A Fully Integrated 20-Gb/s Optoelectronic Transceiver Implemented in a Standard 0.13-μm CMOS SOI Technology

Behnam Analui, *Member, IEEE*, Drew Guckenberger, *Member, IEEE*, Daniel Kucharski, *Member, IEEE*, and Adithyaram Narasimha, *Member, IEEE*

***Abstract*—A dual-channel 10 Gb/s per channel single-chip optoelectronic transceiver has been demonstrated in a 0.13-μm CMOS SOI technology. The transceiver integrates conventionally discrete optoelectronic functions such as high-speed 10-Gb/s electro-optic modulation and 10-Gb/s optical reception on an SOI substrate using a standard CMOS process. The high optical index contrast between silicon ($n = 3.5$) and its oxide ($n = 1.5$) allows for very large scale integration of optical devices, while the use of a standard CMOS process allows these devices to be seamlessly fabricated together with electronics on the same substrate. Such a high level of optoelectronic integration is unprecedented, and serves to substantially reduce system footprint and power dissipation, allowing efficient scaling to higher data rates and broader functionality.**

This paper describes the photonic components, electronic blocks, and architecture of a CMOS photonic transceiver that achieves an aggregate data rate of 20 Gb/s in a dual-channel package, with a BER of less than 10^{-15} and a power consumption of 1.25 W per channel with both channels operating simultaneously.

***Index Terms*—CDR, CMOS transceivers, holographic lens, integrated optoelectronics, Mach–Zehnder interferometer, modulator driver, optical interconnects, optical modulation, optical transceivers, silicon photonics, transimpedance amplifier, VCO.**

I. Introduction

As DATA NETWORKS scale to increase throughput, some of the major obstacles that limit serial data rates in copper channels are signal attenuation, dispersion, and cross-talk. These problems become more pronounced at higher data rates, drastically limiting the reach of copper links at data rates of 10 Gb/s and above [1], [2]. This can be mitigated to some extent with advanced pre- and post-equalization techniques [3], [4], but such solutions require complicated control and adaptation algorithms, consume considerable power, and do not scale well to higher data rates. The THz-scale, low loss bandwidth of single-mode fiber (SMF) in the optical C-band centered at 1550 nm has made it the only viable transmission medium for long distance, broadband communication networks. Recently, SMF has also become the transmission medium of choice in metro access networks and enterprise LAN backbones. The scalability of a single strand of fiber to higher bandwidths through optical techniques such as wavelength division multiplexing (WDM), together with the limited capacity and reach of copper links, suggest that fiber will replace copper even for shorter links, at data rates of 10 Gb/s and above.

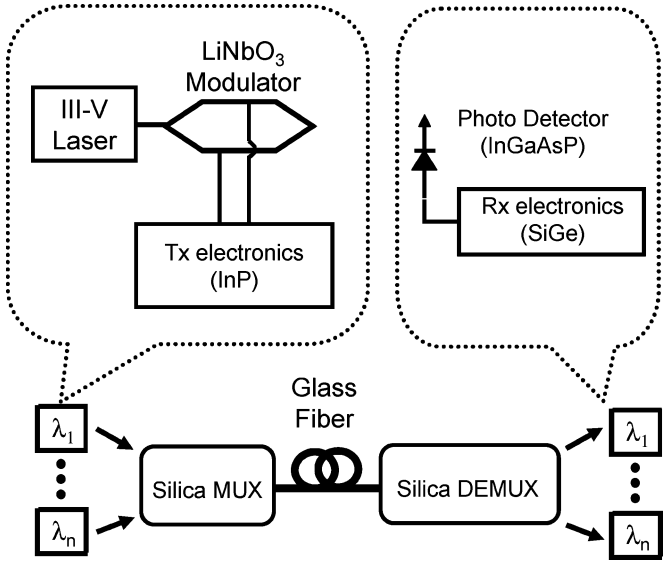

Fig. 1. Conventional multi-wavelength optical communication system.

Optical data transceivers operating at 10 Gb/s have been implemented in CMOS in the past [5]–[7], but they demonstrate a low level of optoelectronic integration. They are based on conventional optics and suffer from the legacy of their telecom counterparts by relying on hybrid solutions that co-package different optical components manufactured in a diversity of material systems using highly customized processes (Fig. 1). Such an approach results in very high cost, dominated by the packaging and testing of these different components. Porting optical and optoelectronic devices to silicon (i.e., "siliconizing" the photonics) within a commercially deployed and high yielding CMOS process leverages the cost models of monolithic electronic integration in silicon, and it reduces size and power dissipation while simultaneously achieving all the other benefits of integration [8].

Elements of such "silicon photonics" technology were introduced in [9]. This paper focuses on an application of the same technology to creating highly-integrated communication systems. It describes the architecture, key building blocks, implementation, and testing of a dual-channel, 10 Gb/s per channel, C-band (1550 nm) optoelectronic transceiver integrated in a silicon-on-insulator (SOI) material system. It demonstrates the integration of relevant optical, electronic and optoelectronic trans-

Manuscript received June 3, 2006; revised August 23, 2006.

The authors are with Luxtera Inc., Carlsbad CA 92008 USA (e-mail: behnam@luxtera.com). Authors are listed alphabetically.

Digital Object Identifier 10.1109/JSSC.2006.884388

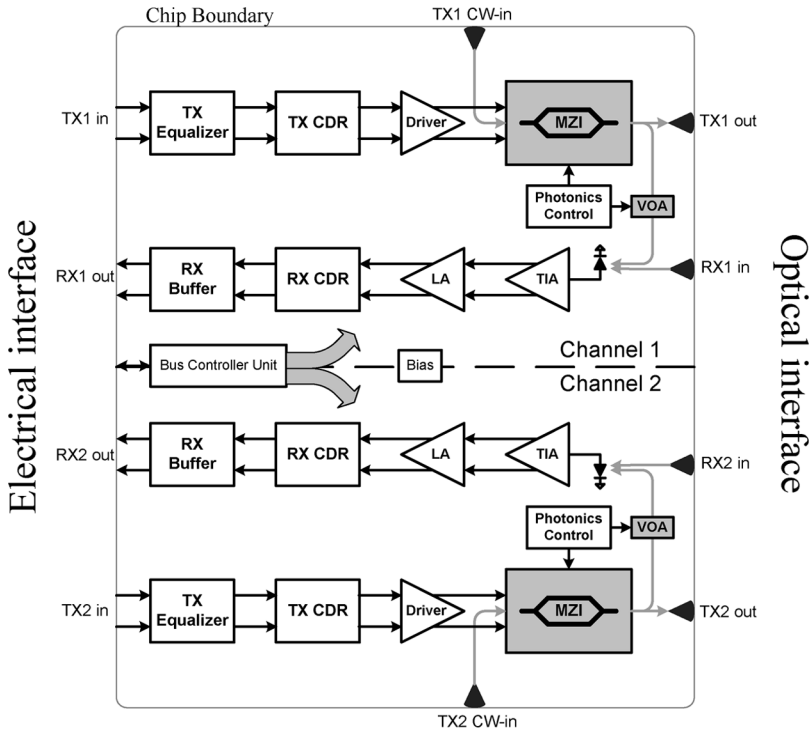

Fig. 2. Dual-channel 10 Gb/s per channel optoelectronic transceiver architecture.

ceiver functions on a single substrate using a standard 0.13-μm CMOS SOI process, achieving an aggregate data rate of 20 Gb/s in a very small form factor. To the best of our knowledge, this is the first fully integrated transceiver system in silicon, incorporating VLSI analog, digital, and RF circuits with photonic components on the same SOI substrate.

This paper is organized as follows. The overall system architecture is presented in Section II, followed by a detailed discussion of the system building blocks in Section III, including separate subsections for the optical, optoelectronic, and electronic components. Experimental results demonstrating system performance are presented in Section IV, followed by conclusions in Section V.

