An Optical Interconnect Transceiver at 1550 nm Using Low-Voltage Electroabsorption Modulators Directly Integrated to CMOS

Jonathan E. Roth, *Member, IEEE*, Samuel Palermo, *Member, IEEE*, Noah C. Helman, David P. Bour, *Fellow, IEEE*, David A. B. Miller, *Fellow, IEEE*, and Mark Horowitz, *Fellow, IEEE*

***Abstract*—A low-voltage, 90-nm CMOS optical interconnect transceiver operating at 1550-nm optical wavelength is presented. This is the first demonstration of a novel optoelectronic modulator architecture (the quasi-waveguide angled-facet electroabsorption modulator) in a system. It features a simple electronic packaging via flip-chip bonding to silicon. Devices have a broad optical bandwidth, are arrayed two dimensionally, and feature surface normal, spatially separated, and misalignment-tolerant optical ports. The modulators are driven with a novel pulsed-cascode driver capable of supplying an output-voltage swing of 2 V (twice the nominal 1-V CMOS supply) without overstressing thin-oxide core CMOS devices. At the receiver side, a sensitivity of −15.2 dBm is obtained with an integrating/double-sampling front end. The transceiver includes clock generation and recovery circuitry that enables a data serialization factor of five. At a maximum data rate of 1.8 Gb/s, the optical transmitter, receiver, and clocking circuitry consume 12.6, 4.5, and 6.5 mW, respectively, for a total link electrical power dissipation of 23.6 mW. To the best of our knowledge, this is the first demonstration of an interconnect transceiver operating at 1550 nm with a III–V output device directly integrated to the CMOS.**

***Index Terms*—Electrooptic transducers, flip-chip devices, integrated optoelectronics, optical interconnection.**

I. Introduction

OPTICAL interconnects have distinct advantages over electrical interconnects, which can be exploited across the range of scales from global telecommunications to buses and clock distribution within computers. The advantages of optical interconnects include low attenuation, dispersion, and jitter, no crosstalk, and simple impedance matching [1].

Manuscript received February 27, 2007; revised August 29, 2007. This work was supported in part by the Defense Advanced Research Projects Agency (DARPA) under a Grant from the GIT MARCO Focus Research Center for Interconnections for Gigascale Integration.

J. E. Roth, D. A. B. Miller, and M. Horowitz are with the Electrical Engineering Department, Stanford University, Stanford, CA 94305 USA (e-mail: jonroth@stanfordalumni.org; dabm@ee.stanford.edu; horowitz@stanford.edu).

S. Palermo was with the Electrical Engineering Department, Stanford University, Stanford, CA 94305 USA. He is now with Intel Corporation, Hillsboro, OR 97124 USA (e-mail: sam.palermo@intel.com).

N. C. Helman was with the Applied Physics Department, Stanford University, Stanford, CA 94305 USA. He is now with the University of California, San Francisco, San Francisco, CA 94158 USA (e-mail: nhelman@cmp.ucsf.edu).

D. P. Bour was with Agilent Technologies, Palo Alto, CA 94304 USA. He is now with BridgeLux, Inc., Sunnyvale, CA 94089 USA (e-mail: davebour@bridgelux.com).

Digital Object Identifier 10.1109/JLT.2007.909334

Optoelectronic modulators using the quantum-confined Stark effect [2] are effective transmitter devices for optical interconnects due to their potential for high-frequency operation and low power dissipation. Most of these devices fall into the category of surface normal or waveguide devices. The surface normal devices typically require a thick multiple-quantum-well (MQW) region to get an adequate contrast and, thus, require a large operating voltage to achieve the necessary electric field for switching. It is possible to reduce the MQW thickness and the operating voltage by containing the MQW region in an asymmetric Fabry–Pérot resonator, but the resulting design is constrained by a narrow wavelength band of operation and a strong temperature dependence. While the waveguide devices do not have these problems, packaging is difficult due to the need to couple to small-area modes, and waveguides can only be arrayed in one dimension of the wafer surface.

This paper demonstrates an optical interconnect transceiver at 1550 nm using a modulator architecture that combines the benefits of both the surface normal and waveguide modulators—the quasi-waveguide angled-facet electroabsorption modulator (QWAFEM). These devices have previously been demonstrated to operate over a wavelength range of 16 nm [3]. They allow for surface-normal access to spatially separated input and output ports and are tolerant to small misalignments of the beam. They have a low drive voltage of 2 V and can directly be flip-chip-bonded to CMOS without a high-speed electrical packaging.

In order to explore the use of these modulators in high-density chip-to-chip optical interconnect applications, 2-D modulator arrays are directly flip-chip-bonded to a CMOS-transceiver chip, thus eliminating the need for high-speed electrical packaging. The transceiver is fabricated in a 90-nm CMOS process and employs a novel pulsed-cascode modulator driver [4] that is capable of supplying an output-voltage swing of 2 V (twice the nominal 1-V supply) without overstressing thin-oxide core CMOS devices. Completing the optical link is a low-voltage integrating and double-sampling receiver front end [5] that eliminates the requirement of a high-bandwidth transimpedance amplifier (TIA).

In this paper, we, to the best of our knowledge, present and discuss the first demonstration of an optical interconnect transceiver operating at 1550 nm with a III–V output device directly integrated to the CMOS. Section II describes the design and operation of the QWAFEM devices. The low-voltage CMOS transceiver is outlined in Section III. In Section IV, we detail the

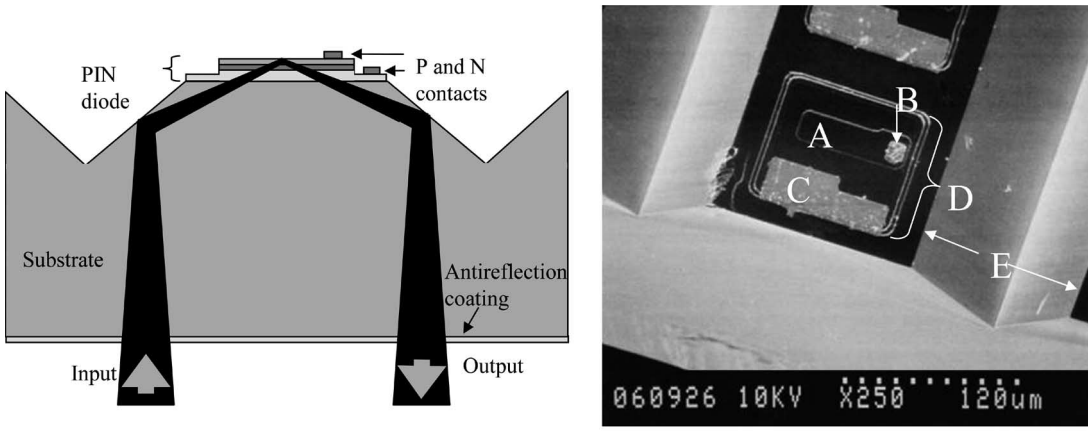

(a) (b)

