Power-efficient dual-rate optical transceiver

Yongrong Zuo, Fouad E. Kiamiley, Xiaoqing Wang, Ping Gui, Jeremy Ekman, Xingle Wang, Michael J. McFadden, and Michael W. Haney

A dual-rate (2 Gbit/s and 100 Mbit/s) optical transceiver designed for power-efficient connections within and between modern high-speed digital systems is described. The transceiver can dynamically adjust its data rate according to performance requirements, allowing for power-on-demand operation. Dynamic power management permits energy saving and lowers device operating temperatures, improving the reliability and lifetime of optoelectronic-devices such as vertical-cavity surface-emitting lasers (VCSELs). To implement dual-rate functionality, we include in the transmitter and receiver circuits separate high-speed and low-power data path modules. The high-speed module is designed for gigabit operation to achieve high bandwidth. A simpler low-power module is designed for megabit data transmission with low power consumption. The transceiver is fabricated in a 0.5 μm silicon-on-sapphire complementary metal-oxide semiconductor. The VCSEL and photodetector devices are attached to the transceiver's integrated circuit by flip-chip bonding. A free-space optical link system is constructed to demonstrate correct dual-rate functionality. Experimental results show reliable link operation at 2 Gbit/s and 100 Mbit/s data transfer rates with ~104 and ~9 mW power consumption, respectively. The transceiver's switching time between these two data rates is demonstrated as 10 μs, which is limited by on-chip register reconfiguration time. Improvement of this switching time can be obtained by use of dedicated input–output pads for dual-rate control signals.

OCIS codes: 200.4650, 200.0200, 250.3140, 250.0250.

1. Introduction

Optical interconnects have dominated for long-haul communications and are increasingly used in short-distance applications such as local area networks. Recent research has shown that optical interconnects are attractive candidates for short-reach links, specifically at the board and backplane levels.[1–3] Compared to their electrical counterparts, they offer inherent advantages in bandwidth–distance product, interconnect density, power consumption, and cross talk. However, with ever-increasing data rates and growing on-chip integration densities, high power density and excessive power dissipation result in increased packaging and cooling costs as well as in potential reliability problems. Additionally, temperature fluctuation induces variations of optoelectronic (OE) device parameters, such as threshold current and slope efficiency of vertical-cavity surface-emitting lasers (VCSELs), which eventually introduce additional jitter in the optical signal and result in degradation of system performance. A low operating temperature, however, has been shown to improve the lifetimes of OE devices such as VCSELs and the reliability of OE systems.[4] Therefore there is a need for low-power optical interconnect designs.

Ongoing research efforts are attempting to achieve low-power characteristics for optical interconnects. These research fields include design optimization for VCSEL drivers[5,6] and optical receivers[7,8]; benchmarking and study of the use of circuit technologies such as complementary metal-oxide semiconductors (CMOS),[9] bipolar,[10] and gallium arsenide[11] technologies; and the integration of various OE devices such as VCSELs[9] and multiple-quantum-well modulators.[12] In addition, low-power design methodologies have been proposed for which the link design parameters and constraints are taken into account and design variables are optimized for a given data rate and system configuration.[13]

All these previous research efforts focused on static design optimization for a target system performance with no consideration given to the dynamic needs of the system, whereas the required system performance may change from time to time in many appli-

The authors are with the Department of Electrical and Computer Engineering, University of Delaware, 114 Evans Hall, Newark, Delaware 19716. Y. Zuo's e-mail address is zuo@ece.udel.edu.

Received 14 April 2005; accepted 18 July 2005.
0003-6935/05/337112-13$15.00/0

cations that we detail in Section 2 below. In this paper we present our research on the design of a dual-rate optical transceiver to achieve dynamic power-efficient operations for optical interconnects in applications with data traffic that varies in intensity. Our transceiver can adjust its data rate dynamically according to performance requirements, allowing for power-on-demand operation. Thus we can save significant power by running the optical link at a lower data rate whenever the performance demand is low. A 0.5 μm silicon-on-sapphire CMOS optical transceiver chip was designed, fabricated, and tested to verify the proposed scheme. This transceiver chip was integrated with VCSEL and photodetector devices by flip-chip bonding. A free-space optical link system was constructed to demonstrate the dual-rate functionality. Experimental results have shown reliable link operation at 2 Gbit/s and 100 Mbit/s data transfer rates with ~104 and ~9 mW power consumption, respectively. The transceiver switching time between these two data rates was demonstrated to be 10 μs, which was limited by on-chip register reconfiguration time. Improvement of this switching time can be obtained by use of dedicated input–output pads for dual-rate control signals. With continual downscaling of CMOS technology, the dual-rate transceiver architecture is expected to scale to higher data rates and lower power dissipation. At the circuit level, the incorporation of dual-rate functionality into a typical gigabit optical transceiver requires 255 additional metal-oxide semiconductor transistors. As a continuation of previously reported research on optical transceivers with fast power-on[14] and accelerated bit-error-rate (BER) test capability,[15] this research facilitates power-efficient operation of optical interconnections.

This paper is organized as follows: In Section 2 we describe our motivation for designing a dual-rate transceiver. Section 3 discusses the design methodology. In Section 4 we present the dual-rate transceiver's architecture and the design of transmitter and receiver circuits. In Section 5 we describe the flip-chip integration of optoelectronic devices (VCSELs and photodetectors) with the CMOS transceiver chip, the integration of the hybrid chip with printed circuit boards (PCBs) for testing, and the free-space optical link setup. In Section 6 we detail the test system's setup, characterize the dual-rate transceiver's performance, demonstrate the functionality of the dual-rate optical link, and present the power-saving results. Finally, the conclusions are given in Section 7.

2. Motivation

Optical interconnects are promising candidates for connections within spaceborne systems. For example, optical links can be used to connect focal plane array sensors to the digital processor in an imaging satellite.[16] Compared to their electrical counterparts, optical links offer lower weight, reduced electromagnetic interference–radio frequency interference sensitivity, improved safety (e.g., complete electrical isolation, immunity to lightning–surge currents, and no sparks) and resistance to countermeasures and eavesdropping.

Most of the imaging satellite platforms today are in near-polar orbits.[17] They follow an orbit (basically north–south) that, in conjunction with the Earth's rotation (west–east), allows them to cover most of the Earth's surface over a certain period of time. Many of these satellite orbits are also Sun synchronous, such that they cover each area of the world at a constant local time of day, so, when they are used to monitor a specific area of interest, the peak performance is required only for brief periods of time when it is on a set orbit of ~90 min. For most of the time they need to remain active only to wait out bad weather or wait for the next orbit over the area of land.[18] During that time the data link operates at much lower rates and carries control and diagnostic information. Because there is limited power available on board such systems, satellite users have a strong desire to run gigabit optical links at megabit rates with low power consumption. Thus a dual-rate optical link that can dynamically lower its power consumption dramatically for data transmission outside the peak period would meet the energy-saving requirement well because the data rate on such optical links would be low most of the time. Such a dual-rate capability becomes a critical factor in selecting optical interconnects in these space-based applications.

This dual-rate concept may also be applied to the optical interconnects that are usually adopted in large networks (e.g., the Internet backbone) to reduce power consumption. The data traffic patterns on these network applications have the characteristics of strong diurnal and weekly cycles.[19,20] The traffic is great during a certain period of the day but otherwise the network traffic is much less. Data traffic can vary by as much as a factor of 5 between peak and minimum use, and it shows an even larger disparity for broadband access networks.[21] Such a clustered nature of the network traffic indicates that the system's power could be reduced dramatically if the network were run at a much lower speed during low data traffic periods. This reduced power consumption also lowers device operating temperatures, improving the reliability and lifetime of OE devices such as VCSELs. These advantages make dual-rate optical transceivers more suitable than traditional optical transceivers for these applications.

