A Single-Chip CMOS-Based Parallel Optical Transceiver Capable of 240-Gb/s Bidirectional Data Rates

Clint L. Schow, Fuad E. Doany, *Member, OSA*, Christian W. Baks, Young H. Kwark, *Member, IEEE*, Daniel M. Kuchta, *Senior Member, IEEE*, and Jeffrey A. Kash, *Fellow, IEEE*

***Abstract*—We report here on the design, fabrication, and high-speed performance of a parallel optical transceiver based on a single CMOS amplifier chip incorporating 16 transmitter and 16 receiver channels. The optical interfaces to the chip are provided by 16-channel photodiode (PD) and VCSEL arrays that are directly flip-chip soldered to the CMOS IC. The substrate emitting/illuminated VCSEL/PD arrays operate at 985 nm and include integrated lenses. The complete transceivers are low-cost, low-profile, highly integrated assemblies that are compatible with conventional chip packaging technology such as direct flip-chip soldering to organic circuit boards. In addition, the packaging approach, dense hybrid integration, readily scales to higher channel counts, supporting future massively parallel optical data buses. All transmitter and receiver channels operate at speeds up to 15 Gb/s for an aggregate bidirectional data rate of 240 Gb/s. Interchannel crosstalk was extensively characterized and the dominant source was found to be between receiver channels, with a maximum sensitivity penalty of 1 dB measured at 10 Gb/s for a victim channel completely surrounded by active aggressor channels. The transceiver measures 3.25 × 5.25 mm and consumes 2.15 W of power with all channels fully operational. The per-bit power consumption is as low as 9 mW/Gb/s, and this is the first single-chip optical transceiver capable of channel rates in excess of 10 Gb/s. The area efficiency of 14 Gb/s/mm^2 per link is the highest ever reported for any parallel optical transmitter, receiver, or transceiver reported to-date.**

***Index Terms*—CMOS analog integrated circuits, crosstalk, driver circuits, optical communication, optical receivers, optoelectronic devices, photodetectors, photodiodes, semiconductor laser arrays, semiconductor lasers.**

I. Introduction

WIDE, high-speed optical data buses integrated on electrical printed circuit boards (PCBs) have the potential to meet the extreme bandwidth demands of future high-performance computing and switch/router systems. Continued CMOS device scaling and the advent of multicore chip architectures have kept on-chip performance gains on a steady growth curve. Unfortunately, packaging technology, including the high-speed data links that interconnect modules on PCBs, has not kept pace. The resulting shortfall between on-chip and off-chip interconnect bandwidth typically forces system-level architectural tradeoffs that could be avoided if greater bandwidth could be achieved in chip-to-chip data links [1]. Employing parallel optics to interconnect modules on PCBs has the potential to elevate the bandwidth of on-board data buses to the Tb/s level [2]. Today's high-performance computers typically use fiber-optic interconnections for multi-Gb/s/channel high-speed data links longer than ∼10 m. Indeed, the superiority of optical interconnects over electrical data links for sending high-speed data over long distances is well established. Over time, optical links have tended to displace electrical links for shorter and shorter interconnects as increasing data rates have made communications using electrical signaling more problematic. However, commercially available parallel optical modules tend to be too bulky, consume too much power, support too few channels, and cost too much on a dollar/Gb/s basis to be useful for chip-to-chip data busses on distance scales < 1 m. Indeed, for on-PCB optical data buses to become viable, transceiver component technology must be developed that possesses many challenging attributes, namely: low-cost, low-power consumption, high area density, a high per-channel data rate, and scalability to high channel counts.

Commercial parallel transmitter (TX), receiver (RX), and transceiver modules typically incorporate 850-nm sources and detectors and are predominately built according to the "SNAP12" multisource agreement [3]. These components are useful for remote I/O over distances from ∼10–300 m using protocols such as Infiniband.[1] However, they offer inadequate bandwidth (currently 5 Gb/s/ch) and density (∼0.17 Gb/s/mm^2), with rather high power consumption (∼2.0 W for a 12-channel TX/RX pair, or 33 mW/Gb/s/link), for board-level optical interconnects. Higher channel rates in the SNAP12 package, up to 10 Gb/s, have been demonstrated using high-speed prototype amplifier chips in both SiGe [4] and CMOS technologies [5]. The per-link power dissipation of the SiGe modules, ∼19.2 mW/Gb/s/link, represented a substantial reduction compared to commercial offerings. In addition, it was shown that replacing the SiGe laser driver with a CMOS version reduced transmitter power consumption by ∼15% [5]. Fiber-coupled, 10 Gb/s/ch, 12 channel modules have also been built with 990-nm optoelectronic (OE) components [6]. These "POSH" modules consumed 5.5 W per TX/RX pair in a footprint of ∼22 cm^2, corresponding to per-link power and area

Manuscript received March 18, 2008; revised May 22, 2008. Current version published April 17, 2009. This work was supported by DARPA under Contract MDA972-03-3-0004.