Fig. 1. (a) Cross-sectional schematic of QWAFEM before bonding to CMOS (not to scale). (b) SEM image of QWAFEM. A: Diode mesa, B: p-contact, C: n-contact, D: Electrically isolated n-doped region, and E: Selectively etched mirror.

experimental setup and results, and in Section V, we summarize this paper.

II. Modulators

The QWAFEM-device architecture is shown in Fig. 1 [3]. The modulators are fabricated on a double-side polished (100) InP wafer, upon which InGaAs/InP epitaxial layers have been grown via metal–organic chemical vapor deposition. The growth consists of a p-doped-intrinsic n-doped (PIN) diode, containing an MQW structure in the intrinsic region and a three-period InGaAsP/InP distributed Bragg reflector (DBR) in the n-doped region. Mesas are etched for the diodes, and metal p- and n-contacts are deposited in an evaporator. On two opposite sides of each mesa, mirrors are selectively etched in the substrate to reveal the {111} planes. Light enters the QWAFEM at normal incidence through the antireflection-coated back surface of the InP substrate. The light undergoes three total internal reflections (two from the selectively etched mirrors and one reflection from the interface between the epitaxially grown InGaAsP and air) before exiting along a path antiparallel to, and displaced from, the incident beam. A resonant cavity is formed around the MQW region by the DBR and the epitaxy–air interface, enhancing absorption. Oblique incidence in the MQW region increases the interaction length of the light with the quantum wells and enhances the performance of the resonator.

Unlike in a waveguide modulator where propagation along the absorbing material is accomplished by coupling to a waveguide mode, there is no modal constraint imposed in coupling through the active region at grazing incidence, easing the optical alignment constraints. Furthermore, the triple-bounce geometry results in a fixed separation between the input and output beams for small translations of the input beam across the modulator substrate surface. In this implementation, the displacement between the input and output ports facilitated the separation of the output beam for detection using a pick-off mirror (POM), which can provide a fourfold improvement in insertion loss compared with using a power beamsplitter (BS) for this purpose. In addition, multiple modulators on a chip could be tested by translating the chip without realigning the detection optics. For future parallel link implementations, it should be possible to align an array of lensed fibers for the input and output couplings to the modulator chip in a single alignment step, provided that the pitches of the fibers and the modulators are matched.

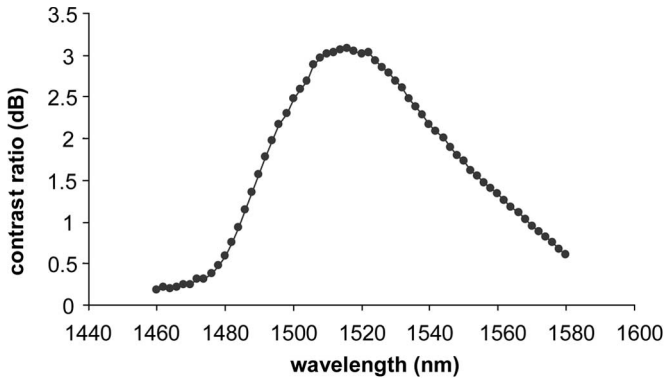

Fig. 2. Static contrast-ratio curve (ratio of output optical power at 0.5-V forward bias divided by the output power at 1.5-V reverse bias) for a typical QWAFEM device (not bonded to CMOS).

Building upon previous work [3], in this implementation of the QWAFEM, the spacing of modulators was changed to match the pitch of the CMOS-transceiver chip, and the diode mesas were resized to reduce capacitance. The current devices range from 20×60 to 40×90 μm in the diode area, corresponding to the designed capacitances ranging from 700 fF to 2.5 pF. While reducing the device size reduces the capacitance, it will also reduce the tolerance to misalignments, and in the limit of small devices, it will reduce the maximum contrast ratio. To avoid stringent growth thickness calibration, three epitaxial wafers were grown with different resonator lengths. Each resonator had two sacrificial layers, of which none, one, or both could selectively be etched to optimize the resonator length. The optimal combination of the wafer and the number of sacrificial layers etched was chosen after experimental comparison of nine fabricated device arrays. A static curve of the contrast ratio versus the wavelength of a typical such QWAFEM is shown in Fig. 2.

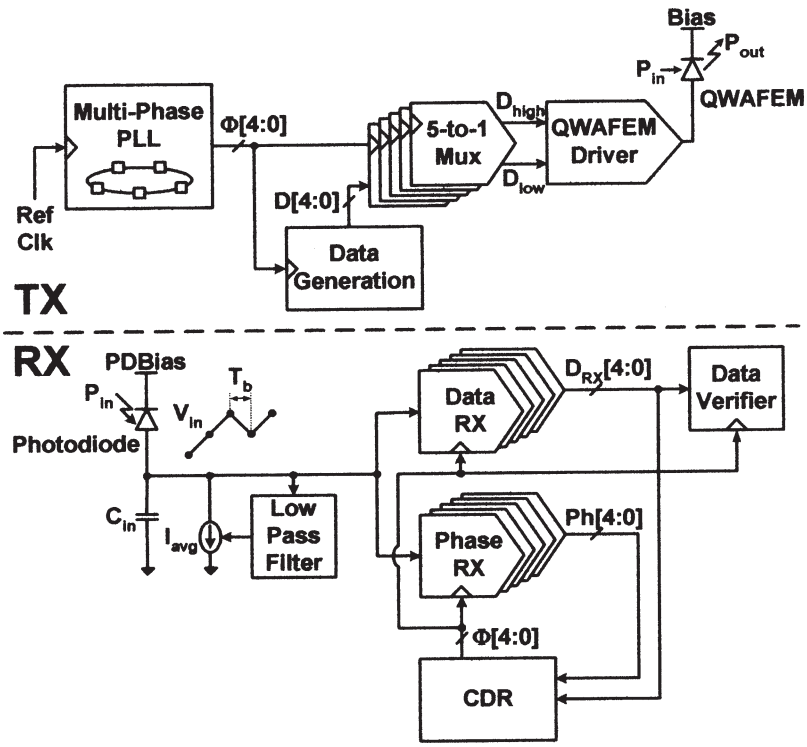

Fig. 3. Optical transceiver architecture.

