A Fully Integrated 4× 10-Gb/s DWDM Optoelectronic Transceiver Implemented in a Standard 0.13 μm CMOS SOI Technology

Adithyaram Narasimha, *Member, IEEE*, Behnam Analui, *Member, IEEE*, Yi Liang, *Senior Member, IEEE*, Thomas J. Sleboda, Sherif Abdalla, *Member, IEEE*, Erwin Balmater, Steffen Gloeckner, Drew Guckenberger, *Member, IEEE*, Mark Harrison, Roger G. M. P. Koumans, Daniel Kucharski, *Member, IEEE*, Attila Mekis, Sina Mirsaidi, Dan Song, and Thierry Pinguet

***Abstract*—Optical and electronic building blocks required for DWDM transceivers have been integrated in a 0.13 μm CMOS SOI technology. Using these building blocks, a 4 × 10-Gb/s single-chip DWDM optoelectronic transceiver with 200 GHz channel spacing has been demonstrated. The DWDM transceiver demonstrates an unprecedented level of optoelectronic system integration, bringing all required optical and electronic transceiver functions together on a single SOI substrate. An aggregate data rate of 40 Gb/s was achieved over a single fiber, with a BER of less than 10^{-12} and a power consumption of 3.5 W.**

***Index Terms*—CMOS optoelectronic transceiver, high-speed integrated circuits, integrated optical interleaver, photonic integrated circuits, silicon photonics, wavelength division multiplexing.**

I. Introduction

THE scaling of data networks towards increased throughput has served to highlight some of the major factors that limit serial data rates in copper channels. The problems of signal attenuation, dispersion, and crosstalk severely limit the reach of copper links at data rates exceeding 10 Gb/s [1], [2], and, while various mitigation techniques have been proposed [3], [4], such techniques are complicated and power-hungry and do not scale well to higher data rates. The combination of these problems has created an opportunity for other transceiver technologies. We have recently proposed and demonstrated a CMOS photonics optoelectronic transceiver approach that addresses some of the above stated issues [5], [6]. CMOS photonics offers an unprecedented level of optoelectronic integration, porting many diverse optical and optoelectronic devices to silicon-on-insulator (SOI) within a commercial and high-yielding SOI CMOS process. The coexistence of optical devices and CMOS transistors in the active silicon not only leverages the cost models of integration in silicon to provide a cheap solution, but also enables simultaneous reduction of system size and power dissipation. Fig. 1, reproduced from [6], shows a technology cross section of the CMOS photonics in SOI.

Manuscript received May 1, 2007; revised August 31, 2007. This work was supported by the Defense Advanced Research Projects Agency under Contract HR0011-05-9-0004.

The authors are with Luxtera Inc., Carlsbad CA 92011 USA (e-mail: anarasimha@luxtera.com).

Digital Object Identifier 10.1109/JSSC.2007.908713

An advantage of CMOS (or, more generally, silicon) photonics is to enable low cost scalability to higher data rates. Trivially, the constituent elements of a 10-Gb/s link may be replicated to create parallel links. This type of scaling to twice the native data rate was demonstrated in [6]. However, as transmission distances increase and additional throughput is desired, cable costs start to dominate the costs of such transceivers, since a separate outbound and inbound single-mode fiber is required to carry each 10-Gb/s data stream. An alternative scaling option is dense wavelength division multiplexing (DWDM), which optically multiplexes multiple 10-Gb/s data streams into a single fiber at the transmitter and optically demultiplexes them at the receiver, thus requiring only a single fiber to carry all of the data. Note that neither of these scaling methods requires an increase in the native data rate supported by a single electrical channel, thus remaining fully compatible with a maximum single-channel data rate of 10-Gb/s.

This paper focuses on the axis of scalability provided by DWDM, and the techniques and devices that enable it within the silicon photonics platform. DWDM systems are common in long-haul telecommunication systems but usually consist of large racks of extremely expensive equipment. In this paper, we show how DWDM functionality can be integrated on a single SOI chip. We detail one of the core DWDM building blocks available within the silicon photonics technology platform followed by a case study describing a fully integrated, four-channel, DWDM transceiver using interleaver optical muxes/demuxes achieving an aggregate data rate of 40-Gb/s over a single fiber at a BER $< 10^{-12}$.

The remainder of this paper is organized as follows. Section II expands on the silicon photonics device library described in [6] by describing an optical interleaver-based DWDM mux/demux integrated in the same technology. Section III describes a silicon photonics 4 × 10-Gb/s DWDM system architecture based on such a device, and Section IV includes experimental results from the measurement of such a system.

II. DWDM in Silicon Photonics

Here, we enhance the CMOS photonics device library described in [6] by adding a DWDM building block to the library. DWDM multiplexers and demultiplexers in silicon photonics are wavelength-selective filters that form the core of the DWDM system and are thus key enabling devices for DWDM

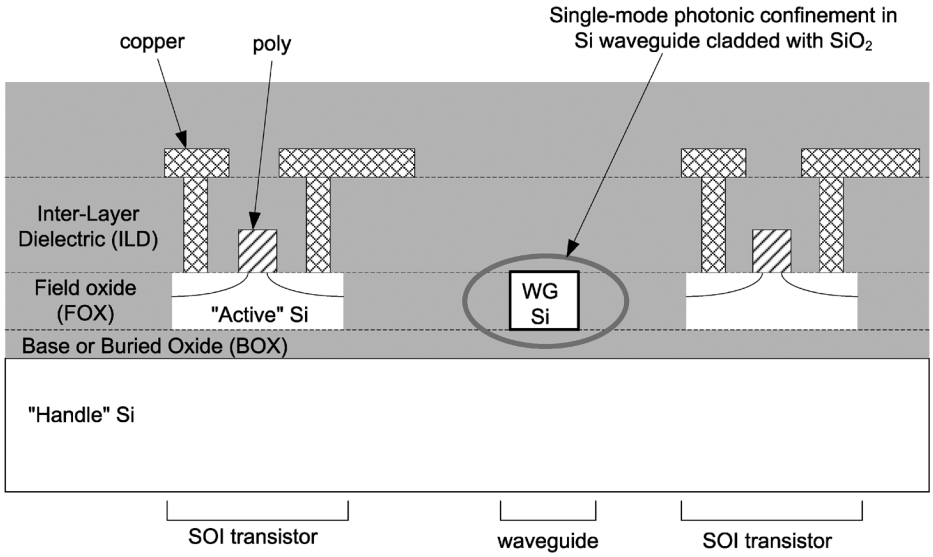

Fig. 1. Conceptional cross section of silicon CMOS photonics technology (not to scale).

integration. A DWDM filter consists of a strategic arrangement of waveguides and splitters that produce a desired filter function by causing desired wavelengths to interfere constructively and rejected wavelengths to interfere destructively within the device. The high-optical-index contrast between silicon and its oxide allows such devices to be very small, even for large channel count.