3. Design Methodology

To implement dual-rate functionality, the transmitter and receiver circuits include separate high-speed and low-power data path modules (see Fig. 1). The high-speed module is designed for gigabit operation and optimized to achieve the maximum bandwidth. For high performance and to reduce the switching noise, a differential current model logic (CML) standard is used for electrical signaling.[22] Because CML power consumption is constant, the high-speed module consumes the same amount of power when it is operating at low or high speed. To reduce the power

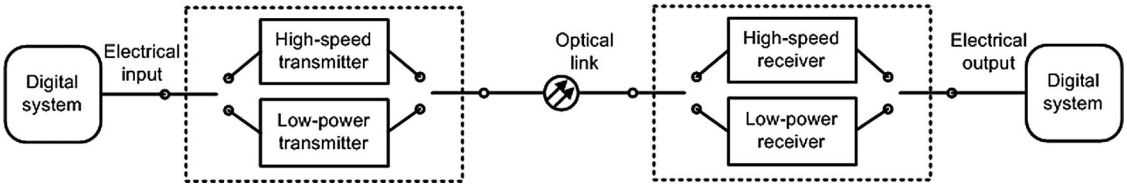

Fig. 1. Diagram of an optical interconnect between digital systems with dual-rate transceivers.

consumption for low-speed data transmission, a separate low-power module is added and, optimized for megabit data transmission; a much simpler circuit design is used. A CML driver with larger output impedance is adopted to lower the driving current while delivering the same voltage swing. To further lower power dissipation at low speed, CMOS logic signaling can be used. To switch dynamically between high-speed and low-power operation, sleep transistors with individual enabling signals are incorporated into both modules to power them on or off.[14,23] One controls these enabling signals by reconfiguring the D flip-flop register chains in the transceiver.

Normally the dual-rate transceiver is integrated within or very close to a digital system, as shown in Fig. 1. However, in most applications the optical connections between two transceivers–systems can be much longer with free-space or fiber optic links. The dual-rate transceiver achieves power saving with a distinct design topology in both the electrical domain and the optical domain for the high-speed and the low-power modules.

A. Power Saving in the Electrical Domain

In the electrical signaling domain, high-speed data transmission is achieved by use of differential signaling with impedance-matched transmission line design. As a rule of thumb, transmission line effects should be considered when the rise or fall time of the input signal (t_r or t_f, respectively) is smaller than 2.5 t_{flight} (t_{flight} is the time of flight).[24] In our case, assuming that the distance of the electrical traces between a digital system and a transceiver is less than 5 cm and that the wave propagation speed on most FR-4 PCBs is ~20 cm/ns, t_{flight} across the electrical traces will be 0.25 ns and 2.5t_{flight} will be 0.625 ns. Here t_r or t_f is estimated as one tenth of duration for a data bit. Thus, for gigabit data transmission (with t_r, t_f = 0.1 ns at most), the electrical trace between the digital system and the transmitter or receiver becomes a transmission line. Therefore the high-speed transceiver module needs proper output impedance or termination impedance to match the characteristic impedance of the transmission line to prevent reflections. However, when the operating frequency is below 100 Mbit/s, t_r or t_f can be 1 ns or more. The electrical trace, therefore, is not considered a transmission line at this rate and can be modeled by a RC circuit. Thus it is possible to design a low-power transceiver module with a larger termination–output impedance value, which requires much less driving current for generating the same voltage swing at the transmitter. Thus less power is consumed with a larger termination or output impedance design. Such power saving can be augmented by use of CMOS signaling instead of differential signaling, as we detail in the following sections.

B. Power Saving in the Optical Domain

Another major source of power savings is the optical transceiver design. At the optical driver side, a VCSEL is used to transmit optical signals with on–off-keyed amplitude modulation. For high-data-rate transmission, usually a VCSEL is biased above threshold current I_{th}, as shown in Fig. 2(a). If it is biased at a current below threshold level, the data-dependent jitter associated with the turn-on delay at the rising pulse edge will dominate and eventually limit its performance.[25] Typically a VCSEL is biased at near threshold when it is sending out data 0's. To get a clean eye diagram and a better BER, a high extinction ratio is usually required. At the low data rate; however, the turn-on delay is significantly shorter than the pulse duration, so the data-dependent jitter becomes negligible. A VCSEL can be turned off completely when it is transmitting 0's and

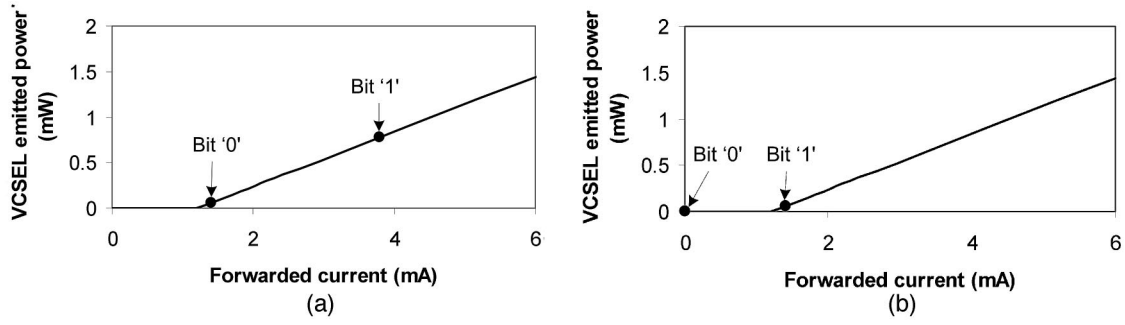

Fig. 2. VCSEL biasing currents for data 1's and 0's (a) at high-speed mode and (b) at low-power mode.

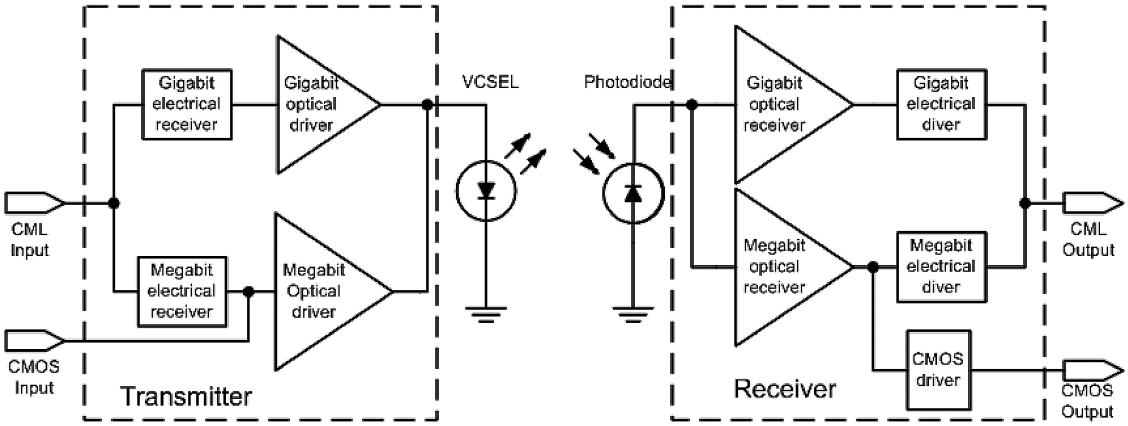

Fig. 3. Diagram of a dual-rate optical transceiver.

can be driven by a small current that is just above I_{th} for transmitting 1's, as shown in Fig. 2(b). This zero-bias on–off switching operation simplifies the driving circuits and leads to the minimum power consumption.[26,27] At the optical receiver side, a transimpedance amplifier (TIA) is adopted for the high-speed module owing to its high bandwidth and high sensitivity. Several fully differential amplifiers with low amplification and high bandwidth are normally required as postamplifier stages to further amplify and recover the received signal. These components exhibit high performance but are also power consuming. They are replaced in the low-power module by a resistor-loaded preamplifier that has relatively bigger load resistors, which converts the received current signal into a voltage signal with sufficient amplitude. Consequently, for all the reasons outlined above, the optical part of the transceiver designed for low-power operation dissipates much less power than its high-speed counterpart.

4. Design of a Dual-Rate Transceiver

A. Transmitter Design

The transmitter module design consists of an electrical receiver, which converts the current input into a voltage signal, and an optical driver, which amplifies the voltage signal that controls the VCSEL bias current (see Fig. 3). Both the transmitter and the receiver comprise a high-speed module and a low-power module. Both the high-speed and the low-power optical drivers drive the same VCSEL, and only one of them is enabled at a time. It is necessary to be able to switch between these two modules quickly for proper transceiver operation. A metal-oxide semiconductor field-effect transistor switch, such as M3 and M4 in Fig. 4, which are called dynamic sleep transistors, is inserted in the path of each current source to power each module on–off. Each transistor is large enough to handle the current that flows through the path with minimum voltage drop. One sleep transistor exists for each current source in the transceiver. Detailed information about the implementation of fast power-on–off with sleep transistors can be found in Ref. 14. Each module in the transmitter is discussed below.