II. Architecture

This high-speed transceiver system is designed for non-return to zero (NRZ) data communication over more than 2 km of SMF. It consists of two parallel high-speed optoelectronic transceivers, each running at 10 Gb/s, integrated together on the same die. The complete architecture of the system is shown in Fig. 2. Each channel consists of a transmit (TX) and a receive (RX) path.

The TX converts a high-speed electrical input into an optical output. The high-speed electrical input to the TX passes through an adjustable post emphasis equalizer. The equalizer can compensate for a moderate amount of high-frequency attenuation due to the electrical channel prior to the TX input. The data is then passed to a clock and data recovery (CDR) circuit, where it is retimed to remove jitter and create a clean, high-amplitude signal for a modulator driver (MD). The multi-stage MD further amplifies this signal to provide a high voltage swing to a high-speed optical modulator integrated within the same silicon. The whole electrical signal path of the TX operates based on differential signaling.

The high-speed modulator operates based on the principle of a Mach–Zehnder Interferometer (MZI), as will be discussed in Section III. Functionally, the MZI is a four-port network. It accepts a continuous wave (CW) optical signal with $\lambda = 1.55\,\mu$m, and high-speed electrical data as inputs. The two outputs are the modulated optical signal and its complement. Even though the system described in this paper was designed for an optical wavelength of 1.55 μm, the design is portable to other wavelengths, including those around 1.3 μm. The optical CW input to the MZI is coupled into the chip by an integrated holographic lens (HL). The HL serves as an "optical pad" and is designed to couple normally incident light in and out of the silicon chip. Once the light is on chip, it is guided through an integrated optical waveguide to the MZI input, as shown in gray in Fig. 2. The optical waveguide allows broadband optical signals to propagate across the chip with low attenuation. The output of the MZI is coupled out of the chip into SMF using another HL. The architecture allows the CW optical source to be either a commercially available external laser, or a semiconductor laser mounted directly on the die and co-packaged with the transceiver.

The RX signal path on the die starts with an HL that couples the received optical data signal from SMF into the chip. An optical waveguide steers the light to a high-speed p-i-n photodetector (PD) that is flip-chip mounted on the die. Light is coupled from the waveguide into the PD using another HL. The PD is mounted on the die in the immediate vicinity of a transimpedance amplifier (TIA). The TIA and a limiting amplifier (LA) together amplify the received signal to levels that are compatible with the RX CDR circuit. The CDR design is

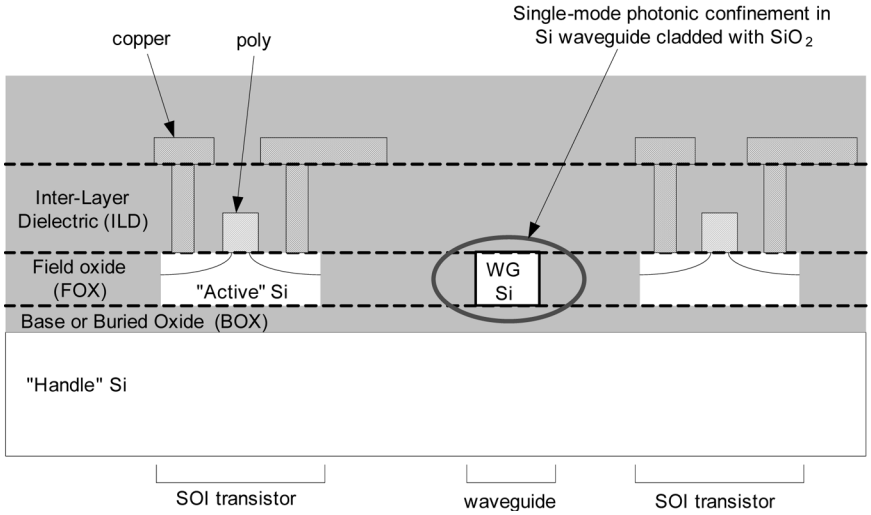

Fig. 3. Conceptual cross section of a CMOS photonics technology (not to scale).

shared between the RX and the TX, and is described in greater detail in Section III.

In addition to chip-to-chip high-speed links, each transceiver supports an on-chip optical loopback function, as shown in Fig. 2. The complementary optical output of the high-speed MZI remains on chip and is routed back to the high-speed PD via an optical waveguide. This feature allows the self-testing of the complete optoelectronic TX/RX chain at the single-chip level. Additionally, the optical loopback demonstrates the feasibility of this technology in the implementation of intra-chip interconnects, which may be useful for applications such as low-skew clock distribution [10], [11]. The loopback feature is enabled by the inclusion of a variable optical attenuator (VOA), which is a part of our silicon photonics library and will be discussed further in Section III. The VOA is enabled during normal operation, suppressing the loopback signal. Disabling the VOA enables the loopback path, as shown in Fig. 2.

The two optoelectronic transceivers share a central bias reference circuit. Additionally, a 16-bit bus controller unit (BCU) is provided for digital communication with an off-chip microcontroller. The BCU enables design for testability (DFT) as a critical function for managing the complexity of an integrated optoelectronic system. Two separate bus architectures on-chip are dedicated to DFT functions.

The simultaneous integration of electronics and photonics in the same active silicon facilitates an unprecedented level of functionality and flexibility in the design of a system-on-a-chip (SoC). For example, process variations in photonics components may be compensated for by using electronic circuits, thus raising the yield of photonic components [9]. The transceiver presented here uses a combination of optical feedback and electronic circuits to automatically set the operating point of the MZI modulator at quadrature, i.e., at the most linear point of its optoelectronic transfer function.

III. System Building Blocks

Fig. 3 shows a cross section of an optical waveguide together with a transistor in the active layer of the CMOS SOI photonics technology. As can be seen from the illustration, optical and electronic devices coexist in the same layer of active silicon. Light propagates on the chip through such waveguides and is manipulated by a combination of passive and/or active optical devices to achieve a variety of system functionalities. Optoelectronic functions such as modulation may be achieved by combining optical waveguides with electrical implants, allowing the optical mode to interact with an imposed electrical field within the semiconductor. The co-existence of electronic and photonic devices in the same silicon layer, created by virtually identical process steps, allows photonic and electronic functions to be integrated seamlessly on the same chip. This section discusses the key optical, optoelectronic, and electronic building blocks used in the development of the 20-Gb/s transceiver described in the rest of the paper.

A. Optical Components

1) Holographic Lens: The ability to couple light in and out of the silicon chip efficiently is essential to the silicon photonics approach. Since a very high index contrast exists between silicon (the waveguide core) and silicon dioxide (the waveguide cladding), optical waveguides in SOI are roughly 50 times smaller than the core of a single-mode fiber. This results in a severe mode mismatch, and consequently a very high insertion loss for simple edge or end-fire coupling. This difficulty is overcome by using holographic lenses, consisting of trenches etched into the active silicon [12], [13]. The trench periodicity is chosen by an optical phase matching condition between the incident light wave and the optically guided wave [13]. The trenches scatter light from a normally incident optical fiber (Fig. 4) into the silicon waveguide. Such a device has a coupling insertion loss of less than 2 dB over a large bandwidth of approximately 3 THz in the optical C-band. The device is reciprocal, allowing light to couple in and out of the chip with equal efficiency. Additionally, the use of normal to surface coupling, as opposed to end-fire coupling, allows for the wafer scale testing of chips and devices using an optical fiber probe as seen in Fig. 5. This is in contrast to conventional optical technologies which require every device to be separately packaged before testing, often at significant cost. The wafer-scale

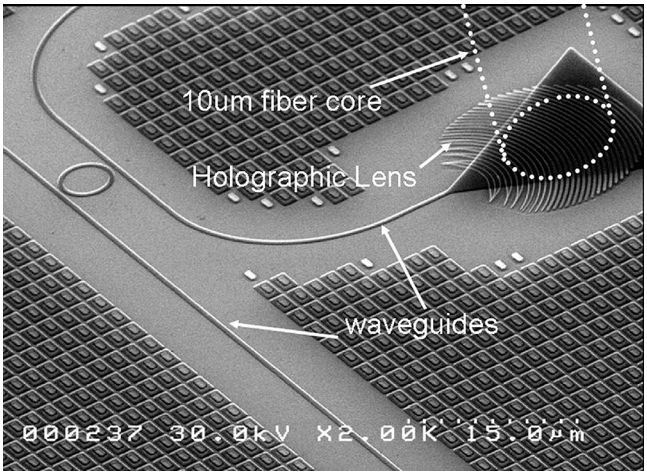

Fig. 4. SEM microphotograph of the holographic lens.