The need to precisely interfere light at optical C-band frequencies ($\sim$193 THz) makes these devices very sensitive to phase. Process-induced errors in the fabrication of the waveguides can cause changes in their structure that can induce random, unwanted phase delays in the optical signal. These unwanted phase delays have the effect of degrading the filter function and introducing unwanted loss and crosstalk. Additionally, the phase delay experienced by the light is sensitive to temperature. Both of these effects require us to electronically compensate for phase deviations from the desired phase.

Phase control in optical muxes/demuxes is achieved using a device called a thermal phase modulator (TPM) [6], [7]. From an electrical perspective, a thermal phase modulator is simply a resistor implanted in an optical waveguide. These resistors heat up the waveguide when a current flows through them. The rise in temperature results in an index change via the thermooptic effect in silicon (the Si index changes by approximately $2 \times 10^{-4}/\mathrm{K}$), causing a phase change $\phi(I)$ in the propagating light wave as a function of the current flowing through the TPM.

An optical DWDM device is thus a subsystem of optical waveguides and splitters in the active silicon, together with thermal phase tuners embedded in these waveguides to correct for fabrication errors and the temperature sensitivities of the device. The phase correction may be implemented using on-chip electronic drivers as part of a completely integrated closed-loop control scheme that allows the device to track a desired filter function. The ability to implement such control to compensate for fabrication inaccuracies and temperature sensitivities is a unique advantage arising out of the integration of optical and electronic functions in the same active silicon. These ideas were illustrated partially through the discussion of an optical AWG in [5]. They are further illustrated in this paper in the context of a DWDM device called an optical interleaver. While mux/demux architectures such as the AWG are advantageous for very high channel counts, the optical interleaver is the ideal choice for low-channel-count DWDM systems in silicon on account of its low insertion loss, size, crosstalk performance, and ease of control.

An optical interleaver consists of a series of cascaded, unbalanced optical interferometers. Such a device achieves a desired filter function by multiplying the spectral response of unbalanced interferometers with different length imbalances. Fig. 2 illustrates the way an unbalanced interferometer works. Such a device creates a periodic spectrally selective transfer function at its two outputs. Light entering the interferometer is split into two arms by an optical directional coupler. The length imbalance in the interferometer allows light to accumulate a substantial amount of wavelength dependent phase in one arm relative to the other arm. When the phase accumulated is an integral multiple of 2π, the two arms interfere in the output directional coupler, creating constructive interference in one port and destructive interference in the other port. When the phase accumulated is an odd multiple of π, the role of the output ports is reversed. The two outputs are complementary, conserving the net optical throughput, thus generating an "interleaving" function as illustrated in Fig. 2.

The electric field transfer function of such a structure is obtained by simply multiplying individual directional coupler transfer functions with the transfer function of the waveguides that connect them and is given by

$$\begin{bmatrix} \overrightarrow{E}_{\mathrm{upper\_out}}(\lambda) \\ \overrightarrow{E}_{\mathrm{lower\_out}}(\lambda) \end{bmatrix} = L_{\mathrm{DC}} \cdot \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{j}{\sqrt{2}} \\ \frac{j}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \cdot \begin{bmatrix} e^{-\alpha \Delta L - j\left(\frac{2\pi n_{\mathrm{eff}}(\lambda)}{\lambda}\Delta L + \phi(I_{\mathrm{upper}})\right)} & \\ & e^{-j\phi(I_{\mathrm{lower}})} \end{bmatrix} \cdot e^{-\alpha L - j\frac{2\pi n_{\mathrm{eff}}(\lambda)}{\lambda}L} \cdot \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{j}{\sqrt{2}} \\ \frac{j}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \cdot \begin{bmatrix} \overrightarrow{E}_{\mathrm{upper\_in}}(\lambda) \\ \overrightarrow{E}_{\mathrm{lower\_in}}(\lambda) \end{bmatrix} \quad (1)$$

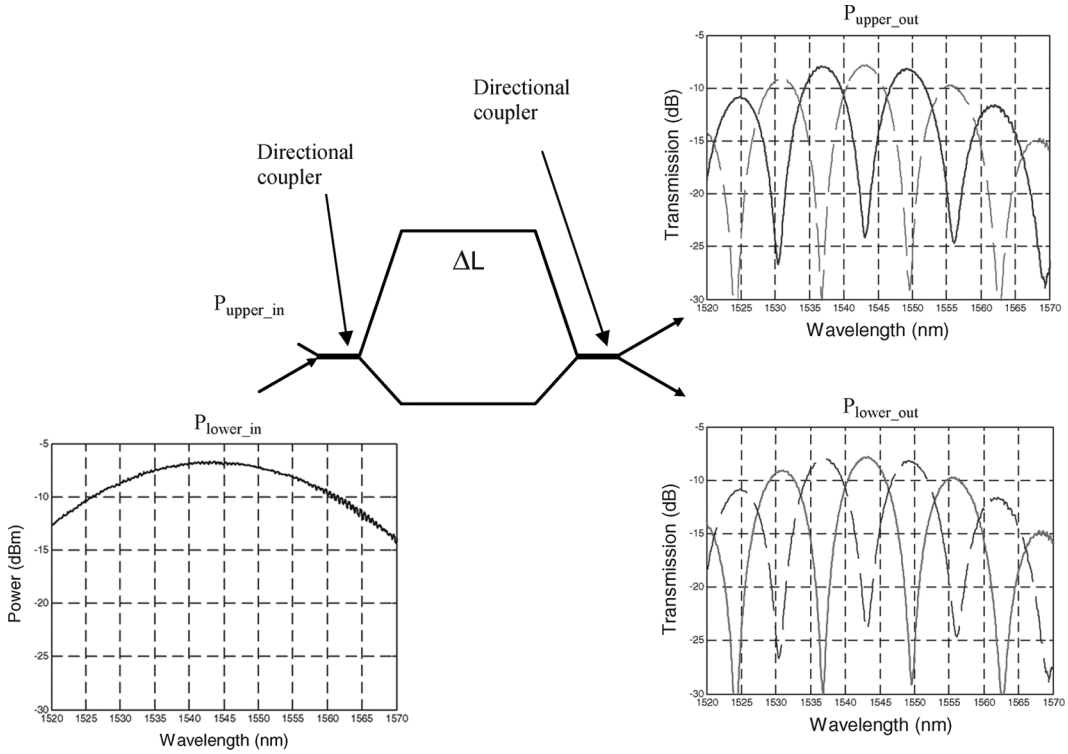

Fig. 2. Functionality of an unbalanced interferometer.

where $E_{\mathrm{upper\_in}}$, $E_{\mathrm{lower\_in}}$, $E_{\mathrm{upper\_out}}$, and $E_{\mathrm{lower\_out}}$ are the electric field vectors as a function of wavelength, corresponding to the optical intensities $P_{\mathrm{upper\_in}}$, $P_{\mathrm{lower\_in}}$, $P_{\mathrm{upper\_out}}$ and $P_{\mathrm{lower\_out}}$, respectively, as indicated in Fig. 2, L_{DC} is the optical power excess loss of a directional coupler, α is the field attenuation coefficient of the optical waveguide per unit length, $L = L_{\mathrm{lower}}$ is the length of the lower interferometer arm, ΔL is the difference between the upper and lower interferometer arms given by $\Delta L = L_{\mathrm{upper}} - L_{\mathrm{lower}}$, $n_{\mathrm{eff}}(\lambda)$ is the effective index of the optical waveguide as a function of wavelength, and $\phi(I)$ is is the phase induced in the upper or lower interferometer arm by a current I driven through the phase control element.