1. *Gigabit Optical Driver*

Two current sources are employed in the high-speed module of the optical driver (Fig. 4). They are implemented with a 6-bit current digital-to-analog converter (DAC). One provides current $I_{bias} + I_{mod}$, and the other provides current I_{mod}. Current source I_{mod} is controlled by a differential pair based on current steering. To transmit a data bit 1, we turn switch M1 on and turn switch M2 off, so I_{mod} is steered to a dummy diode and the VCSEL receives $I_{bias} + I_{mod}$. When the data bit to be transmitted is a 0, however, M2 is turned on and M1 is turned off. I_{mod} is thus connected to the other current source $I_{bias} + I_{mod}$, taking a portion (I_{mod}) of that current source away. So the actual current going through VCSEL is the other portion, I_{bias}. Several differential amplifier stages are placed before the differential pair to amplify the incoming signal from the gigabit electrical receiver and improve the driving strength.

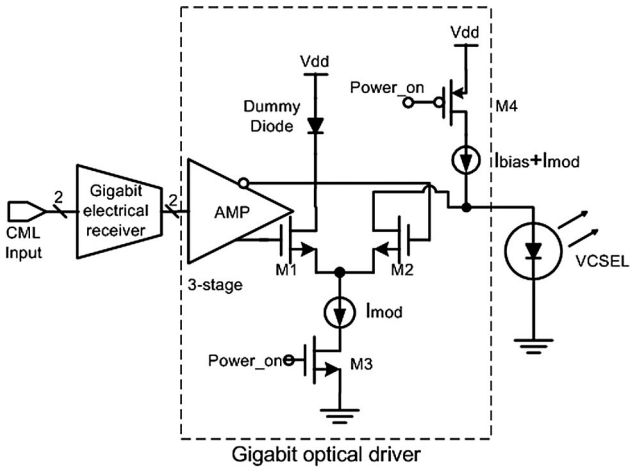

Fig. 4. Diagram of a high-speed transmitter module.

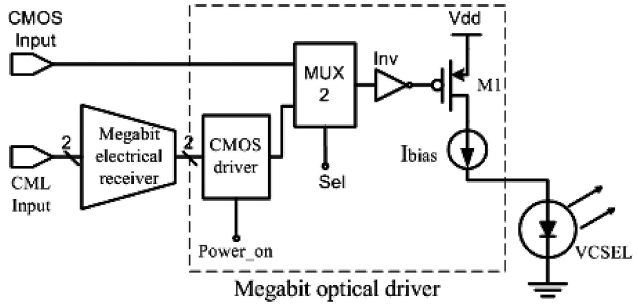

Fig. 5. Diagram of a low-power transmitter module.

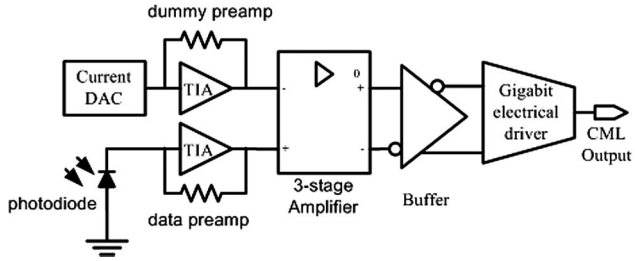

Fig. 7. Diagram of a high-speed receiver module.

2. Megabit Optical Driver

For power saving at the low-speed data rate, a low-power optical driver module is designed as shown in Fig. 5. Multiplexer MUX 2 is used to select one of two optional inputs. CML input or CMOS input. After an inverter, the output controls transistor M1, turning on–off the VCSEL. When the CML input is used, a CMOS driver is used to convert the differential input to a single-ended CMOS signal.

3. Gigabit Electrical Receiver

The design of the high-speed electrical receiver is shown in Fig. 6. By enabling the power-on signal and disabling the power-on-bar signal, we select two 50 Ω resistors as a termination network of the transmission line with a 50 Ω characteristic impedance. The receiver converts the CML current input into a voltage drop to feed the gigabit optical driver.

4. Megabit Electrical Receiver

The design of the low-power electrical receiver is also shown in Fig. 6. We select two 400 Ω impedances by disabling the power-on signal and enabling the power-on-bar signal. A much smaller CML current input (1/8) from the electrical driver is required to generate the same voltage swing at the megabit electrical receiver as its gigabit counterpart. This smaller current reduces the power of the electrical driver that drives this electrical receiver. Details of the electrical driver are discussed in Subsection 4.B. A separate CMOS input pad is added to replace the megabit electrical receiver for the situation when a CMOS input signal is available. The CMOS input pad is especially attractive if the output of a digital system is also a single-ended CMOS signal at a low data rate (<100 Mbits/s). One drawback is that it requires an additional input pad, but ongoing studies indicate that it is possible to reuse one of the CML pads for this input.

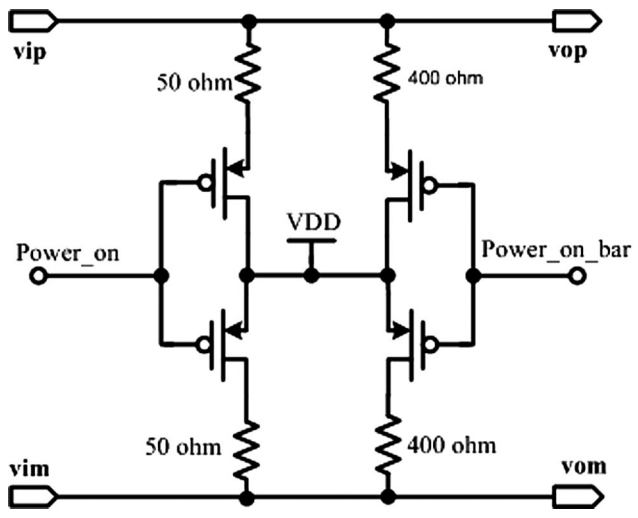

Fig. 6. Electrical receiver, which is used as a gigabit electrical receiver when the power-on signal is enabled or as a megabit electrical receiver when the power-on-bar is enabled. Only one signal can be enabled at a time. VDD, termination voltage.

B. Receiver Design

Our receiver design consists of an optical receiver, which converts the current signal from a photodiode into a voltage signal, and an electrical driver that drives the signal off the transceiver chip. As with the transmitter, there is a high-speed module and a low-power module for each part. For both modules, one dummy preamplifier is designed in addition to one data preamplifier to generate a dc threshold voltage for the decision circuit, as shown in Figs. 7 and 8. Such a dc-coupled receiver is preferred because an ac-coupled receiver usually requires thousands of bits to be transmitted to stabilize the decision threshold voltage by use of a *RC* filter. This long latency prevents the adoptability of our dual-rate transceiver in frequent switching situations. In contrast, our dc-coupled receiver design using the dummy preamplifier eliminates such delay because the threshold value is stored in registers after optimization, mak-

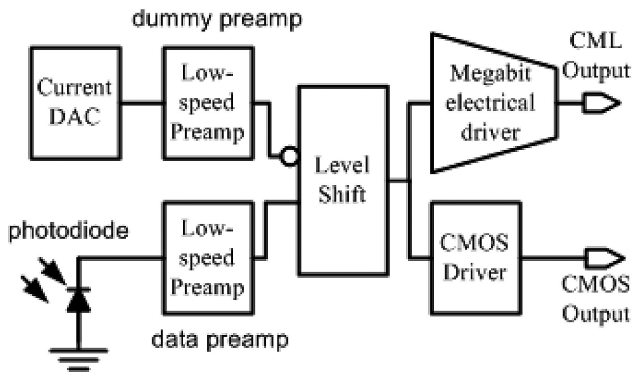

Fig. 8. Diagram of a low-power receiver module.

