computer systems.

John E. Cunningham received the B.S. degree from the University of Tennessee, Knoxville, and the M.S. and Ph.D. degrees from the University of Illinois at Champaign-Urbana, Urbana, all in physics.

He is a veteran research scientist with over 25 years of University, Bell Labs, Sun Labs of Sun Microsystems/Oracle and start-up experience in the area of opto-electronic and semiconductor devices and packaging used within optical networks. Since joining Sun Labs, he has lead advanced packaging initiatives to develop Inter-chip Proximity Communication and worked on Si nano-photonics solutions for data communications within computers. Before joining Sun Microsystems, he was the Chief Scientist at Aralight, where he developed products based on the hybrid integration of vertical cavity surface emitting lasers and photodetectors with CMOS, a technology he codeveloped at Bell Laboratories. While at Bell Laboratories, he also pioneered eight world records on various types of quantum mechanically engineered devices, materials. He has co-authored more than 360 journal papers including some with Nobel Prize Laureates, as well as 40 U.S. patents. Before joining Bell Laboratories, he was a member of the research faculty in the Physics Department, University of Illinois, where he initiated the first Metals Molecular Beam Epitaxy. He is currently a Distinguished Engineer at Sun Labs of Oracle, San Diego, CA, and is the Principal Lead of the Advanced Packaging Group while simultaneously serving as a Co-Principal Investigator for the DARPA UNIC Project.

Ashok V. Krishnamoorthy (M'93–SM'10–F'12) received the B.S. (honors) degree from the California Institute of Technology, Pasadena, the M.S. degree from the University of Southern California, Los Angeles, and the Ph.D. degree from the University of California at San Diego, La Jolla.

He is currently a Hardware Architect and Chief Technologist, Photonics, at Oracle Labs, San Diego, CA, and the Principal Investigator for the DARPA UNIC initiative on silicon “photonics-to-the-processor.” Previously, he was a Distinguished Engineer and Director at Sun Microsystems responsible for advanced optical interconnect and silicon photonics development. He also spent several years as the CTO and President of AraLight, a Lucent technologies spinout developing high-density parallel optical products and technologies. Prior to that, he was an entrepreneur-in-residence at Lucent New Ventures group, and before that a member of technical staff in the Advanced Photonics Research Department at Bell Labs, Holmdel, NJ. He has worked for more than two decades on the integration of photonic devices with silicon CMOS circuits and on building switching and computing subsystems based on these components—including electro-optic modulators on silicon, quantum well devices, vertical cavity surface emitting lasers, and, most recently, Si/Ge photonics. He is the author or coauthor of more than 200 technical papers and eight book chapters and he holds 72 U.S. patents. He has served as a member or chair of more than 30 international conferences, and has been a Guest Editor for several technical journals.

Dr. Krishnamoorthy received the IEEE Distinguished Lecturer Award, the ICO International Prize in Optics, and the Chairman's Award from Sun Microsystems. He is a member of Tau Beta Pi, and is a Fellow of the Optical Society of America.