It is clear from (1) that the transfer function can be modified via the upper and lower interferometer phase control terms $\phi(I)$, allowing us to control the location of the peaks and nulls in the function by applying a suitable current to the phase control elements. This allows for shaping and positioning the nulls and peaks of the transfer function with respect to a desired wavelength, allowing us to compensate for process variations and temperature-induced variations in the transfer function of a given device.

Fig. 2 plots the output optical intensities, $P_{\mathrm{upper\_out}}$ and $P_{\mathrm{lower\_out}}$ corresponding to the fields given by (1) for the $P_{\mathrm{lower\_in}}$ shown in the figure. In this case, $P_{\mathrm{upper\_in}} = 0$. Note the complementary interleaving nature of the two output functions.

The periodicity of the nulls seen in the transfer function depends on the path-length difference between the two arms of the unbalanced interferometer and is given by

$$\Delta\nu = \frac{c}{\Delta L \cdot n_g} \tag{2}$$

where $\Delta\nu$ is the period of the nulls in the transfer function in gigahertz, c is the speed of light in a vacuum expressed in meters per second, ΔL is the path-length difference between the two arms of the unbalanced interferometer in nanometers, and n_g is the group velocity index of the optical waveguide used in the interferometer arms.

Unbalanced interferometers of different ΔL are cascaded together, multiplying out different transfer functions to generate a desired spectral response. A four-channel interleaver employs three unbalanced interferometers, as illustrated in the diagram and die shot in Fig. 3. Random variations in the process manifest themselves as local variations in the waveguide index, resulting in an inability to predict the exact peak of the individual stage transfer functions. Additionally, the imbalance in the interferometer makes its transfer function sensitive to temperature (via an effective index change) due to the thermooptic effect in silicon. These variations serve to degrade the overall interleaver transfer function resulting in crosstalk between the wavelengths. Tuning the interferometer using its control elements induces a change in the effective length of the waveguide, thus modifying one of the $\phi(I)$ terms in (1). To reduce the crosstalk, current is applied to low-speed TPMs in the upper or lower arms of the individual interferometers. Tuning the interleaver during chip startup is simply a matter of adjusting the current through the individual phase-control elements till the power at the desired output port is maximized for a given wavelength. Since the interleaver is designed for an inter-channel spacing of 200 GHz, maximizing one wavelength at a given output port ensures the suppression of other wavelengths in the DWDM wavelength plan. In operation, the three interferometers are continuously tuned by maximizing the output power from each channel using

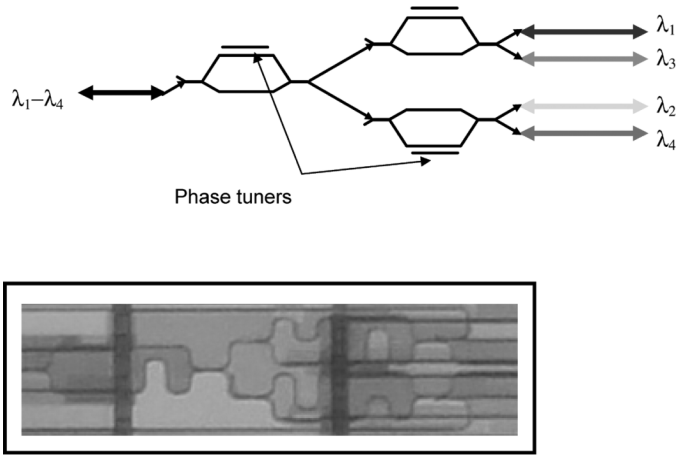

Fig. 3. Die shot and diagram of a four-channel interleaver consisting of three unbalanced interferometers.

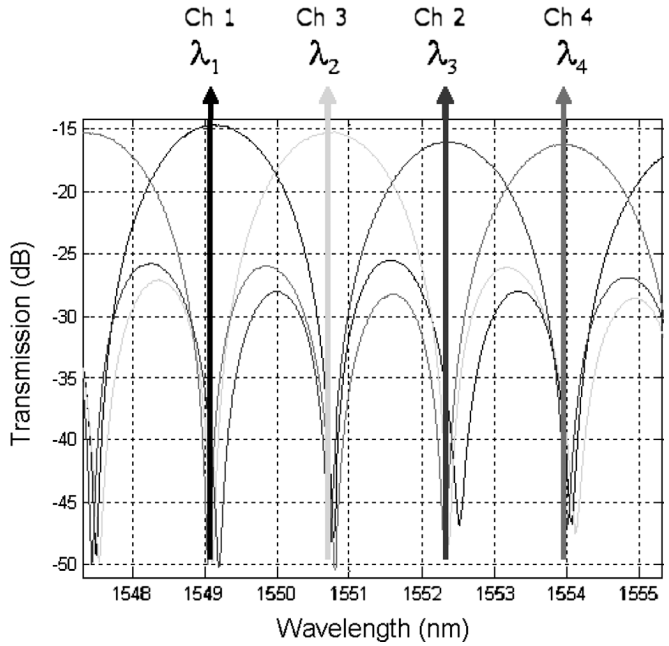

Fig. 4. Experimental tuned interleaver transfer function. Crosstalk suppression of better than 20 dB is achieved with an inter-channel spacing of 200 GHz.

an iterative algorithm that maintains the initially created wavelength plan.

Fig. 4 shows the experimentally obtained spectral response at each of the four output waveguides of a tuned interleaver, when it is operated as a demultiplexer. The channel passband spacing is 200 GHz with an adjacent channel crosstalk suppression of over 20 dB. The crosstalk has been optimized using the thermal control elements in the unbalanced interferometers.