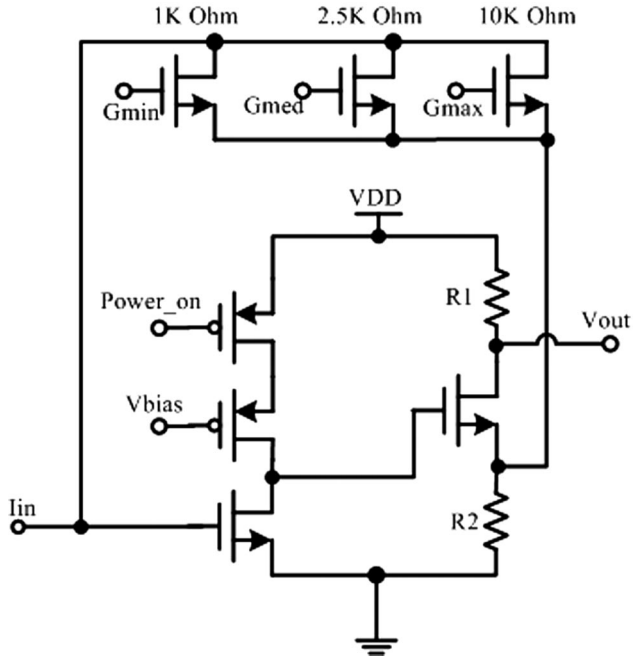

Fig. 9. Transimpedance amplifier with adjustable gain and bandwidth.

ing fast switching between the two operation modes of the dual-rate transceiver possible.

1. Gigabit Optical Receiver

The gigabit optical receiver consists of one preamplifier stage, a three-stage postamplifier, and one buffer stage, as shown in Fig. 7. The preamplifier stage uses a TIA to achieve a high bandwidth and high sensitivity. Its schematic is shown in Fig. 9. The design of the TIA was described before[14,15] and will not be repeated here. Three different gain levels are available to accommodate a wide range of input currents, and one can select them by controlling three digital settings, Gmin, Gmed, and Gmax, in the configuration registers, yeilding 1, 2.5, or 10 KΩ transimpedance gain, respectively. Three fully differential amplifiers are used for the postamplifier stages, and the buffer stage is used to drive an on-chip load for high-speed data transmission.

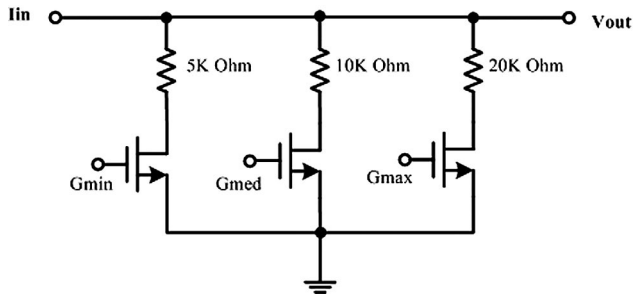

Fig. 10. Preamplifier in a low-power module.

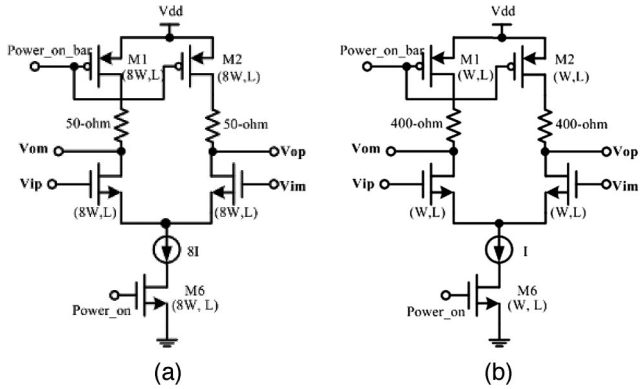

Fig. 11. (a) Gigabit electrical driver with 50 Ω output impedance. (b) Megabit electrical driver with 400 Ω output impedance. The transistors in the gigabit electrical driver are eight times bigger than those in the megabit counterpart. W is the width and L is the channel length of a transistor.

2. Megabit Optical Receiver

Our megabit optical receiver has a resistor-loaded preamplifier stage followed by a level-shift stage, as shown in Fig. 8. Three resistors, with values of 5, 10, and 20 KΩ, are used as low-speed preamplifiers, and we select them by setting the digital control bits as shown in Fig. 10. Compared to the gigabit optical receiver, the megabit receiver, using a relatively larger load resistor, provides higher gain and therefore a large voltage output to the next stage, eliminating the need for additional postamplifier stages.

3. Gigabit Electrical Driver

The gigabit electrical driver is shown schematically in Fig. 11(a). Three sleep transistors, M1, M2, and M6, are used to facilitate high-speed operation and are appropriately sized to have a negligible effect on the output impedance. To match the PCB traces with a characteristic impedance of 50 Ω, which behaves as transmission lines at high speed, 50 Ω resistors are

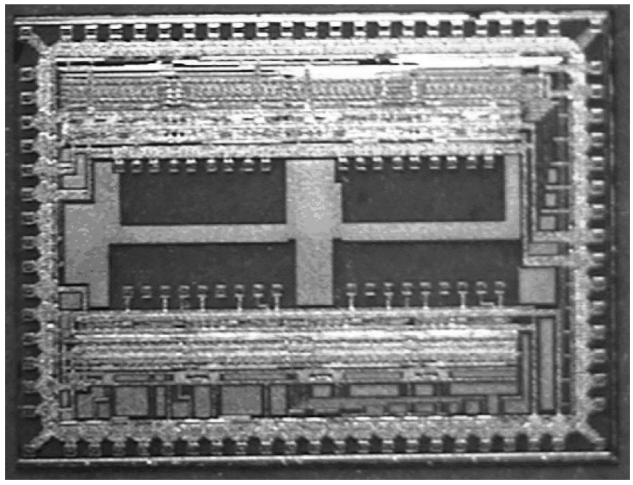

Fig. 12. Micrograph of the CMOS transceiver chip.

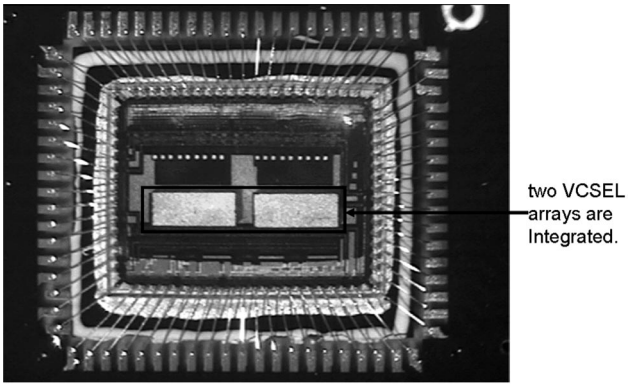

Fig. 13. Hybridized transceiver chip (with VCSELs as shown), wire bonded to the carrier board for optical testing.

used for output impedance to drive the CML current off the chip.

4. Megabit Electrical Driver

For power-saving purposes, the megabit CML driver uses transistors eight times smaller than those in its gigabit counterpart, as shown in Fig. 11(b). The tail current of a different pair is also eight times smaller. To match the design of the megabit electrical receiver in the transmitter we integrated two 400 Ω resistors for output impedance. When it is used to drive the megabit electrical receiver, the current output is only one eighth of the current drawn by the gigabit electrical driver, resulting in approximately seven times less power consumption. To provide maximum flexibility in sending the data off the chip we added a CMOS driver (Fig. 8) besides the CML current output to convert the differential signal into a single-ended CMOS output.

C. Chip Layout

The chip is manually laid out with Cadence Custom integrated circuit design tools front-to-back-design environment, which is a mixed-signal design, including an analog transmitter and receiver part and a digital register chain and control part. To reduce interference between them, we lay them out separately and use separate power-supply rails and input pads. Various analog layout techniques are used to improve matching between transistors. Dummy gates are added to ameliorate the effects of processing nonuniformity. Common centroid geometry is used to minimize gradient-induced mismatches. Large transistors are laid out by use of multiple fingers of unit elements of same length and width, kept in the same orientation and as compact as possible to improve the match.[28] A layout-versus-schematic procedure was performed after layout to ensure matching between the final layout and the schematic design.

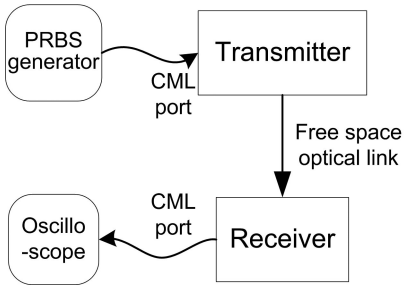

Fig. 15. High-speed testing setup.