III. CMOS Transceiver

A. Transceiver Architecture

The optical interconnect transceiver architecture is shown in Fig. 3. In order to enable short bit periods without consuming excessive area and power in clock generation and distribution, a multiple clock-phase multiplexing architecture is used at both the transmitter and the receiver. In the transmitter frequency synthesis phase-locked loop (PLL), a five-stage ring oscillator provides five sets of complementary clock phases that are spaced a bit period apart. These phases are used to switch a level-shifting multiplexer to produce a serial data stream with a data rate of five times the clock frequency. The multiplexer serial output is then buffered by the modulator-driver output stage [4]. At the receiver side, the input photocurrent is integrated onto the input-node capacitance, and a double-sampling technique is used to resolve the data bits [5], [6]. A demultiplexing factor of five is directly achieved at the input node using five uniform clock phases from the clock and data recovery (CDR) system.

B. Modulator Driver

For modern CMOS technologies, an output swing greater than the nominal power supply is required in order to provide an appropriate contrast ratio with integrated surface normal EMs. This conflicts with the CMOS reliability considerations [7], [8] which constrain the maximum static voltages across a core transistor's gate, source, and drain terminals to be no more than the nominal power supply, whereas the transient voltage spikes must not exceed this limit by more than 20%–30%. Thick-oxide I/O devices that are rated for higher voltage operation could potentially be used to supply the necessary modulator drive voltages, but these thick-oxide devices cannot match the core CMOS devices' speed. Thus, the challenge is to provide at high data rates an acceptable output swing without overstressing the core devices. To address this, a pulsed-cascode output stage is used that reliably supplies a voltage swing of twice the nominal supply and consists of only core devices for maximum switching speed.

Fig. 4 shows the pulsed-cascode output stage which accepts both a "low" input $\mathrm{IN_{low}}$ that swings between the Gnd and the nominal chip Vdd and a "high" input $\mathrm{IN_{high}}$ with the same data value that has been level-shifted to swing between Vdd and Vdd2, where Vdd2 is nominally twice the voltage of Vdd. The level-shifting multiplexer circuitry is detailed in [4]. Static-voltage overstress is eliminated in the output-stage cascode structure by equally splitting the output voltage across the series transistors. Pulsing the gates of the cascode transistors (MN2 and MP2) during transitions with NAND- and NOR-pulse gates, respectively, allows this driver to eliminate the transient drain–source voltage (V_{ds}) overstress present in static-biased cascode drivers [9] and prevents transistor degradation from hot-carrier injection [10].

Fig. 5 shows the simulation waveforms of the pulsed-cascode modulator driver with a nominal CMOS supply of 1 V, providing a 2-V output transition from high to low with an assumed modulator capacitance load of 1 pF. A falling transition from the "low" input switches the bottom nMOS (MN1) to drive node $\mathrm{mid_n}$ to Gnd, and a simultaneous falling transition on the "high" input triggers a positive pulse from the NOR-pulse gate that drives the gate of MN2 from Vdd to near Vdd2 to allow the output to begin discharging at roughly the same time that the MN2 source is being discharged [Fig. 5(a) and (b)].

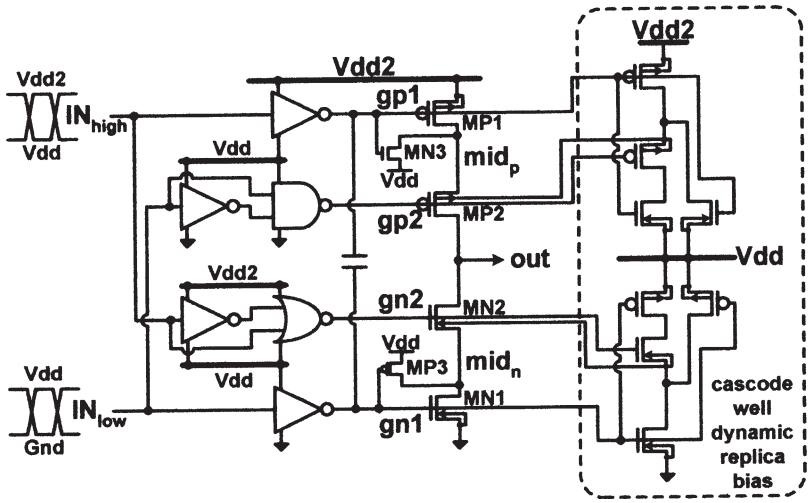

Fig. 4. Pulsed-cascode output stage.

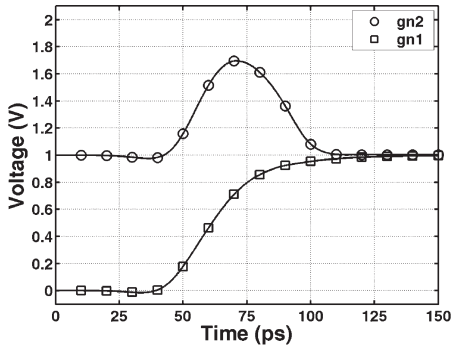

(a)

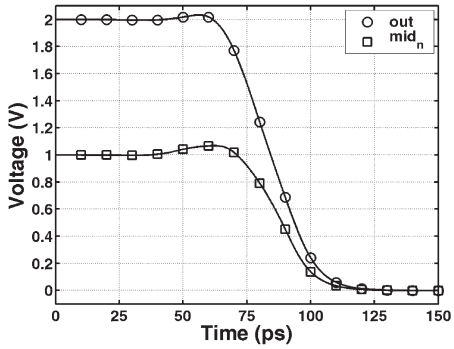

(b)

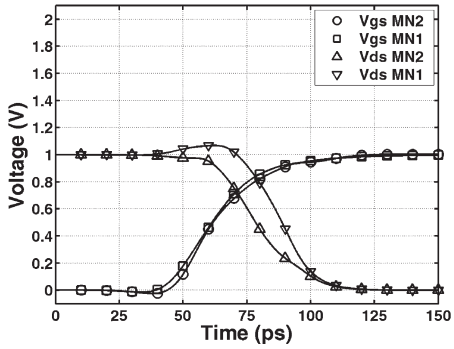

(c)

Fig. 5. Transient simulation of the pulsed-cascode output stage loaded with a 1-pF modulator for the case of a falling transition. (a) nMOS gate voltages, (b) nMOS drain voltages, and (c) nMOS $V_{\rm gs}$ and $V_{\rm ds}$.