III. 40-Gb/s Transceiver Architecture

Here, we describe how the elements from the CMOS photonics library come together with electronics blocks to achieve a 4×10-Gb/s DWDM transceiver. Fig. 5 shows the architecture of the four-channel DWDM transceiver. On-chip and off-chip components and functionalities have been clearly designated in

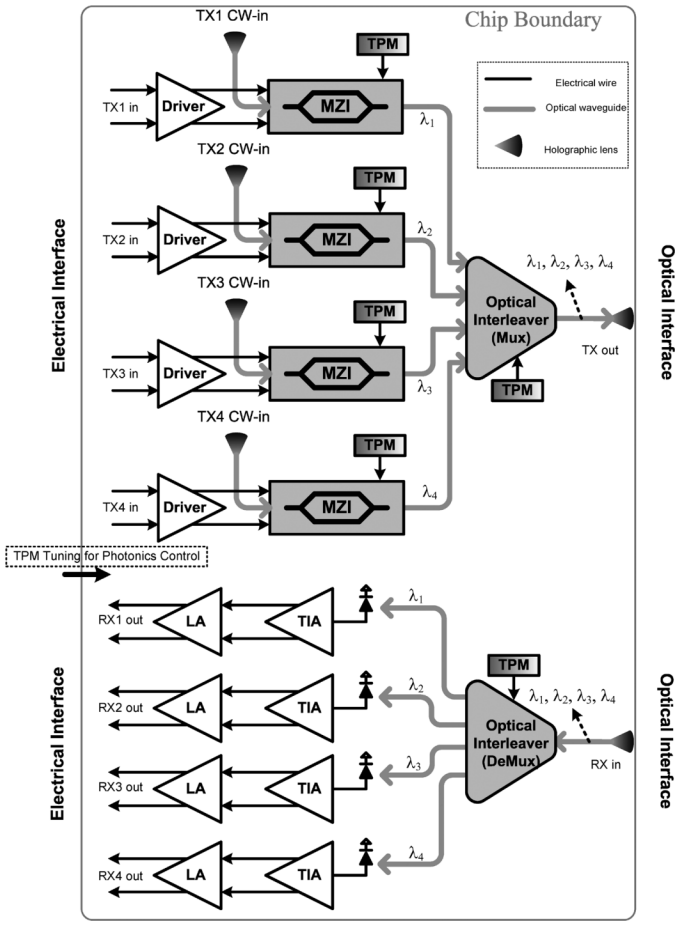

Fig. 5. System architecture of the 4×10-Gb/s.

the figure. In each transmitter channel, a continuous-wave (CW) laser is coupled into an on-chip waveguide using a normal-to-surface holographic lens (HL) [6]. The HL serves as an "optical pad" designed to couple light in and out of the silicon chip. Each optical channel operates at a different wavelength in the C-band.

The wavelength spacing between channels is 200 GHz. Each wavelength is modulated in OOK format using electrical data that is amplified by an on-chip modulator driver. The output of the driver provides a nominal 5 V differential swing to a high-speed optical modulator also integrated within the same silicon. The high-speed modulator uses a free-carrier effect based device in a Mach–Zehnder interferometer (MZI) [5], [6]. The modulated outputs of the four MZIs are combined into a single optical waveguide using a 200 GHz spaced, four-channel interleaver optical multiplexer. The multiplexed optical signal is coupled out of the chip using another HL. The receiver signal path begins with an HL that couples the received optical data containing all four wavelengths from single-mode fiber into the chip. An optical waveguide steers the light to an interleaver optical demultiplexer that separates the optical signal into its four constituent wavelengths with a transfer function identical to that of the interleaver multiplexer (Fig. 4). Each output of the interleaver demultiplexer is guided to a high-speed p-i-n photodetector (PD) that is flip-chip mounted on the die. Light is coupled from the waveguide into each PD using an HL. The PD is mounted on

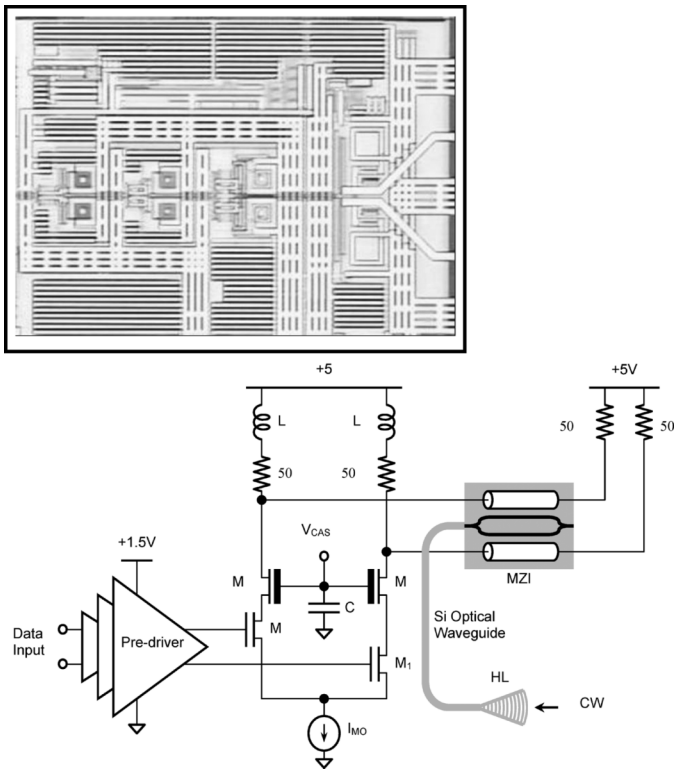

Fig. 6. Die shot and schematic of H/S MZI modulator driver.

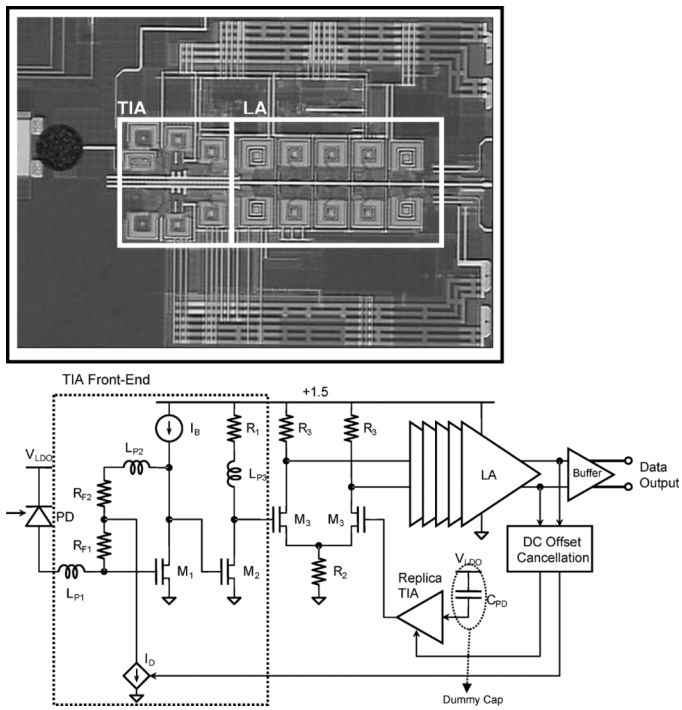

Fig. 7. Die shot and schematic of high-speed receiver front-end.

the die in the immediate vicinity of a TIA. A following limiting amplifier (LA) further amplifies the electrical signal.