D. Fabrication Technology

Our 3 mm × 4 mm chip has 76 perimeter pads (Fig. 12) and was fabricated by use of 0.5 μm CMOS ultrathin silicon-on-sapphire technology from Peregrine Semiconductor Corporation.[29] In this technology, circuits are fabricated in a thin film of silicon grown upon a sapphire wafer. This film brings with it the advantages of low power dissipation, reduced parasitic capacitance, and minimum cross talk. Additionally, sapphire is a good thermal conductor, and its thermal expansion coefficient closely matches that of GaAs, the material of which OE devices are normally made. Thus it is more amenable to hybridization with OE devices than is a silicon substrate. Sapphire is also transparent at wavelengths from 300 to 6000 nm, which makes it possible to use top-emitting OE devices because light can travel through this transparent substrate. These top-emitting devices provide easier integration than bottom-emitting OE

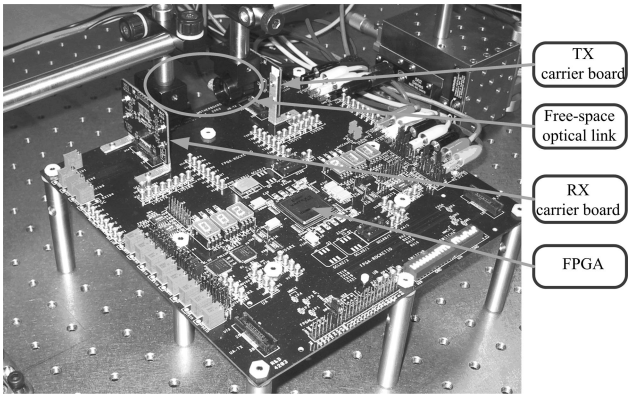

Fig. 14. Demonstration motherboard with two carrier boards on it. TX is transmitter, and RX is receiver.

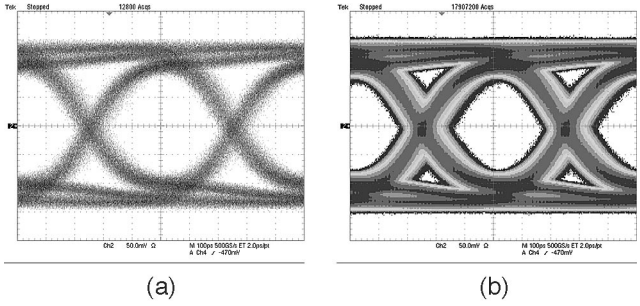

Fig. 16. (a) High-speed link eye diagram measured at 2 Gbits/s. (b) High-speed link eye diagram after the link was run at 2 Gbits/s continuously for 10 h. For both measurements a $2^{15} - 1$ pseudorandom bit sequence was used as source.

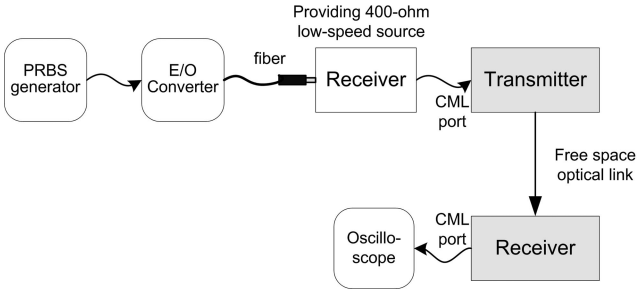

Fig. 17. Low-power testing setup when a CML interface was selected. A second receiver was used to provide a source with 400 Ω output impedance.

devices, which require removal of their substrates when they are hybridized onto CMOS dies.[9]

5. System Integration

A. Flip-Chip Bonding Optoelectronic Devices

The OE devices used in the optical link are 1 × 4 VCSEL arrays and 1 × 4 photodiode arrays from Emcore Corporation. The array sizes are 1.2 mm × 0.45 mm and 1.0 mm × 0.45 mm, respectively. All devices are fabricated in 250 μm pitch. The VCSEL operates at 850 nm and has a typical slope efficiency of 0.4 mW/mA and a typical resistance of 45 Ω at 4–8 mA.[30] The GaAs P–I–N photodiode has a typical responsivity of 0.5 A/W at 850 nm and a 0.4 pF typical capacitance.[31] Both devices are rated for 3.125 Gbit/s applications.

Because flip-chip bonding offers compact connections with low parasitic capacitance, it is used to integrate OE devices onto the transceiver's CMOS die. Both the VCSELs and the photodiodes are top-emitting devices and are bonded face down onto the CMOS die such that light travels through the transparent sapphire substrate.

B. Chip-on-Board Wire Bonding

For testing the hybrid transceiver chip (with the OE devices integrated), a PCB carrier board of size 5.08 cm × 3.81 cm (2 in. × 1.5 in.) was fabricated. Figure 13 shows the hybrid chip wire bonded to the carrier board. A rectangular cutout was made on the board below the chip, allowing the laser to go through. Two carrier boards were placed perpendicularly to the motherboard (see Subsection 5.C below) through high-speed surface mount matched impedance connectors. A free-space optical link was established between them (see Subsection 5.D).

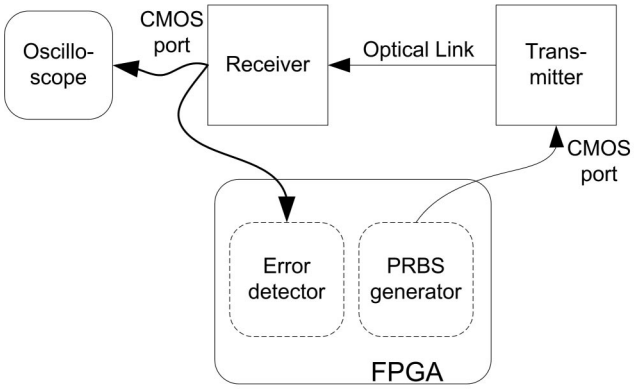

Fig. 19. Low-power testing setup when a CMOS interface was selected.

C. Motherboard

For the demonstration of a complete optical link, an eight-layer motherboard of size 19.05 cm × 21.59 (7.5 in. × 8.5 in.) was designed and fabricated. A Xilinx Virtex II Profield-programmable-gate-array (FPGA; Model XC2VP7) on the motherboard provides control signals to the hybridized chips on the carrier boards through matched impedance connectors. Transceivers can be programmed by the FPGA to select either high-speed mode or low-power mode. Pushbuttons and switches are available to provide manual inputs to the system. HEX displays and LEDs are included on the board for displaying status and results.

D. Free-Space Optical Link Setup

A free-space optical link was set up between two carrier boards through two six-element Universe Kogaku *f*-2.4 lenses, as shown in Fig. 14. These two boards are separated by a distance of 89 mm, perpendicularly connected to the motherboard through matched im-

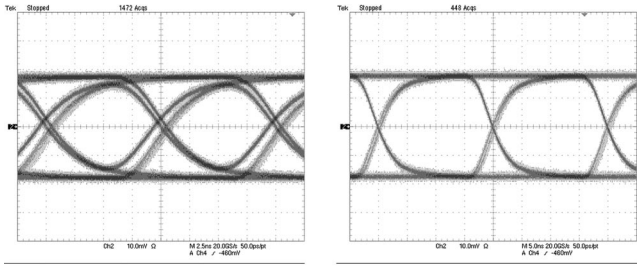

Fig. 18. Eye diagrams for a transceiver working in low-power mode with CML input and output. A $2^{15} - 1$ pseudorandom bit sequence was used. Eye diagrams measured at (a) 100 and (b) 50 Mbits/s.

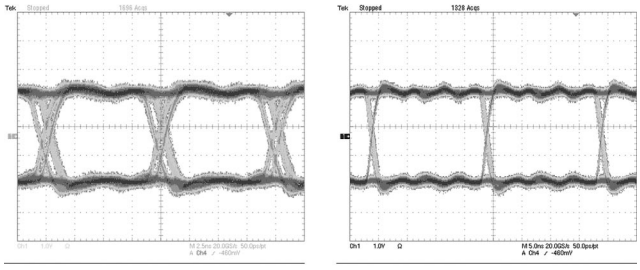

Fig. 20. Eye diagrams when the transceiver worked in the low-power mode with the CMOS input and output. A $2^{15} - 1$ pseudorandom bit sequence was used. Eye diagrams measured at (a) 100 and (b) 50 Mbits/s.