The authors are with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: cschow@us.ibm.com; doany@us.ibm.com; cbaks@us.ibm.com; kwark@us.ibm.com; kuchta@us.ibm.com; jeffkash@us.ibm.com).

Digital Object Identifier 10.1109/JLT.2008.927759

[1][Online]. Available: http://www.infinibandta.org

efficiencies of $\sim$46 mW/Gb/s and 0.05 Gb/s/mm^2. It should be noted that the large footprint of fully packaged modules [4]–[6] is primarily a consequence of the large fiber ribbon connector and heat sink that they incorporate.

In addition to packaged modules, there have been efforts to produce highly integrated chip-like, multi-Gb/s/ch, optical transmitters and receivers. In the MAUI program, separate 48-channel fiber-coupled transmitters and receivers were demonstrated [7]. These modules employed coarse wavelength division multiplexing (CWDM) optics with separate 1×12 VCSEL arrays at four different wavelengths. This CWDM approach allowed 48-channel TX/RX pairs to be connected through a standard 12-channel ribbon fiber for reduced cabling cost. Complete 48-channel data links were demonstrated at 6.25 Gb/s/ch, with single-channel operation up to 10.42 Gb/s reported at a total power consumption of 3 W [8]. This work partially addressed the power and density issues identified above, yielding improved figures of merit of 6 mW/Gb/s/link and 6.25 Gb/s/mm^2. However, the CWDM optics contributed 6–8 dB of optical loss that consumed the bulk of the link budget, and no direct crosstalk measurements were reported.

Large-scale 2-D arrays of optical interconnects have also been explored, with research efforts yielding 256 [9] and 540 [10] element transmitters and receivers operating at 850 nm. Although the channel count was high in this approach, the achieved per-channel data rates were relatively low, on the order of 250 Mb/s. Dense integration of CMOS with $\sim$980 nm OE arrays has also been pursued. A 970-nm, 256-channel transmitter chip has been reported, with 80 channels successfully operating at a data rate of 1 Gb/s [11]. Efforts to commercialize 72-channel highly integrated parallel optical modules have also been undertaken at $\sim$ 2.5 Gb/s per-channel data rates [12]. These highly parallel assemblies [9]–[12] serve to illustrate the potential of large scale OE/CMOS integration for future massively parallel optical interconnects. However, for optical interconnects to displace electrical data links, the bit rate offered by the optics must at least be comparable to, and likely greater than that of state of the art electrical I/O. Therefore, the per-channel data rate of parallel optical data buses must exceed $\sim$10 Gb/s to be considered as a viable potential replacement for electrical links in short-reach (< 1 m) applications.

We report here on the design, assembly and high-speed performance of a novel single-chip CMOS parallel optical transceiver, or Optochip, incorporating 16 transmitter and 16 receiver channels. The Optochip is a chip-like assembly that is designed to be packaged in a manner similar to that employed for today's electronic ICs [13]. A cross-sectional schematic of a board-level optical interconnect based on the transceiver Optochip is presented in Fig. 1. The Optochip, which is comprised of VCSEL and photodiode (PD) arrays that are flip-chip bonded to a CMOS IC, is directly soldered to a high-density organic chip carrier to form an optical module, or Optomodule. The Optomodule is then soldered to an underlying circuit board, or Optocard, that contains dense arrays of polymer optical waveguides, 45° turning mirrors, and lenses for efficient optical coupling. A pair of Optomodules mounted to, and communicating through, an Optocard constitutes a complete optical data bus [14].