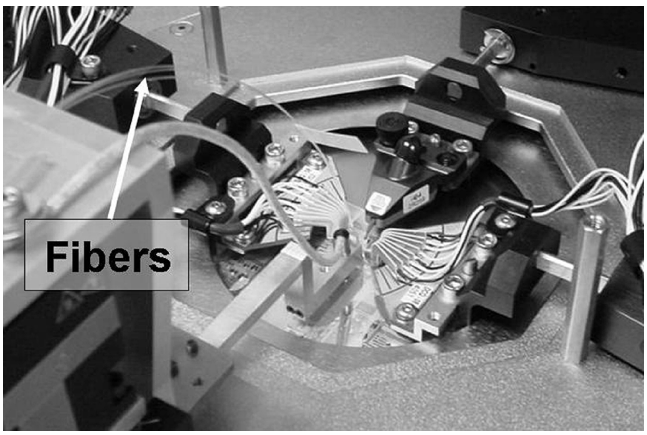

Fig. 5. Wafer probe station with simultaneous photonic and electronic test.

testability afforded by the HL is thus a key advantage of our silicon photonics platform.

2) Optical Waveguide: The optical waveguide is etched into the active top silicon using a planar etch (Fig. 4). Confinement of the optical mode in the silicon is obtained by the index contrast of the active silicon with the underlying buried oxide (BOX) and overlying dielectric layers. The waveguide serves as a low-loss, high-bandwidth interconnect between different optical devices on a single chip. The waveguide is approximately 0.5 μm wide and resides in the same active silicon layer as the CMOS transistors. The high index contrast between Si and SiO_2 allows these waveguides to have bend radii of ~ 30 μm with no additional bending loss, allowing for very compact on-chip routing of optical signals.

B. Optoelectronic Components

The chip demonstrated here uses three distinct types of optoelectronic devices. These are qualitatively compared in Table I, and discussed in more detail below.

1) High-Speed Mach–Zehnder Modulator (Majority Carrier Device): This is the device used to impose high-speed electronic data on the optical carrier wave. A Mach–Zehnder modulator achieves the modulation of optical signals by employing an electro-optic effect [14], [15]. The principle of an interferometer (Fig. 6) is intuitive and simple. Light is split evenly into two arms, and then recombined. An electric field on one of the arms causes a change in carrier density that, in turn, induces a phase shift in the light propagating in this arm. This causes the light to constructively or destructively interfere at the output depending on the field (phase) applied. Accumulation of a phase difference $(\Delta\phi)$ between the two arms causes the recombined light to interfere according to the equation

$$P_{\rm out}/P_{\rm in} = 0.5 + \cos(\Delta\phi)/2. \quad (1)$$

The electro-optic effect used to modulate optical phase in such a device is based on free carrier dispersion [12]. A change in carrier density in the optical waveguide results in a change in the refractive index, thus inducing a change in the optical phase that is proportional to the length of the index modified waveguide section. The key innovation in the phase modulator is operation based on a reverse-biased diode configuration that allows the majority carriers to drift in and out of the optical mode under the influence of electric fields, which is critical for the high-speed capability of this device. This is a much faster mechanism compared to diffusion and recombination to remove minority carriers, as used in modulators based on forward biased p-n diodes, which are generally not suitable for multi-gigabit applications. The core of the electro-optic transducer is thus a reverse-biased lateral p-n diode that is approximately 4 mm long, and the resultant device speed is entirely limited by *RLC* parasitics.

TABLE I
OPTICAL MODULATOR TRADEOFFS

Attribute	Minority Carrier	Majority Carrier	Thermal
Δn / x (index change/length)	High	Low	Moderate
τ (relaxation time)	Moderate (μs)	Fast (fs)	Slow (ms)
IL (insertion loss)	High	Moderate	Low
P_{static} (static power)	Moderate	Low	High
V_{active} (activation voltage)	Low/Medium	High	Low/Medium
$\Delta\varphi$ (Modulation polarity)	+	+	-

2) Thermal Phase Modulator (Thermo-Optic Device): Thermal phase modulators (TPMs) [14] are used as phase tuners in the optoelectronic transceiver system. Such devices are used to set the operating point of the high-speed MZI. From an electrical perspective, a thermal phase modulator is 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 thermo-optic effect in silicon (the Si index changes by approximately 2×10^{-4}/K), causing a phase change in the propagating light wave.

Fig. 7 shows a functional block diagram of the TPM, and its interface with optical and electrical signals. The TPM is approximately 250 μm long and 10 μm wide. These dimensions are optimized to ensure a TPM resistance that allows for maximum possible power dissipation using on-chip drivers. The phase imparted to the optical field by the TPM, $\phi(I_{\rm in})$, is a function of the resistive heating induced by driving the TPM with a current $I_{\rm in}$.

3) Variable Optical Attenuator (Minority Carrier Device): Variable optical attenuators [14] are devices that impart an electrically controlled attenuation to an optical wave. Fig. 8 shows a

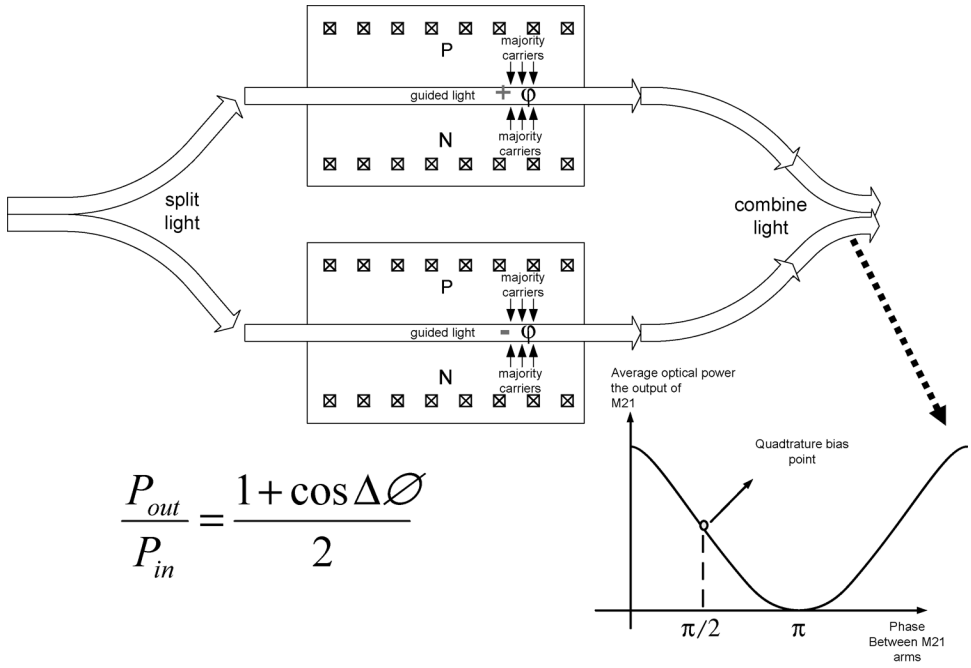

Fig. 6. MZI principle of operation.