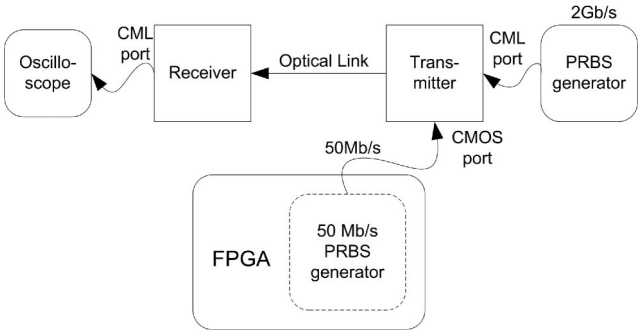

Fig. 21. Mode-switching demonstration diagram. We sent 2 Gbits/s data to a transmitter through a CML interface, and 50 Mbits/s data were sent through a CMOS input.

pedance connectors. Each of the lenses is held with two micropositioning stages for ease of alignment. The lenses are configured in an infinite conjugate ratio imaging system. Each is placed one focal length from its neighboring carrier boards. The first lens collimates the VCSEL beam, and the second lens focuses the light onto the photodiode on the opposite carrier board. For alignment and observation, a tilted reflective neutral-density filter was placed between these lenses as a beam splitter. The filter permits direct observation of the laser spot incident upon its corresponding detector by use of a charge-coupled device (CCD) camera. For optimal fine tuning of the lens alignment, the output voltage of the data preamplifier in the receiver was measured and maximized.

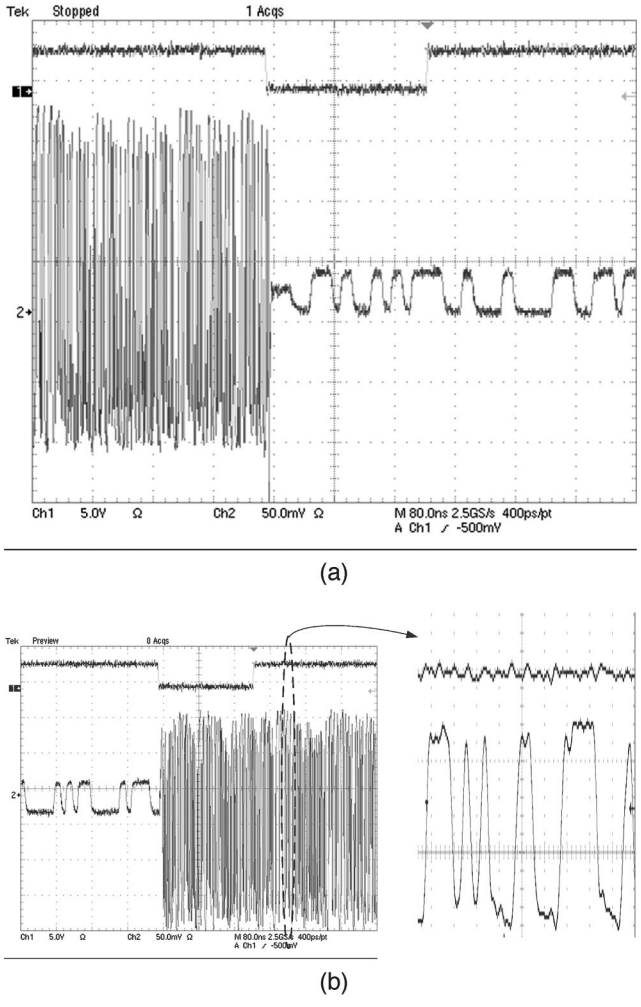

Fig. 22. Mode-switching waveforms when both the CML input and the CMOS input of the transmitter are used. Data transmission switching from (a) 2 Gbits/s to 50 Mbits/s and (b) from 50 Mbits/s to 2 Gbits/s.

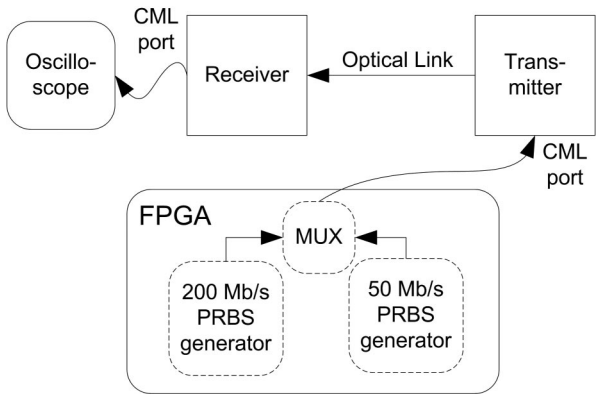

Fig. 23. Another mode-switching demonstration diagram. The CML input of the transmitter was used for both high-speed data and low-power data. MUX, multiplexer.

6. Experimental Results

A. Characterization of Transceiver Performance in High-Speed Mode

In a test of transceiver performance the transceiver was configured to work in the high-speed mode. The testing setup is shown in Fig. 15. Pseudorandom bit sequence (PRBS) data from a pattern generator were sent to the input of the transmitter, and the output of the receiver was connected to a digital oscilloscope to measure the eye diagram. Figure 16(a) shows the eye diagram of the receiver output at the rate of 2 Gbits/s. Bias current I_{bias} and modulation current I_{mod} were set at 0.7 and 3.2 mA, respectively. The TIA was set to minimum gain (Gmin). A BER test was performed, and no error was observed after the link was run for more than 10 h continuously, which indicates that our optical link can run reliably at 2 Gbits/s with a BER of less than 10^{-13}. Figure 16(b) shows the eye diagram after 10 h of continuous operation at 2 Gbits/s.

B. Characterization of Transceiver Performance in Low-Power Mode

The transceiver was reconfigured to work in the low-power mode. Bias current I_{bias} was set to 1.2 mA, and the preamplifier of the receiver was set to minimum gain to use the highest bandwidth available. The two optional interfaces, CML and CMOS, were tested separately for low-power operation.

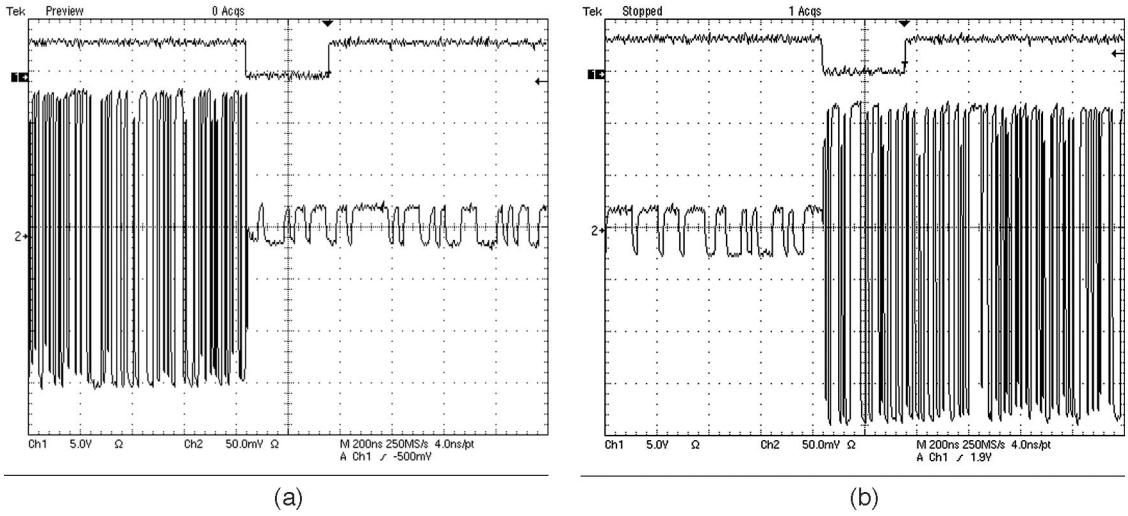

Fig. 24. Mode switching waveforms when only a CML input transmitter was used. Data transmission switching from (a) 200 Mbits/s to 50 Mbits/s and (b) from 50 Mbits/s to 200 Mbits/s.

1. *CML Interface*

Besides the transmitter and the receiver that constructed an optical link as shown in Fig. 17, a second receiver was used to provide a matched data source to the link. A PRBS signal from a pattern generator was converted to an optical signal through an electrical–optical (E/O) converter, and light was butt coupled to the photodiode on the second receiver. Its output was a CML PRBS signal with 400 Ω output impedance.