Fig. 6 shows a die shot and schematic of the modulator driver circuit [6]. It consists of a three-stage predriver that operates from a 1.5 V supply followed by a cascoded differential output driver from a 5 V supply. The output stage drives the 50 Ω MZI transmission lines with 50 Ω far-end terminations. The main challenge in designing the driver is to achieve high voltage swing for the modulator driver while achieving high-speed operation using thin-oxide 0.13 μm transistors. Therefore, a 5 V supply is used to enable large swing while a cascoded thick-oxide transistor is used to protect the switching transistors from over-voltage conditions. The thick-oxide devices also feature a longer channel and can withstand much higher gate and drain-to-source voltages. These thick-oxide transistors are standard to the CMOS SOI process used. Inductive peaking is used in the output stage as well as the predriver to achieve fast transitions. Each driver consumes 575 mW in nominal conditions. The single-ended input signal can be as low as 50 mVp-p, for an optical output with an extinction ratio of about 4 dB.

A die shot and schematic of the receiver front-end is shown in Fig. 7. It consists of a high speed p-i-n PD connected to a TIA followed by a five-stage LA [6], [8]. The PD is flip-chip bonded to the chip. The circular pad visible in the die shot in Fig. 7 corresponds to the anode of the photodiode, which is directly connected to the TIA.

The TIA input stage is based on the resistive feedback structure that achieves low input impedance for high-speed operation. In addition, broadband matching inductors are used in series at the input and in the feedback path. The third stage converts the single-ended input to a differential signal. The second input of this differential amplifier is connected to a replica of the TIA front end to convert the TIA to pseudodifferential architecture for minimizing crosstalk. The input node of this reference branch is terminated via a dummy capacitor $C_{\rm PD}$, which matches the junction capacitance of the PD. The LA consists of resistively loaded differential stages with inductive peaking. The output of the LA is filtered to extract the dc-level information of the signal, which is utilized in an offset-cancellation loop. This loop also compensates for the nonzero average input. The LA drives a 50 Ω output buffer. Each receiver operates from a single 1.5 V supply. The PD is reverse-biased using an integrated low-drop regulator that regulates the 5 V supply. The total power consumption is 120 mW including the output buffer.

Fig. 8 shows a die shot of the transceiver with the functional areas (e.g., modulator drivers, MZIs, optical mux/demux, and TIA/LA) designated. It can be seen from the die shot that the four MZI modulators used have different lengths. This was part of an optimization experiment, allowing us to assess the relative contribution of waveguide length and modulation depth to overall system penalty. It has no impact on the demonstration of the system itself.

IV. System Testing and Results

The 4 × 10-Gb/s chip was wirebonded to a high-speed board with transmission lines and SMA connectors as shown in Fig. 9. The board provides low- and high-speed electrical connectivity to the die as well as adequate thermal sinking for heat generation. All connections to the die are provided via gold wirebonds to the PCB traces around the center cavity. Sixteen edge-mounted Rosenburger SMA connectors allow cabling of the high-speed Tx and Rx electrical signals on the west and east of the board, respectively. The signals are designed as controlled impedance, differentially paired microstrip transmission lines.

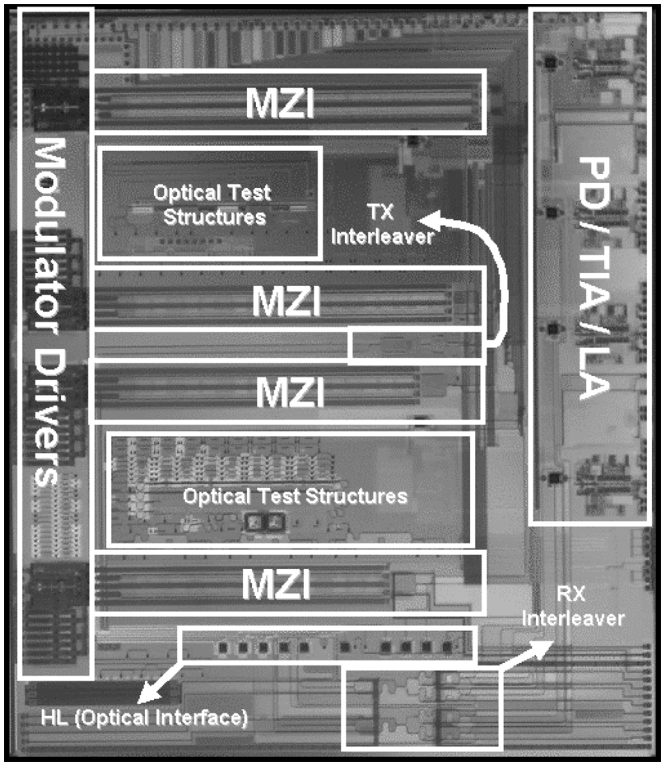

Fig. 8. Die shot of 4 × 10-Gb/s transceiver.

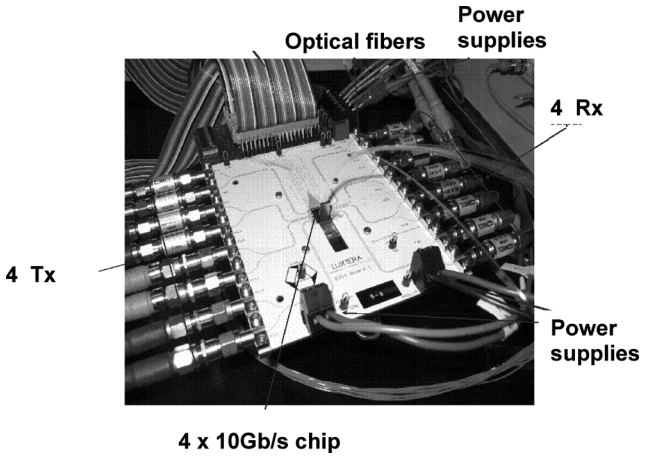

Fig. 9. 4 × 10-Gb/s CMOS photonics tranceiver die wirebonded to high-speed board with SMA connectors.

The PCB is constructed of high-Tg FR-4, with the top layer made out of Rogers 4350. The die is mounted to a copper heat spreader mechanically fixed to the back of the board. A fiber array (FA) is attached to the die over the array of on-die HLs and serves to bring light in and out of the die.

Fig. 10 shows an optical eye diagram at the transmitter output when only one of the four transmitter channels is enabled, showing the functionality of a single transmitter channel. Fig. 11 shows an optical spectrum at the transmitter output optical interface when all four transmitters are on. Four wavelengths are seen in the optical spectrum, demonstrating that all four wavelengths are propagating down a single fiber. To test

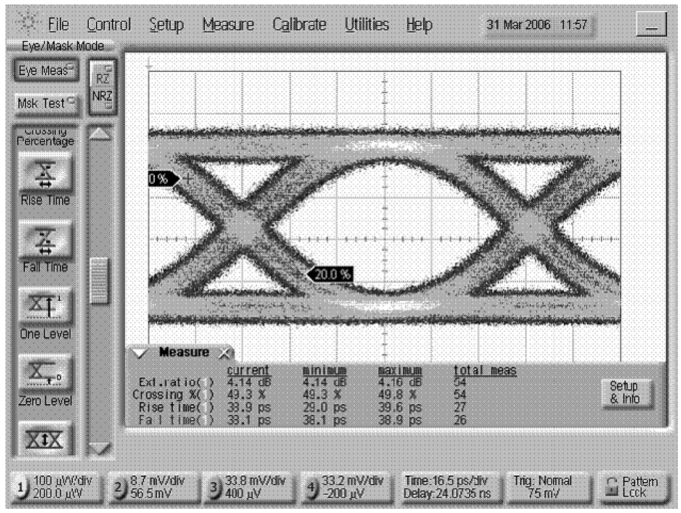

Fig. 10. Optical eye diagram at transmitter output when only one transmitter is turned on.