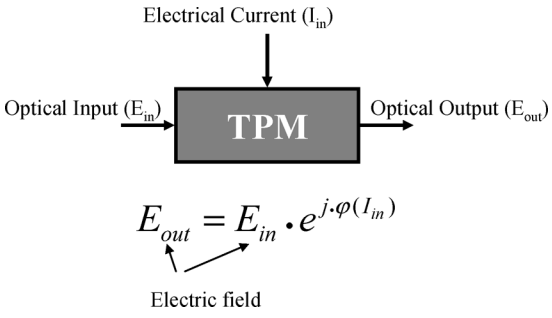

Fig. 7. Thermal phase modulator (TPM) operation.

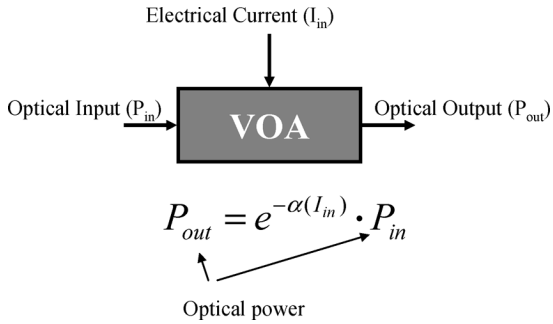

Fig. 8. Variable optical attenuator (VOA) operation.

functional block diagram of the VOA and its interface with optical and electrical signals. The device is approximately 1 mm long and 10 μm wide. It consists of a lateral p-i-n diode integrated into an optical waveguide. Forward-biasing the diode causes carriers to accumulate in the intrinsic region of the p-i-n diode, thus resulting in their overlap with the optical waveguide mode. The variation in carrier density causes a change in the optical absorption coefficient of the waveguide. An optical attenuation of up to 40 dB may be induced with approximately 100 mA of applied current. VOAs may be used to time-division multiplex between many low-speed optical signals simultaneously incident on the same low-speed photodetector.

C. Electronic Components

1) Optical Transmitter: The optical transmitter (Fig. 9) consists of a modulator driver circuit connected to the MZI. Due to the large size of the MZI, the electrical signal is distributed to the individual optical modulating elements using a pair of 50-Ω coplanar transmission lines. Each optical modulating element consists of a reverse-biased p-n junction straddling a section of a waveguide. Each junction presents a lumped capacitive load. To avoid creating impedance discontinuities, a large number of small modulating elements are uniformly distributed along the transmission lines. The dimensions of the transmission lines were optimized to provide a characteristic impedance of 50 Ω when loaded with this distributed junction capacitance. Additionally, care was taken to match the propagation velocity of the electrical signal to the speed of light in the waveguide. Dual termination is used to avoid inter-symbol interference (ISI) in the optical output due to signal reflections. The light is provided to the MZI by a continuous-wave laser whose output beam is coupled into an on-chip waveguide via an HL.

The two arms of the MZI can exhibit a phase offset due to process tolerances. This is compensated for using a control loop consisting of an integrated optical tap which directs 5% of the output light to a monitor PD connected to a 9-bit ADC with a low-impedance input. The ADC uses successive approximation architecture and is designed to operate at 10 Msamples per second. Its power consumption is about 5 mW. The ADC acts as a low-speed TIA and produces a binary word representing the average output power. This information is then used to adjust the currents in thermal phase modulators in each of the MZI arms. The control loop chooses the thermal phase modulator current to

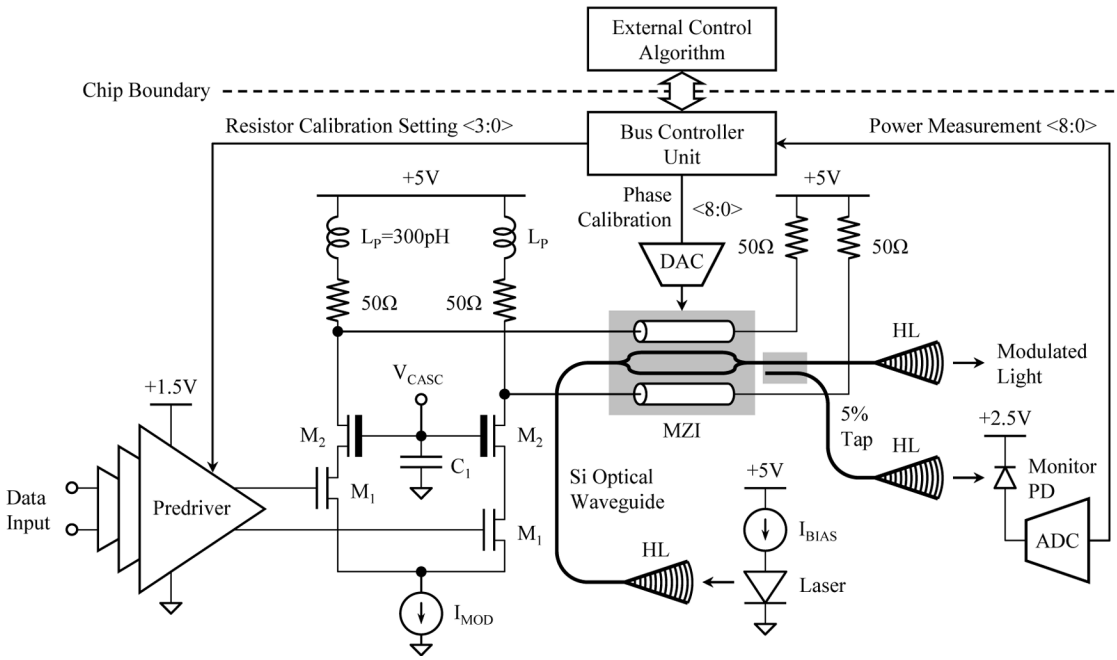

Fig. 9. Optical transmitter schematic.

keep the average output power of the MZI at 3 dB below its maximum output power. This is called the "quadrature point" of the MZI. The device displays its most linear transfer characteristics when modulated around this operating point. When the control loop is operating, the MZI exhibits a much lower sensitivity to process and temperature variations compared to directly-modulated lasers, which typically require automatic modulation control due to significant variations in their slope efficiency. Consequently, constant-swing modulation is sufficient in case of the MZI-based optical transmitter.

As shown in Fig. 9, the modulator driver circuit consists of a three-stage differential predriver operating from 1.5-V supply followed by a 50-Ω output stage, which is connected to the MZI transmission line and uses a 5-V supply. A significant design challenge associated with driving an integrated MZI is related to its high voltage swing requirement. To achieve 10-Gb/s operation with a reasonable extinction ratio at the MZI output, the driver needs to provide a differential swing of at least 5 V peak-to-peak with 20%–80% transition times of about 25 ps. Such large voltage can severely stress the breakdown limits of 0.13-μm CMOS transistors. This challenge is not unique to silicon modulators. Most hybrid solutions utilize InP HBTs or AlGaAs/GaAs devices to avoid the tradeoff between speed and the breakdown voltage [16]–[19]. To enable monolithic integration, a cascoded MZI driver was developed. It consists of a differential pair M_1 implemented with fast-switching thin-oxide 0.13-μm nFETs with a pair of thick-oxide cascode devices M_2, whose purpose is 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. They are available in most advanced digital CMOS technologies, and are intended for use in mixed-voltage I/O interfaces. The output stage is biased with a constant tail current $I_{\mathrm{MOD}} = 100$ mA. Shunt inductive peaking is used in the output stage as well as the predriver to achieve fast transitions [20]. Symmetrically placed gate decoupling capacitance C_1 assures low impedance looking into the source terminals of the cascode transistors, which reduces the voltage swing at the drain terminals of the switching differential pair M_1 and helps in neutralizing its Miller capacitance [21]. The predriver accepts current-mode logic (CML) signals, and utilizes a resistor calibration procedure to desensitize the predriver common-mode output voltage and swing to variations in the process. This assures a consistent switching operation of the output stage. The procedure is executed at power-up and uses a combination of fixed and process-dependent currents, which are generated based on high-precision external components. The fixed current is forced into a bank of resistors controlled by a state machine. These resistors are sequentially switched and their voltage drops are compared to a reference branch conducting a process-dependent current. The result is a 4-bit word which controls switched resistor banks loading the differential predriver stages. Compared to a simple tail current adjustment, this approach yields a constant voltage swing and power dissipation across process corners.