Thus, the cascode nMOS drain–source voltage does not overly exceed the nominal supply voltage [Fig. 5(c)]. The NOR-pulse gate is sized such that the gate of MN2 does not swing all the way to Vdd2 and that the edge rate of the pulse signal also matches the falling rate of $\text{mid}_{\rm n}$. Therefore, during the transition, a gate–source voltage that does not overly exceed the nominal supply is developed across MN2. The "high" input also activates a pull-down nMOS (MN3) to drive node $\text{mid}_{\rm p}$ from Vdd2 to Vdd to prevent excessive $V_{\rm ds}$ stress on MP2. Similarly, during an output transition from low to high, the "high" input switches the top pMOS (MP1) to drive node $\text{mid}_{\rm p}$ to Vdd2, and the "low" input triggers a negative pulse from the NAND-pulse gate that drives the gate of MP2 transistor from Vdd to near Gnd. For ratios of $C_{\rm out}/C_{\rm midn}$ from 1.3 (unloaded) to 15.5, no voltage spikes between the gate, source, and drain terminals of any output devices exceed more than 20% above the supply voltage.

It is important that the cascode transistors have a similar drive strength as the top or bottom transistors to reduce the $V_{\rm ds}$ stress during transients. Thus, in order to minimize the body voltage effect on the cascode transistors, they are placed in separate wells that are dynamically biased with replica circuitry to track their source voltages. This reduces the cascode transistors' threshold voltages, resulting in a similar voltage drop across the two series driving transistors. The increased drive strength of the cascode transistor also serves to reduce the modulator driver's output transition time. Little power and area overhead is necessary for the replica-bias circuitry, as the replica transistors are sized to be less than 10% of the main driver transistors.

C. Integrating and Double-Sampling Receiver

While receiver circuitry power and area may not be a primary issue for traditional telecom applications which demand high sensitivity, in high-density optical interconnect applications, performance parameters such as sensitivity must be balanced with power and area constraints. A receiver front-end architecture that reduces the number of linear gain elements, and

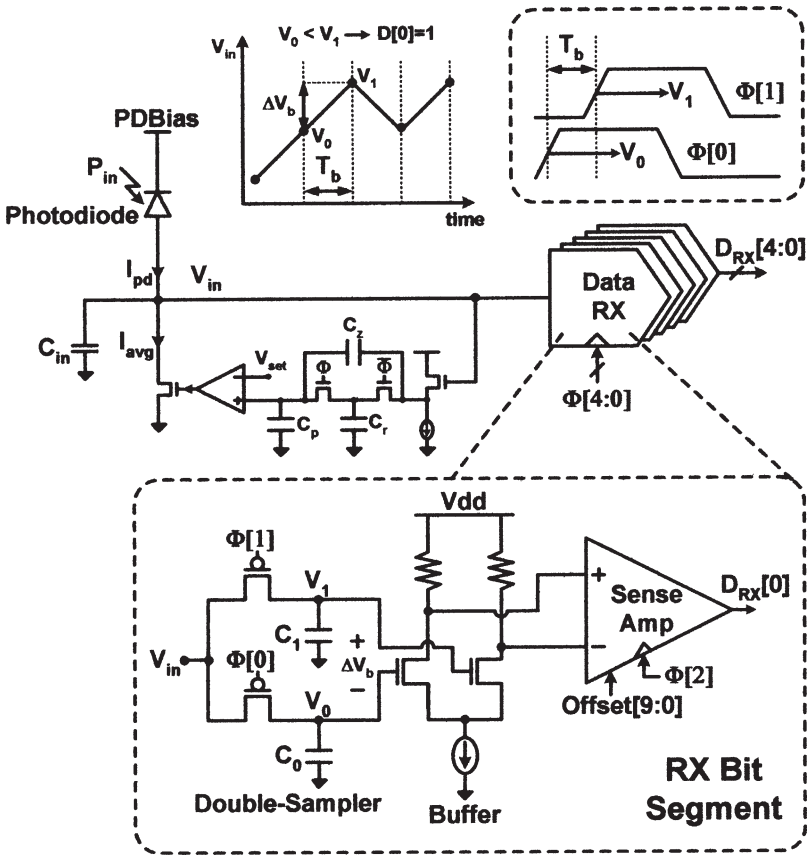

Fig. 6. Integrating and double-sampling receiver front end.

thus is less sensitive to the reduced gain in modern CMOS processes, is the integrating and double-sampling front end [6]. An absence of high-gain amplifiers allows for savings in both the power and the area and makes the integrating and double-sampling architecture more suitable for the chip-to-chip optical interconnect applications.