Figure 18 shows the eye diagram measured at 100 Mbits/s with 1.4 ns of jitter and at 50 Mbits/s with 0.7 ns of jitter. We note that the amplitude of the eye diagrams is only ~40 mV, which can be explained by the fact that the signal driven by a driver with an output impedance of 400 Ω is mismatched with the 50 Ω input impedance of the oscilloscope. The actual signal amplitude when the oscilloscope provides 400 Ω input impedance should be roughly eight times bigger, ~320 mV.

2. *CMOS Interface*

Figure 19 shows the testing setup when CMOS input and output were selected in the low-power mode. The PRBS generated by the on-board FPGA was sent to the CMOS input port of the transmitter. The CMOS output of the receiver was sent back to the FPGA for verification or to an oscilloscope for eye diagram measurement. Figure 20 shows the eye diagrams of the receiver output at the rate of 100 Mbits/s and 50 Mbits/s, respectively. The measured amplitude was ~3 V because CMOS interface was selected.

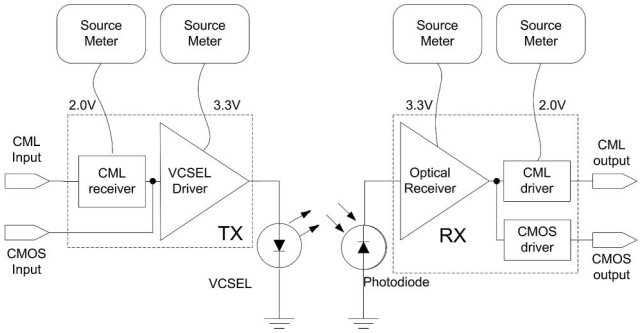

Fig. 25. Four Keithley source meters were used to provide a voltage supply for each part of the transceiver, and we measured their currents to calculate the power dissipation.

C. Mode-Switching Demonstration

The digital control signals that set the operation mode of the dual-rate transceiver are stored in registers that are connected in series as a chain in the transceiver. One can switch the operation modes by shifting data through the register chain, which has a total of 86 registers. Although it is necessary to reconfigure only 6 of these registers for mode switching, it takes 86 clock cycles to reconfigure all the registers. That corresponds to ~9.2 μs when a clock of 9.375 MHz is used. In this test, the transceiver was reconfigured every 10 μs, switching between the high-speed mode and the low-power mode. If a faster switching speed is desired for some applications, six dedicated input pads may be added to the transceiver design for these power-on control signals such that mode switching can be controlled directly instead of by shifting data through the whole register chain. In that case the switching speed will be limited by the power on–off transition times of each module and will be of the order of a few nanoseconds, as was shown previously.[14]

Mode switching was first demonstrated with the CML interface used for high-speed input and the CMOS interface used for low-power input in the transmitter. The testing diagram is shown in Fig. 21. The FPGA was used to generate a 50 Mbit/s PRBS CMOS signal, which was connected to the CMOS input port of the transmitter as a low-power data source. A signal generator created a 2 Gbit/s PRBS signal, which was connected to the CML input port of the transmitter as a high-speed data source. To ob-

Table 1. Power Dissipation of Transmission and Receiver

Mode	Speed	Unit	Transmitter: Electrical Receiver (2.0 V)	Transmitter: Optical Driver (3.3 V)[a]	Receiver: Optical Receiver (3.3 V)	Receiver: Electrical Driver (2.0 V)	Total Power
High speed	2.0 Gbits/s	Current (mA)	0	16.66	10.06	7.82	
		Power (mW)	0	54.9	33.2	15.6	103.8
Low power (CML)	100 Mbits/s	Current (mA)	0	1.60	0.60	0.75	
		Power (mW)	0	5.3	2.0	1.5	8.8
Low power (CMOS)	100 Mbits/s	Current (mA)	0	1.44	0.97	0	
		Power (mW)	0	4.75	3.20	0	8.0

[a]The electrical receiver transmitter (200 V) in no case dissipated power.

serve the waveform of this switching process between these two modes, we fed received data out of the CML port at the receiver to an oscilloscope. Figure 22 shows the mode-switching waveforms. The waveform at the top in each figure is the switching signal that is effective at low speed. The waveform at the bottom is the actual data received. Figure 22(a) shows that 2 Gbit/s high-speed data were received at first, and then we switched to the 50-Mbit/s low-power data after the switching signal was established. The amplitude of 50 Mbits/s is relatively small on the oscilloscope owing to impedance mismatch. Figure 22(b) shows the switching from 50 Mbits/s to 2 Gbits/s. The extended waveform for 2 Gbits/s is also shown, at the right in Fig. 22(b).

Mode switching was also demonstrated with only the CML port used for both the high-speed mode at 200 Mbits/s and the low-power mode at 50 Mbits/s as shown in Fig. 23. Here the FPGA was used to source the data at both speeds, and the high-speed data rate (200 Mbits/s) was limited by the performance of the FPGA. After a 2:1 multiplexer, these two data sources were connected to the CML input of the transmitter. The 2:1 multiplexer switched the data sources by mode switching the signal. Figure 24 shows the transition waveforms.

D. Power Dissipation Measurement

It is commonly known that the power in synchronous CMOS digital systems is dominated by dynamic power dissipation. However, for analog circuits such as a dual-rate transceiver, power dissipation is set by the static current consumed, such that $P_{\text{static}} = V_{\text{dd}} \times I_{\text{static}}$,[32] where V_{dd} is the supply voltage and I_{static} is the measured current.

To monitor the power dissipation for the transceiver we used four Keithley source meters to provide power supplies and measure the currents drawn by the four sections of the transceiver: the electrical receiver and the optical driver on the transmitter side and the optical receiver and the electrical driver on the receiver side. Figure 25 shows the setup. The measured currents and the calculated power dissipation in both modes are summarized in Table 1. For low-data-rate operation, the power consumption for the CML interface and for the CMOS interface was measured separately.

Because the electrical receiver in the transmitter consists only of a termination network, it should not consume power and, accordingly, no power consumption was observed. As we can see from Table 1, this transceiver consumes a total power of 103.8 mW when it is running at 2 Gbits/s but only 8.8 mW for the CML interface or 8.0 mW for the CMOS interface at 100 Mbits/s. Thus a power saving of more than 11 times is obtained when the transceiver is switched from high-speed mode to low-power mode. With the downscaling of CMOS technology that permits a more-compact circuit layout and lower voltage supply,[33] it is reasonable to expect this dual-rate transceiver to work at higher data rates and with lower power dissipation.

E. Die Area and Power Consumption Overhead

The power saving is obtained at the expense of addition of another low-power data path, which means more die area consumed. The additional logic that controls the mode-switching activity adds to the power consumption overhead. We now examine how much cost is incurred to reap this benefit.

The die area per channel including both a high-speed module and a low-power module is 630 μm × 250 μm for the transmitter and 670 μm × 200 μm for the receiver. The low-power transmitter module contains 99 transistors and consumes 280 μm × 100 μm of the die area. The low-power receiver module has 156 transistors and occupies 340 μm × 100 μm of the die area. So the total cost for adding a single-channel low-power module is the addition of 255 transistors, which corresponds to 620 μm × 100 μm of die area on this 0.5 μm CMOS chip, or approximately a 21.2% increase.

The control logic to switch the operation modes was implemented as six registers in this design. As we know, CMOS digital circuits draw a significant level of current only during a switching event. The main power consumption for these digital circuits is dominated by dynamic power dissipation, $P_{\text{dyn}} = f V_{\text{dd}}^2 C_{\text{out}}$,[34] where f is the switching frequency, V_{dd}^2 is the supply voltage, and C_{out} is the parasitic capac-

itance. So the power consumed by this additional logic will be increased with more-frequent switching. When the switching frequency is extremely high, the power overhead from the additional control logic may cancel the power saving from the dual-rate operation. To circumvent this problem, one could replace the registers of the control logic with six dedicated input–output pads to take the switching control signals directly into the transceiver chip if the chip is not input–output limited. This would remove the power consumption overhead of the control logic, and the switching speed would not be limited to the register reconfiguration time.