2) Optical Receiver: The optical receiver front-end, shown in Fig. 10, consists of an external high-speed p-i-n PD with a responsivity of 0.9 A/W and a capacitance of 150 fF that is connected to a TIA and followed by a five-stage limiting amplifier (LA). The PD is flip-chip bonded to the transceiver chip and receives incident light via an on-chip waveguide and an HL. The cathode terminal of the PD is biased via a low drop-out (LDO) regulator to provide voltage adjustability and improve isolation from the main power supply rail. Since the PD is flip-chip bonded onto the die, it can be located in close physical proximity to the TIA, minimizing the parasitic inductance of the interface. However, series inductive peaking is a useful technique for extending TIA bandwidth as it helps to offset the effect of

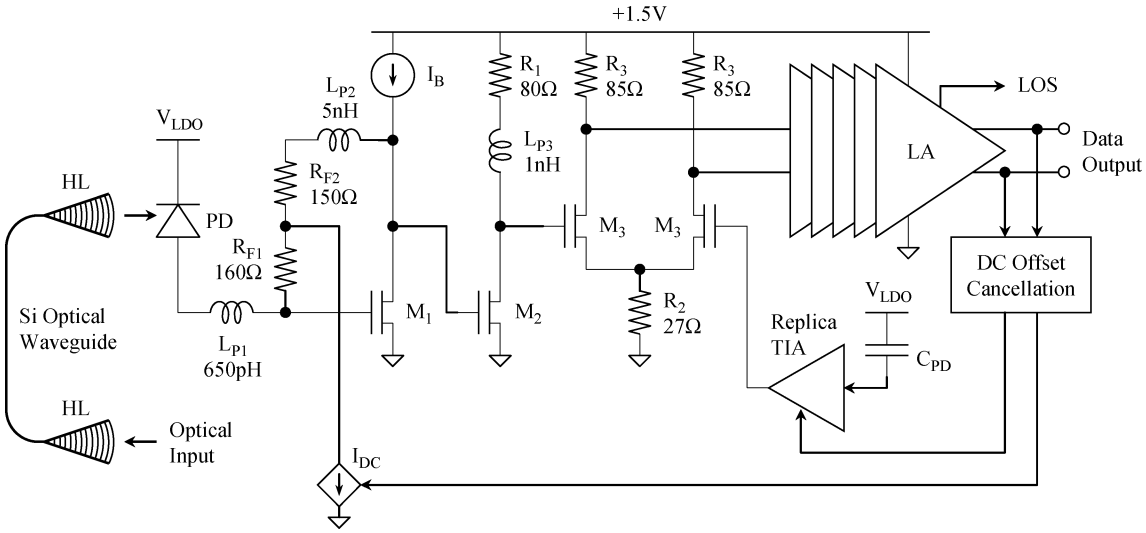

Fig. 10. Optical receiver schematic.

the PD capacitance [22], [23]. A spiral inductor $L_{\rm P1}$ was added to the input terminal to accomplish this goal. The input branch of the TIA comprises a common-source stage M_1 with negative feedback to lower the input impedance [24]. The feedback resistance was split into two elements $R_{\rm F1}$ and $R_{\rm F2}$ to isolate the load introduced by a DC offset cancellation circuit. A spiral inductor $L_{\rm P2}$ was also added in the feedback path to achieve bandwidth extension. This increases the impedance of the feedback path at high frequencies, which introduces a zero in the frequency response and helps to offset the effect of poles that are created by parasitic capacitances including the PD capacitance. It also reduces input-referred noise contribution of the feedback resistors [25]. The first stage also uses an active load $I_{\rm B}$ to enable high gain without sacrificing voltage headroom. It allows biasing M_1 at a higher current level compared to using a simple resistor with the same 1.5-V supply. Undeniably, it is a compromise between gain and noise performance; however, a high-gain stage is desired for a proper operation of a feedback TIA. The additional noise and capacitive loading introduced by the active load were considered in the design process. The first common-source stage M_1 is followed by a second, resistively loaded common-source amplifier with a shunt peaking inductor $L_{\rm P3}$. The third stage is a differential amplifier comprising transistors M_3 and resistors R_3. It converts the single-ended input to a differential signal, which is then sent to the LA. A resistor R_2 is used as a tail current source to reduce noise, and to provide a process stabilized gain. Both R_2 and drain loads R_3 are implemented with the same type of resistors and use similar layout geometries to assure process tracking. The second side of the differential amplifier is connected to a replica of the first two TIA stages. The input node of this reference branch is terminated at the LDO output via a dummy capacitor $C_{\rm PD}$, which matches the junction capacitance of the PD. This pseudo-differential arrangement helps reduce crosstalk sensitivity, since most disturbances are introduced into the circuit as common-mode signals.

The limiting amplifier consists of identical resistively-loaded differential stages with shunt inductive peaking to assure a wide bandwidth, and nFET-based tail current sources to provide good common-mode rejection. It also incorporates a level detector, which flags a loss of signal (LOS) condition. The level detector compares the received signal average power against a locally created reference. The LOS is flagged when the signal amplitude drops below a pre-defined level. The flag is then propagated to the CDR to activate the frequency acquisition loop. The LA provides sufficient gain to convert the minimum expected photocurrent into a CML-compatible voltage swing suitable for driving the CDR. The output of the LA is filtered to extract the DC level information, and the offset cancellation loop subtracts their difference from the inputs of both TIA branches.

3) Common Clock and Data Recovery Circuit: The CDR circuit, used in both the TX and RX, is shown in Fig. 11. It consists of a data recovery loop and a frequency acquisition loop. The data recovery loop uses conventional full-rate architecture with a linear Hogge phase detector, which provides intrinsic data retiming [26]. The proportional and reference pulses of the phase detector are applied to a single-ended charge pump and filtered by a loop filter. The output of the CDR loop filter is connected to a dedicated varactor bank in the voltage-controlled oscillator (VCO). The locking range of the data recovery loop, however, is rather limited, and needs to be augmented with a dedicated frequency acquisition loop comprising a phase-frequency detector (PFD) with a second charge pump, a loop filter, and a dedicated VCO varactor bank. The frequency-acquisition phase-locked loop (PLL) uses a conventional third-order architecture, and all of its components, except for the charge pump, are implemented using 0.13-μm static CMOS logic gates operating from a 1.2-V supply. The PFD uses a topology described in [27]. At startup or if data is temporarily lost, the PFD compares the divided VCO output to a 160-MHz reference clock and steers the VCO towards the expected frequency. The output of the PFD is monitored by a lock detector circuit. The lock detector hands off control of the VCO from the frequency acquisition loop to the data recovery loop if the frequency offset between the VCO and the reference clock becomes less than 500 ppm while data is detected at the input. The two loop filters

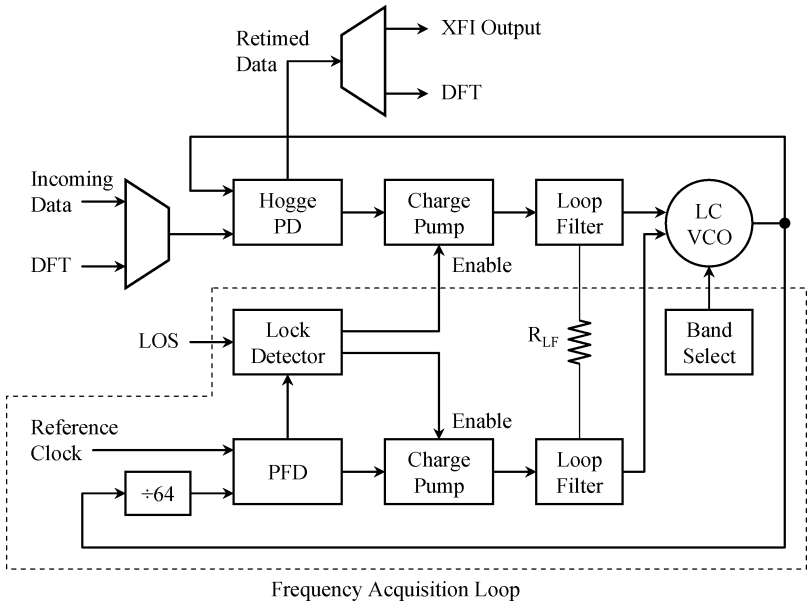

Fig. 11. CDR diagram.