The integrating and double-sampling receiver front end [5], as shown in Fig. 6, demultiplexes the incoming data stream with five parallel segments that include a pair of input samplers, a buffer, and a sense amplifier. Two current sources at the receiver input node, namely, a photodiode current and a current source that is feedback-biased to the average photodiode current, supply and deplete charge from the receiver input capacitance, respectively. For data encoded to ensure dc balance, the input voltage will integrate up or down due to the mismatch in these currents. A differential voltage $\Delta v_{\rm b}$ is developed in each receiver segment by sampling at the beginning and end of a bit period defined by the rising edge of the recovered clocks $\Phi[n]$ and $\Phi[n+1]$, respectively. While, in a previous implementation [8], $\Delta v_{\rm b}$ was directly applied to an offset-corrected StrongArm latch [11] used as a sense amplifier for data regeneration, the reduced supply voltage that comes with scaling technologies causes the integrating input to exceed the sense-amp input range. In order to fix the sense-amp common-mode input level and to buffer the sensitive sample nodes from kickback charge, a differential buffer is inserted between the samplers and the sense amp. The power penalty of the additional buffer is quite small (250-μW per segment), as the buffer gain is low to avoid sense-amp offset saturation, and the bandwidth requirements are relaxed due to input demultiplexing. The use of pMOS samplers provides a receiver input range from 0.6 to 1.1 V. Demultiplexing directly at the input allows the sense-amp sufficient time (five times the bit period) for data regeneration and precharging, thus eliminating the requirement for a TIA operating at the bit rate.

IV. Experiment

Three experimental configurations were used, as shown in Fig. 7. In the first configuration, laser light was free-space-coupled onto the modulators, and the output was collimated and coupled onto a large-area photodetector for dc contrast-ratio measurements. In the second configuration, light exiting the modulators was coupled into a fiber for transition speed measurements with a high-speed oscilloscope. In the final configuration, light was coupled into high-speed detectors (HSDs) to complete the transceiver link.

Common to all of the configurations, an array of InP QWAFEMs was flip-chip-bonded to the CMOS-transceiver chip with eight transmit channels sized for varying drive strengths and two receive channels, as shown in Fig. 8(a). The transceiver was fabricated in a 1-V 90-nm CMOS process.

The chip was placed in an open-cavity surface-mount package on a test board mounted on a three-axis translation stage. An HP8133A pulse generator supplied the reference clock

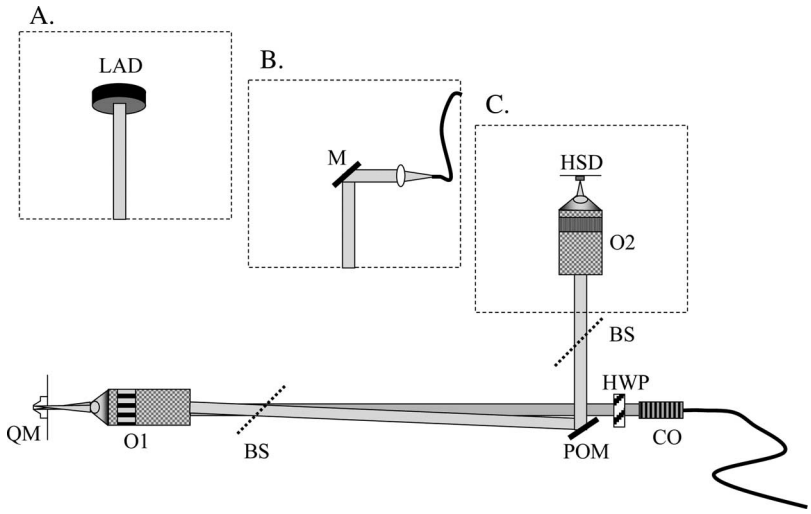

Fig. 7. Pulsed-transceiver link schematic showing three configurations A–C for different measurements. Linear-polarized light enters through the fiber and free-space collimator (CO) and the rotatable half-wave plate (HWP). Focusing onto the QWAFEM (QM in the figure) is accomplished by a microscope objective (O1), and the spatially displaced output beam of the modulator is reflected by the pick-off mirror (POM). In A, for the dc contrast-ratio measurements, collimated light is absorbed by a large-area detector (LAD). In B, for the high-speed rise- and fall-time measurements, the light is reflected off a second mirror and focused into a single-mode fiber. In C, for the full transceiver link, the light is focused by a second microscope objective (O2) onto a high-speed detector (HSD). Alignment of the beam on the modulator and the detector is accomplished with an IR camera (not shown), an LED for illumination, and two removable pellicle beamsplitters (BS).

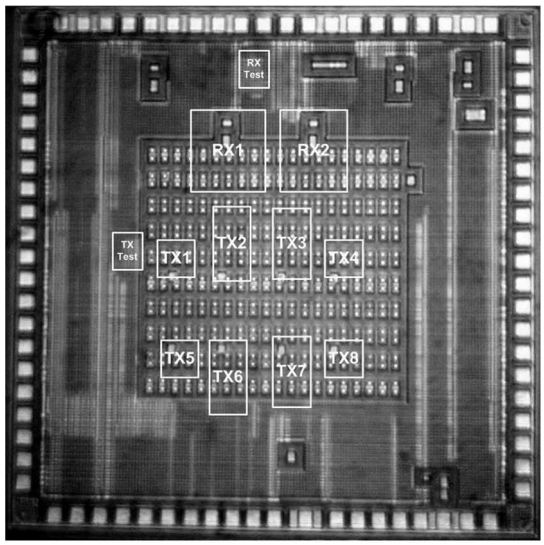

(a)

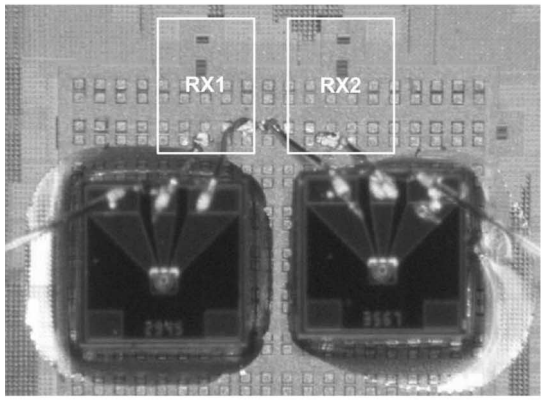

(b)

Fig. 8. (a) Die micrograph of CMOS transceiver. (b) 1550-nm photodiodes wirebonded to the receivers.

to the transmitter PLLs, and the transmit data sequence was controlled with an on-chip 20-b register that can be programmed with a computer via a serial testing interface.