7. Conclusions

We have described a dual-rate optical transceiver designed for power-efficient optical connections within and between modern high-speed digital systems. The transceiver has a capability to adjust its data transfer rate dynamically according to performance requirements, allowing for power-on-demand operation. A 0.5 μm silicon-on-sapphire CMOS optical transceiver chip was designed, fabricated, and tested to verify the proposed scheme. The design of a dual-rate transceiver was described in detail, together with its two operation modes. The integration of the transceiver chip with OE devices by flip-chip bonding was presented. A free-space optical link system was set up for the demonstration of the dual-rate functionality. Experimental results have shown that the transceiver can run reliably at 2 Gbits/s for high-speed transmission with power dissipation of ~104 mW and consume less than 9 mW of power at 100 Mbits/s for low-power operation, resulting in a more than 11-times power savings. The overhead for adding a low-power module is the addition of 255 more transistors. The transceiver's switching time between high-speed mode and low-power mode was demonstrated as 10 μs, which was limited by the register reconfiguration time. Using dedicated input pads for mode-switching signals can shorten this time for more-frequent switching situations.

The dual-rate transceiver was designed as a proof of concept. To our best knowledge, this is the first reported dual-rate optical transceiver that has successfully demonstrated the trade-off between speed and power in a complete optical link. With innovative design and careful layout of the transceiver, together with the downscaling of CMOS technology, it is possible to further increase the speeds of both the gigabit and the megabit modules and to decrease their power dissipation. For more flexibility, we may consider designing optical transceivers with more optional modules to fill in the operating frequency gap that exists in the current prototype of the dual-rate transceiver. By applying the design techniques for the dual-rate transceiver described here, it is possible to build optical links with more power efficiency by virtue of the ability to trade off the power and link performance such that power is conserved whenever possible.

The authors thank Matthew Roughan of the University of Adelaide and Liangping Ma of the University of Delaware for discussions of network traffic and the possible applications of the dual-rate optical transceiver.

References

1. D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," Proc. IEEE **88,** 728–749 (2000).
2. A. Z. Shang and F. A. P. Tooley, "Digital optical interconnects for networks and computing systems," J. Lightwave Technol. **18,** 2086–2094 (2000).
3. M. W. Haney, M. P. Christensen, P. Milojkovic, J. Ekman, P. Chandramani, R. Rozier, F. E. Kiamilev, Y. Liu, and M. Hibbs-Brenner, "Multichip free-space global optical interconnection demonstration with integrated arrays of vertical-cavity surface-emitting lasers and photodetectors," Appl. Opt. **38,** 6190–6200 (1999).
4. B. Hawkins and B. Hawthorne, "2.5 Gpbs oxide isolated VCSEL reliability report" (Honeywell International, 2003).
5. L. Chen, M. Li, C. Chang-Hasnain, and K. Lau, "A low-power 1-Gb/s CMOS laser driver for a zero-bias modulated optical transmitter," IEEE Photon. Technol. Lett. **9,** 997–999 (1997).
6. R. Tao, M. Berroth, and Z. Wang, "Low power 10 Gbit/s VCSEL driver for optical interconnect," Electron. Lett. **39,** 1743–1744 (2003).
7. A. Emami-Neyestanak, D. Liu, G. Keeler, N. Helman, and M. Horowitz, "A 1.6 Gbps, 3mW CMOS receiver for optical communication," in *IEEE Symposium on VLSI Circuits* (Institute of Electrical and Electronics Engineers, 2002), pp. 84–87.
8. M. Ingels and M. S. J. Steyaert, "A 1-Gb/s, 0.7-m CMOS optical receiver with full rail-to-rail output swing," IEEE J. Solid-State Circuits **34,** 971–977 (1999).
9. D. Plant, M. B. Venditti, E. Laprise, J. Faucher, K. Razavi, M. Chateauneuf, A. G. Kirk, and J. S. Ahearn, "256-Channel bidirectional optical interconnect using VCSELs and Photodiodes on CMOS," J. Lightwave Technol **19,** 1093–1103 (2001).
10. T. Nagahori and N. Suzuki, "An analog front-end chip set employing an electro-optical mixed design on SPICE for 5-Gb/s/ch parallel optical interconnection," IEEE J. Solid-State Circuits **36,** 1984–1991 (2001).
11. O. Qasaimeh, W. Zhou, P. Bhattacharya, D. Huffaker, and D. G. Deppe, "Monolithically integrated low-power phototransceivers for optoelectronic parallel sensing and processing applications," J. Lightwave Technol. **19,** 546–552 (2001).
12. A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, T. K. Woodward, J. E. Ford, G. F. Aplin, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, L. M. F. Chirovsky, and D. A. B. Miller, "3-D integration of MQW modulators over active submicron CMOS circuits: 375 Mb/s transimpedance receiver-transmitter circuit," IEEE Photon. Technol. Lett. **7,** 1288–1290 (1995).
13. O. Kibar, D. A. Van Blerkom, C. Fan, and S. C. Esener, "Power minimization and technology comparisons for digital free-space optoelectronic interconnections," J. Lightwave Technol. **17,** 546–555 (1999).
14. P. Gui, F. E. Kiamilev, X. Q. Wang, X. L. Wang, M. J. McFadden, M. W. Haney, and C. Kuznia, "A 2-Gb/s 0.5μm CMOS parallel optical transceiver with fast power-on capability," J. Lightwave Technol. **9,** 2135–2148 (2004).
15. X. Wang, F. Kiamilev, P. Gui, J. Ekman, G. C. Papen, M. McFadden, M. W. Haney, and C. Kuznia, "A 2-Gb/s optical transceiver with accelerated bit error ratio test capability," J. Lightwave Technol. **32,** 2158–2167 (2004).
16. N. Gata, S. Subramaniana, S. Rossa, C. LaBawb, and J. Bond, "Thermal infrared imaging spectrometer (TIRIS) status report," in *Infrared Technology & Applications XXIII,* B. F.

Andresen and M. Strojnik, eds., Proc. SPIE **3061,** 284–291 (1997).
17. Canadian Center For Remote Sensing online tutorial: http://www.ccrs.nrcan.gc.ca/ccrs/learn/tutorials/fundam/chapter2/chapter2_2_e.html.
18. P. Fricker, "Airborne imaging technology: today's reality," Directions Mag. (19 April 2004).
19. J. W. Roberts, "Traffic theory and the Internet," IEEE Commun. Mag. **39,** 94–99 (2001).
20. M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates, and Y. Zhang, "Experience in measuring Internet backbone traffic variability models, metrics, measurements and meaning," presented at *the 18th International Teletraffic Congress,* Berlin, Germany, 2003.
21. M. Roughan and C. Kalmanek, "Pragmatic modeling of broadband access traffic," Computer Commun. **26,** 804–816 (2003).
22. M. Anis, M. Allam, and M. Elmasry, "Impact of technology scaling on CMOS logic styles," IEEE Trans. Circuits Syst. II **49,** 577–588 (2002).
23. J. W. Tschanz, S. G. Narendra, Y. Ye, B. A. Bloechel, S. Borkar, and V. De "Dynamic sleep transistor and body bias for active leakage power control of microprocessor," IEEE J. Solid-State Circuits **38,** 1838–1845 (2003).
24. J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits,* 2nd ed. (Pearson Education, 2003), p. 168.
25. J. Tatum and J. Guenter, "Modulating VCSELs" (Honeywell International, 1998).
26. L. Zei, K. Petermann, R. Jäger, and K. J. Ebeling, "Operation range of VCSEL-interconnect links with 'below-threshold' biasing," J. Lightwave Technol. **18,** 477–481 (2000).
27. L. P. Chen and K. Y. Lau, "Regime where zero-bias is the low-power solution for digitally modulated laser diodes," IEEE Photon. Technol. Lett. **8,** 185–187 (1996).
28. A. Hastings, *The Art of Analog Layout* (Prentice-Hall, 2001).
29. P. McAdam and B. Goldberg, "UTSi CMOS SOS for mixed signal ICs" (Peregrine Semiconductor Corporation, 2001).
30. Emcore Corporation, "Spec, Oxide Array Die, 1 × 4 3.125 Gb (U-lot), 8685-1402," data sheet (Emcore Corporation, 2001).
31. Emcore Corporation, "Spec, PIN Array Die 1 × 4 3.125 Gb, S8485-1405," data sheet (Emcore Corporation, Albuquerque, 2001).
32. M Pedram and J. Rabaey, *Power Aware Design Methodologies* (Kluwer Academic, 2002), Chap. 8.
33. H. Iwai, "CMOS technology-year 2010 and beyond," IEEE J. Solid-State Circuits **34,** 357–366 (1999).
34. J. P. UyeMuru, *CMOS Logic Circuit Design* (Kluwer Academic, 1999), Chap. 3.