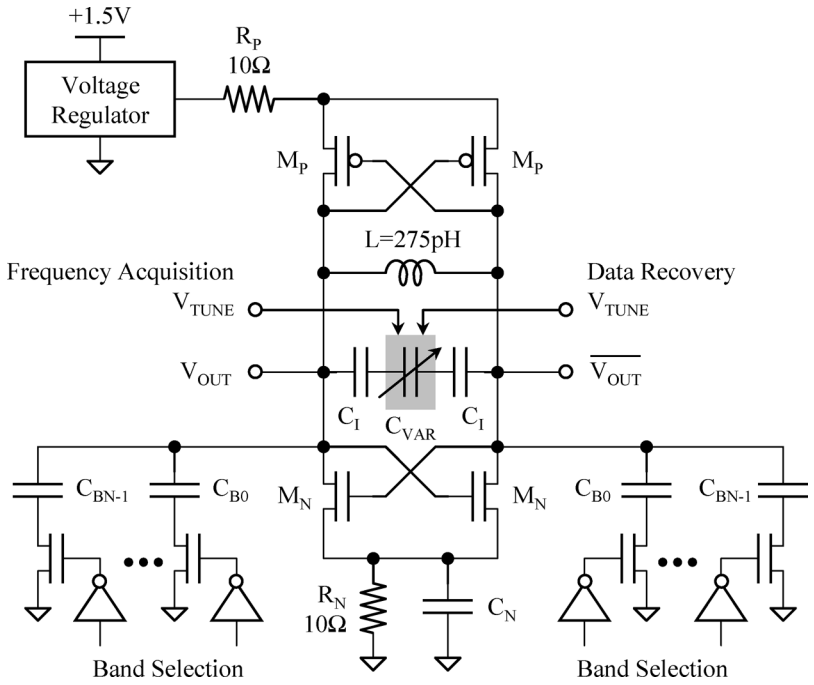

Fig. 12. VCO schematic.

are connected via a large resistor $R_{\rm LF}$, which makes the total tuning range of the VCO available in either mode of operation without altering the dynamic behavior of each loop. The input and output data streams are multiplexed to enable a number of loop-back modes. This is a part of a design for testability (DFT) methodology incorporated into the chip. The data path of the CDR is fully differential and consists of 1.5-V CML gates and buffers. The frequency acquisition loop is comprised of static CMOS gates operating from a 1.2-V supply, which is internally regulated from the main 1.5-V rail.

The CDR uses a single-phase 10-GHz *LC* VCO with complementary cross-coupled transistor pairs, as shown in Fig. 12. This arrangement maximizes trans-conductance and voltage swing at any given bias current, which translates to a reduced power consumption and a lower phase noise [28]. The VCO tank comprises a single spiral inductor L and a varactor bank $C_{\rm VAR}$ connected to the negative resistance cells via DC-isolation capacitors C_I. The inductor measures 70 μm×70 μm and uses a single, top metal layer to achieve a quality factor of about 8.5. The VCO varactor bank is split into two sections. One section is controlled by the frequency acquisition PLL and the other one is controlled by the CDR loop. This allows independent optimization of loop parameters. A resistor $R_{\rm N}$ and a capacitor $C_{\rm N}$ are biasing the source terminals of nFETs $M_{\rm N}$. Unlike an active tail current

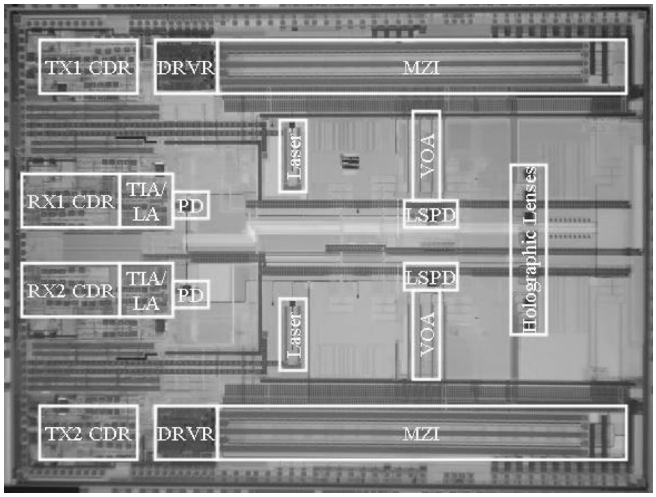

Fig. 13. Chip microphotograph.

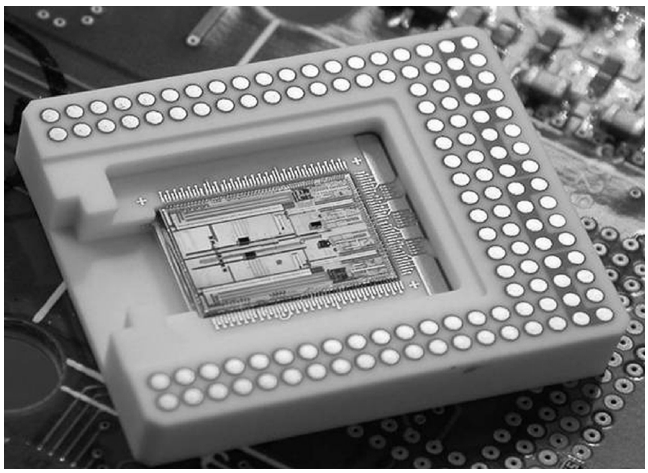

Fig. 14. Packaged dual-channel transceiver system.

source, they do not contribute any 1/f noise [29]. The source terminals of pFETs M_P are biased with a voltage regulator and a series resistor R_P to improve isolation from the main 1.5-V supply rail. A wide tuning range of 2 GHz is achieved without an excessively large tuning sensitivity $(K_{\rm VCO})$ by using banks of digitally controlled, switched fringe capacitors connected to the tank [30]. The switching is accomplished using nFETs operated in triode or cutoff regions, as shown in Fig. 12. Their sizing is a compromise between series resistance in triode and parasitic loading of the *LC* tank in cutoff. The tuning range is centered around 10.5 GHz, and is split into four overlapping sub-bands using the switched capacitors. Each of the sub-bands spans about 700 MHz. This roughly corresponds to $C_{\rm min}$ of 0.7 pF and $C_{\rm max}$ of 1 pF.

IV. Experimental Results

Fig. 13 shows a microphotograph of the system chip with the two parallel 10-Gb/s channels indicated. A photograph of the fully assembled package is shown in Fig. 14. The package is compression mounted to a printed circuit board (PCB) for testing, with a backside heat sink. The die is 8 mm $\times$ 5.6 mm and fits within a 17 mm $\times$ 20 mm $\times$ 4 mm package. The two-channel transceiver supports an aggregate data rate in excess of 20 Gb/s and consumes a total power of only 2.5 W. All high-speed testing was performed using an Agilent N4903A J-BERT and an Agilent 86100 series DCA. Fig. 15 shows an electrical eye obtained at the RX output, demonstrating the receiver's XFI mask compliance. A transmitter output optical eye is shown in Fig. 16. A typical extinction ratio of 5–6 dB is achieved at the output of the TX with rise and fall times of 28 ps.