Light from an Agilent 81680A tunable laser with a range of 1457–1584 nm was coupled via polarization-maintaining fiber into a free-space collimator. The collimator was followed by a rotatable half-wave plate, which is used to ensure that the linear-polarized light in the QWAFEM's resonator would be transverse electric (i.e., in the plane of the quantum wells) for optimal performance. The collimated beam was focused onto the modulator array with a Mitutoyo infinity-corrected $10\times$ near-infrared (NIR) objective with a free-space focal spot diameter of about 12 μm. Between the collimator and the objective, a removable pellicle BS was used to allow imaging of the beam on the modulator array to aid the beam alignment. The light entry and exit points on the array's substrate were displaced by 200 μm. After collimation by the microscope objective, the beam exiting the modulator was separated for detection by a POM. A photodetector was placed in the beam path for the dc contrast-ratio measurements. The default high-speed modulation of the devices was bypassed by setting each bit in the 20-b sequence to the same state, and the modulators were switched by changing the bias voltage applied to the p-contacts of all driven modulators. For each working device, the optimal combination of the bias voltage and the wavelength was chosen to maximize the contrast ratio.

The maximum contrast ratio measured on a device bonded to CMOS was 3.86 dB and measured at 1528 nm for a 2-V swing, which is somewhat exceeding the performance of the unbonded device in Fig. 2. Upon coupling the beam into a single-mode fiber, the same device yielded a peak contrast ratio of 5.53 dB for the same conditions. Modulation of the beam in this device could also result in change of shape of the beam because different angular components in the beam would differently interact with the resonator. Any such change of shape effectively

corresponds to coupling light into higher order modes. Such higher order modes would not propagate in the fiber; therefore, any such power in those modes would be lost, hence actually possibly increasing the contrast of the modulator in the system and explaining the larger contrast ratio observed after coupling into the single-mode fiber. It was found that the contrast ratio decreased as the optical power in the system was increased, which may be due to the photogenerated carriers screening the field applied across the MQW region. A test of the electrical contacts on the modulator chip indicated that the metal contacts to the devices were nonohmic, which may be responsible for an inefficient sweepout of the carriers, and such imperfect contacts may also limit the response time of the modulator.

Rise and fall times of the QWAFEMs were measured using an 86109A 30-GHz oscilloscope. After the POM, the setup was modified such that the beam was reflected off a second mirror and into a single-mode fiber. The fiber-coupled light passed through an erbium-doped fiber amplifier and a variable attenuator and into the oscilloscope. The devices were set to send a pattern of ten sequential bits on and then ten sequential bits off to measure the rise and fall times. The fastest transmitter had a rise time of 1.2 ns and a fall time of 900 ps, which are measured from 10% to 90%. The device's estimated capacitance was 1.5 pF. The device with the highest contrast ratio, which was used in the transceiver link, had a rise time of 3.8 ns and a fall time of 3.9 ns, and its estimated capacitance was 1.8 pF.

As these transition times exceed the anticipated values by over an order of magnitude, a combination of electrical and optical testing was performed to determine the root cause. High-speed electrical operation of the driver loaded with a bonded modulator is verified by using on-chip samplers to subsample the output voltage and convert it to a proportional current-driven off-chip and viewed on the oscilloscope. The on-chip driver has an adequate bandwidth for 10-Gb/s operation. However, due to the excessive series contact resistance, which is estimated to be on the order of 1 kΩ, this output drive voltage is filtered at the actual modulator. Thus, the resulting optical waveform has increased transition times, as shown in the 1-Gb/s optical oscilloscope waveforms of Fig. 9.

For the high-speed transceiver link, the test setup was modified such that the output light from the POM was coupled via a Mitutoyo infinity-corrected 20$\times$ NIR objective into a 20-μm diameter high-speed InGaAs/InP photodetector (PDCS20T, Albis Optoelectronics, Switzerland). The photodetectors are attached to the receivers on a second identical CMOS-transceiver chip via short wirebonds [Fig. 8(b)]. This chip is also packaged and attached to a test board mounted to a three-axis translation stage. To enable measurements over a wider range of optical power, the output of the Agilent 81680A tunable laser was coupled via nonpolarization-maintaining fiber through an erbium-doped fiber amplifier, a variable fiber attenuator, and a polarization controller and into the free-space collimator. In this configuration, it was possible to optimize the phase and bias of the detectors. The received data are verified with an on-chip 20-b register whose output can either be scanned out to a computer or also be observed on an oscilloscope. The bit error rate (BER) of individual worst-case bit sequences was measured as the input optical power and the detection phase were adjusted.

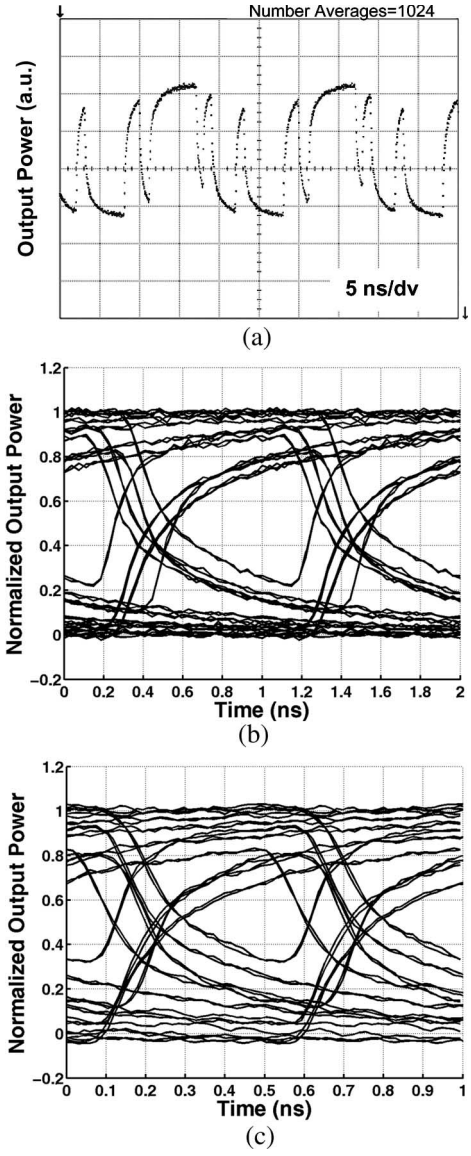

Fig. 9. Optical transmission waveforms. (a) 1-Gb/s, 20-bit data pattern (01111110100010000011). (b) 1-Gb/s pseudoeye diagram created by post-processing oscilloscope data. (c) 2-Gb/s pseudoeye diagram. The waveforms are averaged due to the low SNR after the single-mode-fiber coupling.