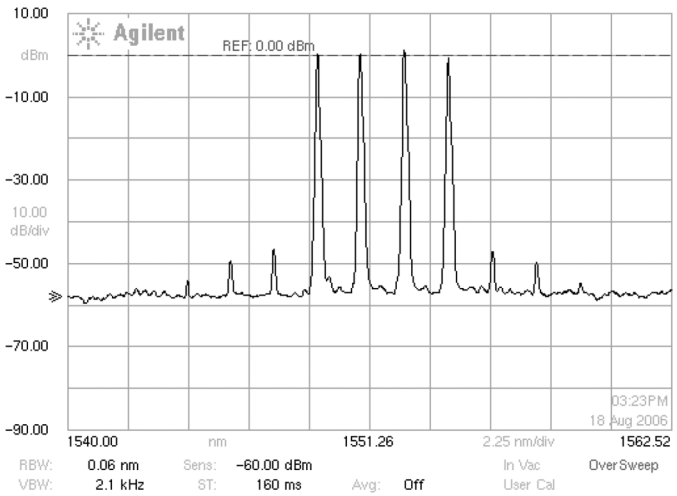

Fig. 11. Optical spectrum at transmitter output when all four transmitters are turned on. All four channels are presented and are transmitted down a single fiber.

the full optoelectronic operation, this DWDM optical output is externally looped back to the DWDM receiver input. All of the four channels are operating simultaneously (each at 10-Gb/s) with uncorrelated data applied to all the four transmitter inputs.

At the receiver, the light is coupled into the chip through an HL and is optically demultiplexed. Receiver channel 4 is monitored using an on-chip optical tap that is routed through an optical waveguide to an HL. Fig. 12 shows the optical spectrum at the channel–4 monitor tap. The spectrum shows that channels 1–3 have been suppressed by more than 20 dB by the interleaver optical demux.

Fig. 13 shows a bathtub curve for Rx 3, indicating an eye opening of better than 0.6 UI at a BER of 10^{-9}. Fig. 14 shows receiver sensitivity curves for all four channels during full DWDM operation. The TIA in each channel achieves an optical sensitivity of better than -15 dBm average power for BER $= 10^{-12}$. The data stream is a 10-Gb/s, $2^{31} - 1$ pseudorandom bit sequence (PRBS). One receiver output was

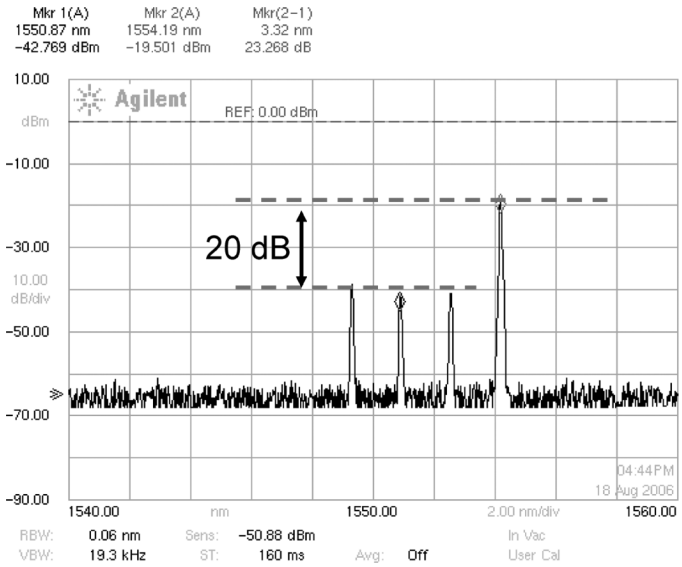

Fig. 12. Optical spectrum at channel–4 optical tap, showing suppression of channels 1–3 by the interleaver optical demultiplexer.

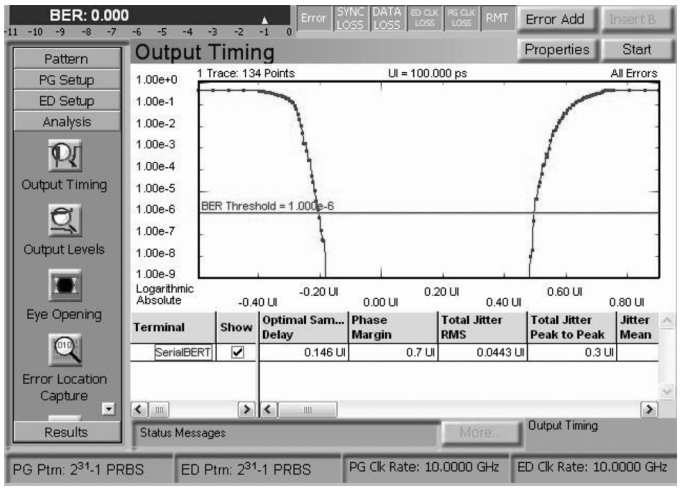

Fig. 13. Bathtub curve for Rx 3 at 10-Gb/s.

examined, and error-free operation was observed when measured overnight $(\text{BER} < 10^{-14})$. The measured optical power penalty caused by optoelectronic crosstalk between channels is 0.6 dB in this test. The penalty observed is due to the combined effect of optical crosstalk (within the optical multiplexer and demultiplexer) together with electronic crosstalk on the die as well as the test board.

The clustering of Rx1 and Rx4 data relative to the data for Rx2 and Rx3 in Fig. 14 is a consequence of the modulator length experiment described in Section III. While Rx1 and Rx4 have very different modulation penalties on account of different modulator lengths, these are compensated for by the significantly different optical waveguide lengths within the transceiver, which accounts for their different optical loss. Similarly, Rx2 and Rx3 have similar but also compensating modulation penalties and waveguide lengths, however, these are different from Rx1 and Rx4.

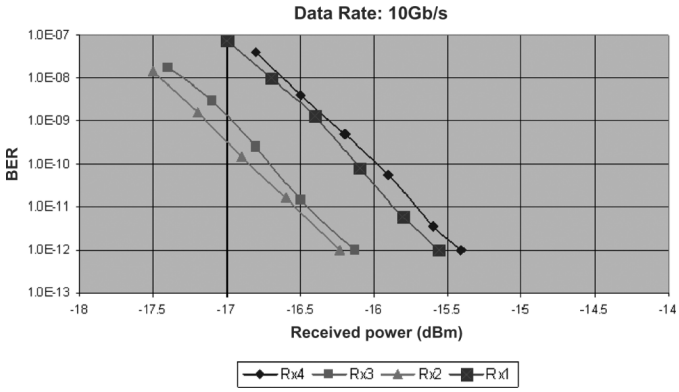

Fig. 14. Receiver sensitivity curves for four-channel 4 × 10-Gb/s DWDM transceiver.