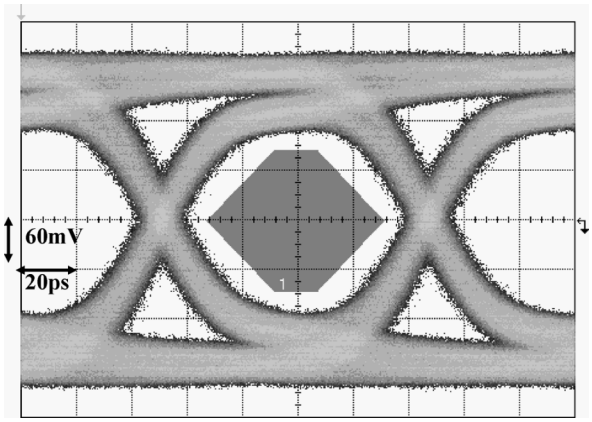

Fig. 15. Single-ended XFI compatible RX output eye diagram. (Input Pattern = PRBS $2^{31}-1$ @ 10.3125 Gb/s; $P_{\rm in} = -10$ dBm @ $\lambda = 1545$nm).

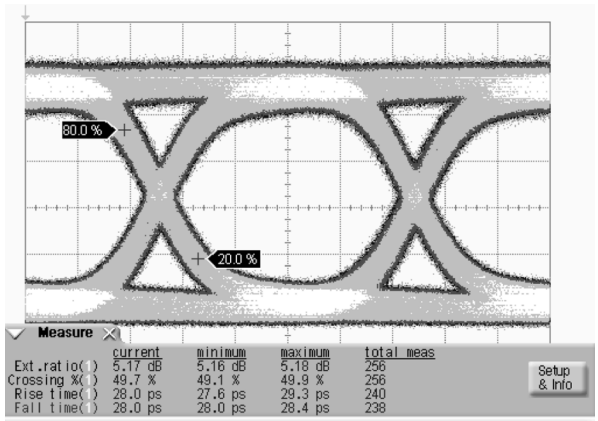

Fig. 16. TX optical output eye diagram at 10.3125 Gb/s, PRBS $2^{31}-1$.

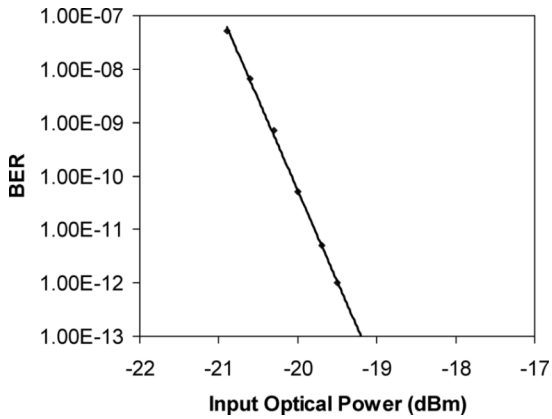

Fig. 17. BER versus optical power at the photodetector.

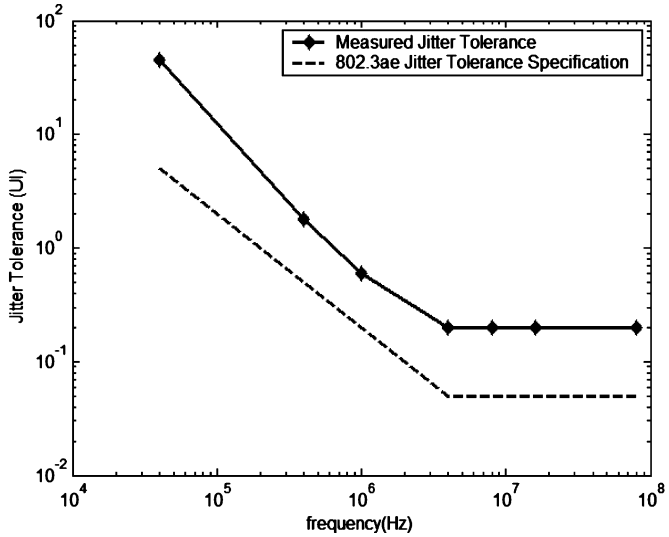

Fig. 18. Receive CDR jitter tolerance.

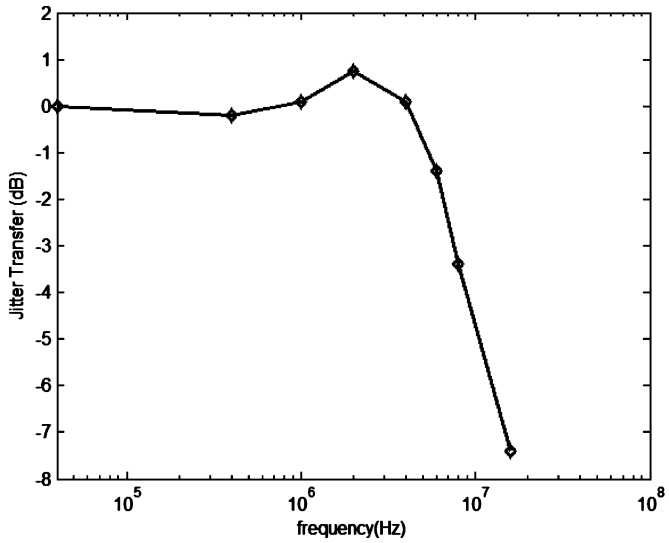

Fig. 19. Receive CDR jitter transfer.

The receiver bit-error rate (BER) versus average optical power incident on the photodetector is illustrated in Fig. 17 for a data rate of 10.3125 Gb/s, using a $2^{31}-1$ PRBS pattern. This demonstrates a receiver sensitivity of -19.5 dBm for a BER of 10^{-12}. The receiver overload level is 0 dBm, resulting in a dynamic range of almost 20 dB. The transmitter exhibits an electrical input sensitivity of 25 mVpp (diff).

The CDR jitter tolerance for the receiver is shown in Fig. 18, relative to the 802.3ae specification. The CDR jitter transfer curve is shown in Fig. 19. The CDR jitter transfer bandwidth is 7 MHz, and there is less than 1 dB of jitter peaking.

To demonstrate complete system functionality, a bidirectional link was constructed in the lab, as illustrated in Fig. 20. The link consisted of two chip assemblies communicating on two adjacent channels through 2000 m of single-mode fiber (SMF). This assembly ran for more than 50 hours without an error before the test was concluded. This result demonstrates the robustness of this system and equates to a BER of better than 10^{-15}. The complete system-level performance is summarized in Table II.

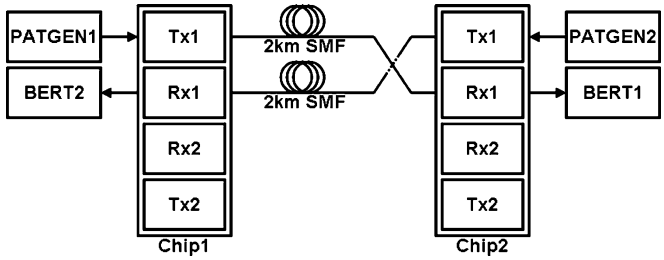

Fig. 20. Bidirectional link experiment setup.

TABLE II
SYSTEM PERFORMANCE SUMMARY

Die Size	**8mm x 5.6mm**
Wavelength	**1535-1555nm**
Link Reach	**>2000m**
Channel Data Rate Range	**9.5-11Gb/s**
Chip Power Consumption (Total for two channels)	**2.5W**

V. CONCLUSION

A complete 20-Gb/s optical transceiver system has been demonstrated, integrating key optical and electrical components on a single substrate using a 0.13-μm CMOS SOI process. Transceiver chips operating at 10 Gb/s have been implemented in CMOS in the past, but never before have they demonstrated such an unprecedented level of optoelectronic integration. The advantages of this integration include reduced size, power consumption, and package integration complexity. To the best of our knowledge, this is the first fully integrated optical transceiver system reported in a silicon technology platform.