While the CMOS transceiver was designed to nominally operate at 5–16 Gb/s, the contact-resistance-limited transition times of the transmitter did not permit operation at that speed. When the chip was too slowly triggered, its performance degraded due to the limited voltage-controlled-oscillator (VCO) range. Thus, in order to get meaningful results from the transceiver link, we synthesized a repeating 10-b pattern by specifying the 20-b sequence in pairs of bits to allow the VCO to operate at a higher frequency. Since the receiver chooses the decision threshold based on the average current at the photodetector, it was necessary to send signals with an equal number of ones and zeros. We tested several bit patterns, attempting to generate the worst-case detection scenario available with 10 b. By taking a histogram of each of the worst-case bits in the pattern, we were able to estimate the error rate. The transmission of 10-b sequences was tested over a range

TABLE I
TRANSMISSION CHARACTERISTICS OF OPTICAL LINK

Data Rate (Gbps)	Sensitivity (dBm)	Bit timing margin
1	-17.6	0.31
1.1	-17.1	0.35
1.2	-19.4	0.63
1.5	-17.7	0.63
1.6	-16.5	0.66
1.7	-15.6	0.68
1.8	-15.2	0.47

Sensitivity is the minimum difference between on and off power at the receiver to attain an estimated BER=10^{-10}. The bit timing margin represents the fraction of a bit period by which the sampler phase may be changed and still achieve an estimated BER=10^{-9}.

of 1–1.8 Gb/s. At 1.8 Gb/s with an average detected power of -15.2 dBm, the BER estimated from the histogram was 10^{-10}. The bit timing margin was such that the BER was estimated at less than 10^{-9} over a total range of phase shift of the receiver clock of 47% of the period of 1 b. Table I shows the collected results of our BER test. The lower data rates required uncharacteristically more power because the link's performance was degraded as the speed was decreased far below the chip's designed clock rate.

The measured 10%–90% rise and fall times correspond to an $\sim$1.7-ns time constant (assuming a simple one-pole system), which, in turn, corresponds in our simulations to $\sim$1.3-Gb/s maximum data rate. Our measured rates of up to 1.8 Gb/s are somewhat higher than this. It is possible that changing the bias voltage on the device during the optimization of the signal in the link may have also changed the rise and fall times from the values measured, hence giving a different data-rate limit.

Loss was measured for the optical path. From the laser source to the free-space collimated beam, the loss was 0.7 dB. Between there and focusing through the microscope, reflecting off the V-grooves, and the separation by the POM, the loss was an additional 7.8 dB. Focusing on the device mesa (in the "pass" state of the device) incurred a loss of 2.3 dB. In focusing the beam on the detector, the loss calculated from a photocurrent measurement was 2.6 dB.

The total transceiver electrical power dissipation is 23.6 mW at 1.8 Gb/s. Transmitter power, including clock generation, is 15.2 mW, with 3.8 mW to drive the QWAFEM, 8.8 mW for the multiplexer and buffers, and 2.6 mW for the TX PLL. The receiver consumes 8.4 mW, including the clock recovery, with 4.5 mW from the integrating/double-sampling front end and 3.9 mW from the CDR circuitry. Total transceiver area is 0.092 mm^2, with 0.017 mm^2 for the transmitter and 0.075 mm^2 for the receiver.

By assuming a resolution of the contact-resistance issue and a nominal increase in driver sizing for the 1.8-pF modulator, the estimated driver power consumption at 16 Gb/s is 69.0 mW, including multiplexing and buffering. At this data rate, the other link components would also consume more dynamic switching power. At 16 Gb/s, the TX PLL consumes 23.0 mW, the RX front-end consumes 23.0 mW, and the RX CDR consumes 35.0 mW, for a total link power of 150 mW [5]. At 10 Gb/s, the total power reduces to 97.8 mW due to the reduced dynamic switching power.

Improvements in power efficiency are possible in parallel I/O systems where there is a potential to easily share the transmit clock-generation PLL among several channels with proper clock distribution and a potential use of phase-correction circuitry. This allows the TX PLL power to be amortized among the channel number, which could range as high as 12 or 20. In a typical 12-channel system, the implemented 1.8-Gb/s transceiver power consumption would drop from 23.6 to 21.2 mW per channel. If the aforementioned contact-resistance issues are resolved, then the power at 10 Gb/s would drop from 97.8 to 84.6 mW. While the implemented per-channel clock recovery does not allow easy sharing of the receiver clock recovery, alternate clocking architectures, such as a source-synchronous forwarded-clock system, allow for similar sharing of receiver clocking circuitry at the expense of an extra dedicated channel for the clock.

V. CONCLUSION

We believe that this paper is the first demonstration of an optoelectronic interconnect transceiver at 1550 nm using an output device directly bonded to the CMOS. Integrating the modulators with CMOS is practical due to their low drive voltage of 2 V and the simplicity of packaging via solder bonds. The QWAFEM architecture is a good candidate for optical interconnect systems due to its surface normal input and output ports, its ability to be arrayed in two dimensions, its misalignment-tolerant optical ports, and its broad bandwidth of operation.

A transceiver is implemented with a pulsed-cascode driver that reliably supplies an output-voltage swing of twice the nominal CMOS supply to the QWAFEMs, allowing for compatibility with present and future scaled CMOS technologies. In addition, the integrating and double-sampling receiver front end allows for demultiplexing of the data directly at the input and eliminates the need for a high-bandwidth TIA. The good power efficiency and the small area of the transceiver make it suitable for high-density optical interconnect applications.

REFERENCES

[1] D. A. B. Miller, "Physical reasons for optical interconnection," *Int. J. Optoelectron.*, vol. 11, no. 3, pp. 155–168, 1997.

[2] D. A. B. Miller, D. S. Chemla, T. C. Damen, A. C. Gossard, W. Wiegmann, T. H. Wood, and C. A. Burrus, "Electric field dependence of optical absorption near the band gap of quantum-well structures," *Phys. Rev. B, Condens. Matter*, vol. 32, no. 2, pp. 1043–1060, Jul. 1985.