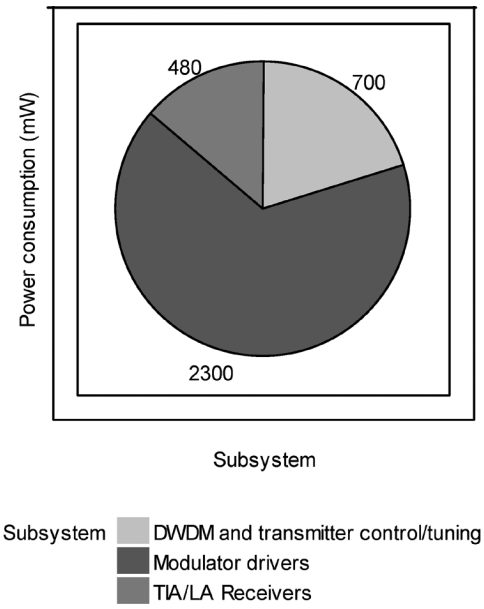

Fig. 15. Pie chart showing subsystem-wise division of power consumption for 4 × 10-Gb/s DWDM transceiver.

Fig. 15 shows a typical power consumption pie-chart for the different subsystems of the chip discussed in this paper. The typical power consumption under nominal conditions is ∼3.5 W, the majority of which is consumed in the modulator drivers.

V. Conclusion

A fully integrated 4 × 10-Gb/s DWDM transceiver system incorporating optical interleavers has been demonstrated within a SOI CMOS photonics platform. The system achieved an aggregate data rate of 40 Gb/s with a BER of better than 10^{-12}. All key optical and electronic components were integrated on a single SOI substrate. To the best of our knowledge, this is the first full integration of a DWDM system in a silicon technology platform.

Acknowledgment

The authors would like to thank V. Sadagopan and G. Masini for assistance with chip layout, B. Ranjbaran for CAD related issues, and C. Gunn for helpful discussions.

REFERENCES

[1] J. W. Goodman, F. J. Leonberger, S.-Y. Kung, and R. A. Athale, "Optical interconnections for VLSI systems," *Proc. IEEE*, vol. 72, no. 7, pp. 850–866, Jul. 1984.

[2] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronics chips," *Proc. IEEE*, vol. 88, no. 6, pp. 728–749, Jun. 2000.

[3] S. Bates and K. Iniewski, "10 GBPS over copper lines—State of the art in VLSI," in *Proc. 5th Int. Workshop Systems-on-Chip for Real-Time Applications*, Jul. 2005, pp. 491–494.

[4] T. Beukema *et al.*, "A 6.4-Gb/s CMOS SerDes core with feed-forward and decision-feedback equalization," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2633–2645, Dec. 2005.

[5] A. Huang *et al.*, "A 10 Gb/s photonic modulator and WDM MUX/DEMUX integrated with electronics in 0.13-μm SOI CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2006, pp. 244–245.

[6] B. Analui, D. Guckenberger, D. Kucharski, and A. Narasimha, "A fully integrated 20-Gb/s optoelectronic transceiver implemented in a standard 0.13-μm CMOS SOI technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2945–2955, Dec. 2006.

[7] G. T. Reed and A. P. Knights, *Silicon Photonics: An Introduction*. New York: Wiley, 2004.

[8] F. Pera and S. P. Voinigescu, "An SOI CMOS, high gain and low noise transimpedance-limiting amplifier for 10 Gb/s applications," in *Proc. IEEE Radio Frequency Integr. Circuits Symp.*, Jun. 2006.

Adithyaram Narasimha (S'99–M'04) received the B.S. degree in electrical and electronics engineering from the Birla Institute of Technology and Science, Pilani, India, in 1998, and the Ph.D. degree in electrical engineering from the University of California, Los Angeles (UCLA), in 2004.

He is currently a Senior Staff Engineer with Luxtera Inc., Carlsbad, CA, where he is part of the Product Development Team working on the company's first high-speed optoelectronic transceiver. Previously, at Luxtera, he worked on porting integrated DWDM device technologies to silicon photonics and on optoelectronic transceiver system and subsystem architectures. While at UCLA, he conducted research on photonic transmission systems and diffractive photonic devices in SOI. Prior to UCLA, he was with Siemens Components Pte. Ltd. (now Infineon), Singapore, where he worked on DRAMs in the Signal Processing IC Group.

Dr. Narasimha was a corecipient of the 2006 IEEE Lewis Winner Best Paper Award.

Behnam Analui (S'97–M'05) received the B.S. and M.S. degrees in electronics engineering from the Sharif University of Technology, Tehran, Iran, in 1998 and 2000, respectively, and the Ph.D. degree in electrical engineering from the California Institute of Technology (Caltech), Pasadena, in 2005.

He was with the Mixed-Signal Communications IC Design Group, IBM T. J. Watson Research Center, Yorktown Heights, NY, in 2003, where he developed an eye-opening monitor for 10-Gb/s multimode fiber links. He is now a Senior Staff Engineer with Luxtera Inc., Carlsbad, CA, as part of the team that is leading the development of the world's first silicon photonics-based transceivers. His technical interests include signal integrity issues in high-speed serial links, jitter, and integrated silicon photonics systems.

Dr. Analui was the Silver Medal winner of the National Mathematics Olympiad in 1994. He was the recipient of Caltech's Atwood Fellowship in 2000 and the Analog Devices Outstanding Student Designer Award in 2002.

Yi Liang (SM'02) received the Ph.D. degree in electrical engineering from Northern Jiaotong University, Beijing, China, in 1988.

She then taught fiber optics and researched several state and federal government-funded programs of optical fiber applications at Northern Jiaotong University from 1989 to 1994. She joined the Ultrafast Switching Group, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, in 1995, where she focused on the research of the 100-Gb/s all-optical switching program for more than three years. From 2000 to 2004, she was with Sorrento Networks Inc., San Diego, CA, where she was involved with the development of optical switching systems and DWDM optical transport systems. She joined Luxtera Inc., Carlsbad, CA, in October 2004, where she focuses on the characterization and validation of CMOS photonic devices and systems.

Thomas J. Sleboda received the B.S. degree in materials engineering and the M.S. degree in metallurgical engineering from the Georgia Institute of Technology, Atlanta, in 1994 and 1996, respectively. His graduate studies were conducted at the Packaging Research Center.

He is currently a Senior Staff Engineer with Luxtera Inc., Carlsbad, CA, where he is involved with packaging development of products that leverage silicon photonics technology. Prior to joining Luxtera, he was with Cameron Health, developing packaging solutions for subcutaneous implantable defibrillators. Prior to this position, he worked with Big Bear Networks developing packaging technology for OC-768 transceiver products. Prior to Big Bear, he was with Motorola (now Freescale) Semiconductor Products Sector, developing packaging solutions for 3G cellular applications, PowerPC (IBM Somerset partnership), and the world's first flip-chip plastic ball grid array (FC-PBGA) product.