ACKNOWLEDGMENT

The authors thank S. Gloeckner, A. Mekis, S. Mirsaidi, and T. Pinguet for their contribution to architecture definition and chip design, E. Balmater, P. De Dobbelaere, M. Harrison, R. Ingram, Y. Liang, M. Peterson, and D. Song for packaging and testing, R. Koumans, G. Masini, V. Sadagopan, and B. Ranjbaran for assistance with chip layout verification, and S. Abdalla, L. Gal, 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 on System-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] H. S. Muthali, T. P. Thomas, and I. A. Young, "A CMOS 10-Gb/s SONET transceiver," *IEEE J. Solid-State Circuits*, vol. 39, no. 7, pp. 1026–1033, Jul. 2004.

[6] H. Werker *et al.*, "A 10-Gb/s SONET-compliant CMOS transceiver with low crosstalk and intrinsic jitter," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2349–2358, Dec. 2004.

[7] L. Henrickson *et al.*, "Low-power fully integrated 10-Gb/s SONET/SDH transceiver in 0.13- μm CMOS," *IEEE J. Solid-State Circuits*, vol. 38, no. 10, pp. 1595–1601, Oct. 2003.

[8] C. Gunn, "CMOS photonics for high-speed interconnects," *IEEE Micro*, vol. 26, no. 2, pp. 58–66, Mar.–Apr. 2006.

[9] 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*, Feb. 2006, pp. 244–245.

[10] A. V. Mule, E. N. Glytsis, T. K. Gaylord, and J. D. Meindl, "Electrical and optical clock distribution networks for gigascale microprocessors," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 10, no. 5, pp. 582–594, Oct. 2002.

[11] M. W. Haney, M. Iqbal, and M. J. McFadden, "Optical interconnects for intrachip global communication: motivation and validation," in *Proc. IEEE LEOS 2005*, Oct. 2005, pp. 206–207.

[12] M. L. Dakss, L. Kuhn, P. F. Heindrich, and B. A. Scott, "Grating coupler for efficient excitation of optical guided waves in thin films," *Appl. Phys. Lett.*, vol. 16, pp. 523–525, Jun. 1970.

[13] D. Taillaert, W. Bogaerts, P. Bienstman, T. F. Krauss, P. Van Daele, I. Moerman, S. Verstuyft, K. De Mesel, and R. Baets, "An out-of-plane grating coupler for efficient butt coupling between compact planar waveguides and single-mode fibers," *IEEE J. Quantum Electron.*, vol. 38, no. 7, pp. 949–955, Jul. 2002.

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

[15] R. A. Soref and B. R. Bennett, "Electrooptical effects in silicon," *IEEE J. Quantum Electron.*, vol. 23, no. 1, pp. 123–129, Jan. 1987.

[16] T. Y. K. Wong, A. P. Freundorfer, B. C. Beggs, and J. E. Sitch, "A 10 Gb/s AlGaAs/GaAs HBT high power fully differential limiting distributed amplifier for III-V Mach-Zehnder modulator," *IEEE J. Solid-State Circuits*, vol. 21, no. 10, pp. 1388–1393, Oct. 1996.

[17] Z. Lao *et al.*, "40-Gb/s high-power modulator driver IC for lightwave communication systems," *IEEE J. Solid-State Circuits*, vol. 33, no. 10, pp. 1520–1526, Oct. 1998.

[18] A. Long, J. Buck, and R. Powell, "Design of an opto-electronic modulator driver amplifier for 40-Gb/s data rate systems," *J. Lightw. Technol.*, vol. 20, no. 12, pp. 2015–2021, Dec. 2002.

[19] Y. Baeyens *et al.*, "High gain-bandwidth differential distributed InP D-HBT driver amplifiers with large (11.3 V_{pp}) output swing at 40 Gb/s," *IEEE J. Solid-State Circuits*, vol. 39, no. 10, pp. 1697–1705, Oct. 2004.

[20] B. Razavi*, Design of Integrated Circuits for Optical Communications*. New York: McGraw-Hill, 2003.

[21] P. R. Gray and R. G. Meyer*, Analysis and Design of Analog Integrated Circuits*, 3rd ed. New York: Wiley, 1993.

[22] M. Neuhauser, H.-M. Rein, H. Wernz, and A. Felder, "13 GB/s Si bipolar preamplifier for optical front ends," *Electron. Lett.*, vol. 29, no. 5, pp. 492–493, Mar. 1993.

[23] B. Analui and A. Hajimiri, "Bandwidth enhancement for transimpedance amplifiers," *IEEE J. Solid-State Circuits*, vol. 39, no. 8, pp. 1263–1270, Aug. 2004.

[24] H. H. Kim, S. Chandrasekhar, C. A. Burrus, Jr., and J. Bauman, "A Si BiCMOS transimpedance amplifier for 10 Gb/s SONET receiver," *IEEE J. Solid-State Circuits*, vol. 36, no. 5, pp. 769–776, May 2001.

[25] H. Tran, F. Pera, D. S. McPherson, D. Viorel, and S. P. Voinigescu, "6-kΩ, 43-Gb/s differential transimpedance-limiting amplifier with auto-zero feedback and high dynamic range," in *IEEE GaAs IC Symp. Tech. Dig.*, Nov. 2003, pp. 241–244.

[26] C. R. Hogge, "A self-correcting clock recovery circuit," *IEEE J. Lightw. Technol.*, vol. 3, no. 6, pp. 1312–1314, Dec. 1985.

[27] N. H. E. Weste and K. Eshragrian*, Principles of CMOS VLSI Design*, 2nd ed. Reading, MA: Addison Wesley, 1993.

[28] A. Jerng and C. G. Sodini, "The impact of device type and sizing on phase noise mechanisms," *IEEE J. Solid-State Circuits*, vol. 40, no. 2, pp. 360–369, Feb. 2005.

[29] S. Levantino *et al.*, "Frequency dependence on bias current in 5-GHz CMOS VCOs: Impact on tuning range and flicker noise upconversion," *IEEE J. Solid-State Circuits*, vol. 37, no. 8, pp. 1003–1011, Aug. 2002.

[30] A. Kral, F. Behbahani, and A. A. Abidi, "RF-CMOS oscillators with switched tuning," in *Proc. IEEE Custom Integrated Circuits Conf.*, May 1998, pp. 555–558.

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 at the IBM T. J. Watson Research Center, Yorktown Heights, NY, in the summer of 2003, where he worked on 10Gb/s multi-mode fiber links and developed an eye-opening monitor for adaptive equalization of serial links. He is now with Luxtera Inc., Carlsbad, CA. His technical interests include signal integrity issues in high-speed serial links 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.

Drew Guckenberger (S'00–M'05) received the B.A.Sc. degree from the University of Waterloo, Waterloo, ON, Canada, in systems design engineering 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, where he worked mainly on the development of high-speed optical receivers and receiver arrays. He is currently a Staff Engineer at Luxtera, Inc., Carlsbad, CA, designing high-speed circuits and optoelectronic systems using silicon photonics. 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 and was a member of the team that placed third in the 2002–2003 Semiconductor Research Council International SiGe Design Contest.

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 IBM Thomas 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.

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 M.S. and Ph.D. degrees in electrical engineering from the University of California at Los Angeles in 2001 and 2004, respectively. He spent most of his time in graduate school investigating high spectral efficiency fiber optic transmission systems and the design and fabrication of diffractive optical devices for efficient coupling into SOI waveguides.

Since 2004, he has been with Luxtera Inc., Carlsbad, CA, where he is a photonic designer and system architect for Luxtera's WDM transceivers.