[3] N. C. Helman, J. E. Roth, D. P. Bour, H. Altug, and D. A. B. Miller, "Misalignment-tolerant surface-normal low-voltage modulator for optical interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 11, no. 2, pp. 338–342, Mar./Apr. 2005.

[4] S. Palermo and M. Horowitz, "High-speed transmitters in 90 nm CMOS for high-density optical interconnects," in *Proc. Eur. Solid-State Circuits Conf.*, Sep. 2006, pp. 508–511.

[5] S. Palermo, A. Emami-Neyestanak, and M. Horowitz, "A 90 nm CMOS 16 Gb/s transceiver for optical interconnects," in *ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 44–586.

[6] A. Emami-Neyestanak *et al.*, "A 1.6 Gb/s, 3 mW CMOS receiver for optical communication," in *Symp. VLSI Circuits*, Jun. 2002, pp. 84–87.

[7] R. Moazzami and C. Hu, "Projecting gate oxide reliability and optimizing reliability screens," *IEEE Trans. Electron Devices*, vol. 37, no. 7, pp. 1643–1650, Jul. 1990.
[8] A. Emami-Neyestanak *et al.*, "CMOS transceiver with baud rate clock recovery for optical interconnects," in *Proc. Symp. VLSI Circuits*, Jun. 2004, pp. 410–413.
[9] T. Woodward *et al.*, "Modulator-driver circuits for optoelectronic VLSI," *IEEE Photon. Technol. Lett.*, vol. 9, no. 6, pp. 839–841, Jun. 1997.
[10] Y. Leblebici and S. Kang, "Modeling and simulation of hot-carrier-induced device degradation in MOS circuits," *IEEE J. Solid-State Circuits*, vol. 28, no. 5, pp. 585–595, May 1993.
[11] J. Montanaro *et al.*, "A 160 MHz, 32 b, 0.5 W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1703–1714, Nov. 1996.

Jonathan E. Roth (M'07) received the B.S. degree in biomedical engineering from Case Western Reserve University, Cleveland, OH, in 2000. From 2000 to 2001, he was a Fulbright Scholar at Lund Institute of Technology, Lund, Sweden, where he worked on tissue optics and photodynamic therapy. He received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 2007, with a thesis titled "Electroabsorption Modulators for CMOS Compatible Optical Interconnects in III–V and Group IV Materials."

He is currently with the Electrical Engineering Department, Stanford University.

Samuel Palermo (S'98–M'07) received the B.S. and M.S. degrees in electrical engineering from Texas A&M University, College Station, in 1997 and 1999, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 2007.

From 1999 to 2000, he was with Texas Instruments, Dallas, TX, where he worked on the design of mixed-signal integrated circuits for high-speed serial data communication. He is currently with Intel Corporation, Hillsboro, OR, working on high-speed optical and electrical I/O architectures. His research interests include high-speed electrical and optical links, clock recovery systems, and techniques for device variability compensation.

Noah C. Helman received the A.B. degree in physics from Harvard University, Cambridge, MA, in 1998 and the Ph.D. degree in applied physics from Stanford University, Stanford, CA, in 2005.

He is currently researching synthetic biology under Prof. W. Lim at the University of California, San Francisco. His work focuses on engineering the mitogen-activated protein-kinase mating pathway in yeast.

David P. Bour (F'00) was born in Pittsburgh, PA. He received the B.S. degree in physics from the Massachusetts Institute of Technology, Cambridge, in 1983 and the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, in 1988.

He is currently the Chief Scientist of BridgeLux, Inc., which is a high-brightness LED manufacturer based in Sunnyvale, CA. Prior to this, he was Engineering Director with the New Business and New Products Group of Applied Materials, working on the development of metal–organic chemical-vapor-deposition growth technology for compound semiconductors. From 1987 to 1991, he was a member of the Technical Staff with the Sarnoff Corporation, Princeton, NJ. From 1991 to 1999, he was a Principal Scientist with the Electronic Materials Laboratory, Xerox Palo Alto Research Center, Palo Alto, CA, fabricating nitride blue laser diodes and phosphide red laser diodes for laser printing. From 1999 to 2005, he was an Agilent Fellow with the Photonics and Electronics Research Laboratory, Agilent Laboratories, where he worked on the development of semiconductor lasers for electronics, CD and DVD data storage, communications, and projection displays. He has published more than 200 articles and several book chapters. He is the holder of 60 U.S. patents.

David A. B. Miller (M'84–SM'89–F'95) received the B.Sc. degree in physics from St. Andrews University, Fife, U.K., and the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 1979.

From 1981 to 1996, he was with Bell Laboratories, where he was a Department Head, latterly of the Advanced Photonics Research Department since 1987. He is currently the W. M. Keck Professor of electrical engineering and the Director of the Solid State and Photonics Laboratory, Stanford University, Stanford, CA. He is the author or coauthor of more than 250 scientific papers. He is the holder of 62 patents. His current research interests include quantum-well optoelectronic and nanophotonic physics and devices and fundamental and applications of optics in information, sensing, switching, and processing.

Dr. Miller has served as a board member for both the Optical Society of America (OSA), the IEEE Lasers and Electrooptics Society (LEOS), and in various other society and conference committees. He was the President of the IEEE LEOS in 1995. He is a Fellow of the Royal Societies of London and Edinburgh, OSA, and the American Physical Society, and he holds honorary degrees from the Vrije Universiteit Brussel, Brussels, Belgium, and the Heriot-Watt University. He was the recipient of the Adolph Lomb Medal and the R. W. Wood Prize from the OSA, the International Prize in Optics from the International Commission for Optics, and the IEEE Third Millennium Medal.

Mark Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978 and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Yahoo Founders Professor with the Electrical Engineering Department, School of Engineering, Stanford University. In 1990, he took leave from Stanford to help start Rambus Inc., which is a company that designs high-bandwidth memory interface technology. His current research includes multiprocessor design, low-power circuits, high-speed links, and new graphical interfaces.

Dr. Horowitz is a Fellow of the Association for Computing Machinery and a member of the National Academy of Engineering. He was the recipient of the 1985 Presidential Young Investigator Award, the 1993 International Solid-State Circuits Conference Best Paper Award, the International Symposium on Computer Architecture 2004 Most Influential Paper of 1989, and the 2006 IEEE Donald Pederson Award in Solid State Circuits.