Sherif Abdalla (S'98–M'02) received the B.Sc. degree in electrical engineering and communication from Cairo University, Cairo, Egypt, in 1992, and the M.S. degree in electrical engineering from the University of Toronto, Toronto, ON, Canada, in 2002.

From 2002 to 2005, he was with Snowbush, Toronto, working on multi-Gb/s serial links. Since 2005, he has been with Luxtera Inc., Carlsbad, CA, where he is the Director of Engineering responsible for electronics IC design activities. Prior to his IC design career, he has worked on different technologies in robotics, machine vision, and oil-field measurement systems.

Erwin Balmater, photograph and biography not available at the time of publication.

Steffen Gloeckner received the Ph.D. degree in physics from Friedrich-Schiller-University, Jena, Germany.

He is a Director of Engineering with Luxtera Inc., Carlsbad, CA, where he is responsible for managing the Photonics Components and System Design Group. Prior to Luxtera, he held design and product engineering positions with the Fraunhofer Institute for Applied Optics and Precision Engineering, Jena, Germany, and OMM, a manufacturer of MEMS-based optical switches.

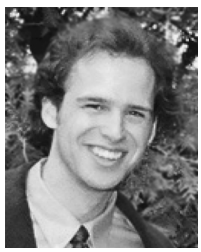

Drew Guckenberger (S'00–M'05) received the B.A.Sc. degree in systems design engineering from the University of Waterloo, Waterloo, ON, Canada, in 2000, and the M.S. and Ph.D. degrees in electrical and computer engineering from Cornell University, Ithaca, NY, in 2003 and 2005, respectively.

He spent the summers of 2002–2004 at IBM's T. J. Watson Research Center, Yorktown Heights, NY, where he worked mainly on the development of high-speed optical receivers and receiver arrays. He has been with Luxtera Inc., Carlsbad, CA, since

2005. His research interests focus mainly on the development of low-power high-performance analog circuits for communication system applications.

Dr. Guckenberger was the recipient of the Cornell IEEE Teaching Assistant of the Year Award in 2000-2001 and has been the recipient of an IBM Ph.D. Fellowship (2001-2003), a Qualcomm Design Fellowship (2004), and an Intel Ph.D. Fellowship (2004–2005). He also received the 2004 IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM) Best Student Paper Award.

Mark Harrison, photograph and biography not available at the time of publication.

Roger G. M. P. Koumans was born in Heerlen, The Netherlands, on January 17, 1971. He received the B.S. and M.S. degrees (*cum laude*) in electrical engineering from Eindhoven University of Technology, Division of Optical Telecommunications, Eindhoven, The Netherlands, in 1994, and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 1996 and 2000, respectively. His Ph.D. dissertation focused on the modeling, characterization, and applications of mode-locked semiconductor and fiber lasers as well as on wavelength-division multiplexing, ultrashort pulse measurements and pulse shaping, wavelength conversion, fiber Bragg gratings, and nonlinear optics.

In partial fulfillment of his master's degree from Eindhoven University of Technology, he worked for Philips Optoelectronics Centre, Eindhoven, on mode-locked semiconductor lasers. After obtaining the Ph.D. degree, he joined Orbits Lightwave, Inc., Pasadena, where he concentrated his research efforts on the development of ultrastable, low-noise, single-frequency fiber lasers that now find application in military, LIDAR, and subterranean applications. In December 2003, he joined Luxtera Inc., Carlsbad, CA. His main research efforts focus on the modeling, design, and testing of novel CMOS photonic devices and systems, semiconductor and Raman lasers, electrooptical modulators, and other silicon-on-insulator (SOI) devices. He has been awarded several U.S. and international patents for his scientific contributions in his field.

Daniel Kucharski (S'00–M'05) received the B.S. degree in electrical engineering from the Rochester Institute of Technology, Rochester, NY, in 1998, and the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, in 2005.

From 1998 to 2000, he was with IBM Microelectronics, Burlington, VT, where he worked on ASIC library development. He spent the summers of 2001–2004 at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he worked on parallel optical data links in SiGe and CMOS. In 2005 he joined Luxtera Inc., Carlsbad, CA, where he is currently involved in development of CMOS circuits for silicon photonics applications.

Attila Mekis received the B.S. degree in physics from Yale University, New Haven, CT, and the Ph.D. degree in theoretical condensed matter physics from the Massachusetts Institute of Technology (MIT), Cambridge, in 2000,

He is currently a Principal Engineer with Luxtera Inc., Carlsbad, CA, where he is currently responsible for the development of integrated optoelectronic components. Prior to joining Luxtera in 2004, he was with Inplane Photonics and Clarendon Photonics, developing integrated optics components for waveguide amplifiers and reconfigurable OADMs, respectively. While at MIT, he conducted research on photonic crystal-based optical components and worked as a consultant for Bell Labs designing photonic crystal lasers.

Sina Mirsaidi received the B.S. degree in electrical engineering from Tehran University, Tehran, Iran, in 1989, and the M.S. and Ph.D. degrees in electrical engineering from École Supérieure d'Électricité (SUPELEC), Paris, France, in 1993 and 1997, respectively.

From 1997 to 1998, he was a Research Engineer with the University of Toronto, Toronto, ON, Canada, where he pursued his interests in the areas of speech compression, adaptive signal processing for nonuniformly sampled signals, as well as detection/estimation strategies for wireless receivers. In 1999, he joined Nortel Networks, Ottawa, ON, Canada, as a Member of the Technical Staff, where he worked on design and development of broadband wireless access products for residential and commercial applications. From 2000 to 2001, he was actively involved in product definition and design of Nortel Networks' Long Haul (LH) fiber-optic-based DWDM systems. From 2001 to 2004, he was a System Architect with Ceyba Inc., Ottawa, where he was in charge of system design and implementation of large/high-capacity end-to-end core optical network systems. In 2004, he joined Luxtera Inc., Carlsbad, CA, where he is currently managing the system design and development of fiber-optic-based high-speed transceivers. During 2001, he was also a System Engineering consultant for Cisco Systems.

Dan Song, photograph and biography not available at the time of publication.

Thierry Pinguet received the engineering degree and the M.S. degree from the Ecole Polytechnique, Palaiseau, France, and the Ph.D. degree in applied physics from Stanford University, Stanford, CA.

He is currently a Principal Engineer with Luxtera Inc., Carlsbad, CA, where he has since 2002 been leading the development of the company's CMOS photonics technology. While at Stanford, his research concentrated on the design, fabrication, and testing of guided-wave and bulk nonlinear optical devices in GaAs for applications ranging from military countermeasures to optical communications. His work has led to numerous breakthroughs in photonic design and fabrication techniques, and he is the owner of 23 patents.