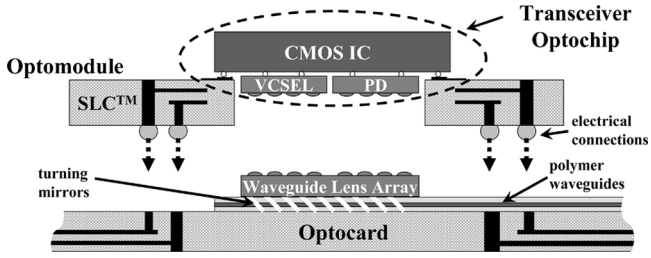

Fig. 1. Cross-sectional schematic of a board-level optical interconnect: Optomodules communicating through optical waveguides integrated in an Optocard.

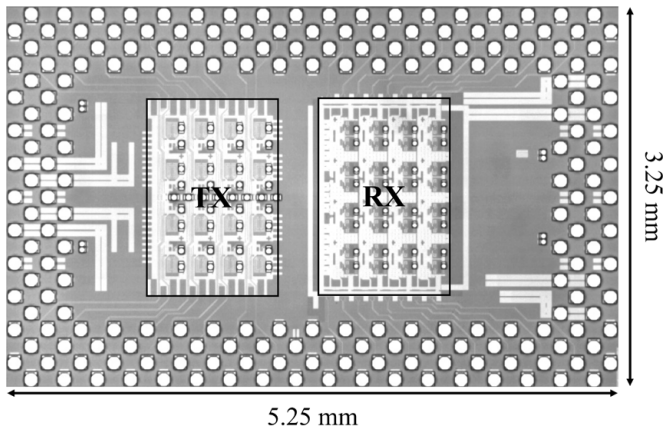

Fig. 2. Photomicrograph of the transceiver integrated circuit illustrating the TX and RX blocks.

The Optochip is designed to possess key attributes required of future optical components to enable dense integrated optical data buses. The Optochip is designed for low-cost by incorporating: 1) unmodified production bulk CMOS technology; 2) VCSEL emitters; 3) flip-chip solder packaging similar to electronic ICs; and 4) operation at 985-nm enabling collimating lenses to be integrated into the substrates of the OE arrays. The power consumption of the Optochip is minimized through: 1) employing low-power CMOS designs for the receiver and transmitter circuits; 2) using efficient VCSELs and PDs for electrooptical conversion; and 3) reducing total optical link loss by employing an efficient two-lens optical system and low-loss polymer waveguides [15]. Finally, in addition to offering per-channel data rates well in excess of 10 Gb/s, the 2-D array architecture of the Optochip provides very high channel density and straightforward scalability to higher channel counts.

II. Transceiver Chip Architecture

The fundamental building block of the transceiver Optochip is the underlying CMOS IC that contains two independent 4×4 arrays of receiver and laser diode driver (LDD) circuits. The chip was fabricated by IBM in the CMOS8RF technology, a 0.13–μm process with 8 levels of interconnect metal. This variant of IBM's CMOS technology is specifically targeted for analog/RF applications through the inclusion of two thick metal layers at the top of the metal stack. The receiver and LDD arrays both consist of 16 independent channels in a 4×4 floor plan with a pitch of 250×350 μm. Fig. 2 shows a micrograph of the fabricated 3.25×5.25 mm transceiver IC with the transmitter and receiver blocks highlighted. A single

ground plane is utilized, with separate power supplies for the transmitter and receiver sections applied through bond pads on the short sides of the IC. The TX and RX sections are both subdivided into two power domains containing eight channels each. Due to somewhat undersized distribution networks, the power connections to the transmitter and receiver blocks have 2–3 Ω of series resistance that result in a parasitic on-chip voltage drop when power is applied to the IC. Finally, the transmitter and receiver blocks are separated by 500 μm. This separation not only serves to minimize TX-to-RX interchannel crosstalk, but also keeps the spacing between the OE elements at 250 μm in the long chip dimension for consistency with the complete Optomodule and Optocard package design described in [13, Section I].

The high-speed data I/O are routed to the long sides of the IC through on-chip 100-Ω differential microstrip transmission lines. The signal lines are routed on the final thick metal layer to minimize loss. At the edges of the chip, the transmission lines terminate in 100-μm octagonal bondpads arranged in a ground-signal-signal-ground (GSSG) configuration. All of the perimeter bondpads are arrayed at a 200-μm pitch to facilitate the direct flip-chip bonding of the transceiver to a high-density organic chip carrier in a manner similar to that employed for conventional electronic ICs (e.g., the IBM C-4 process [16]). As clearly evidenced in Fig. 2, the size and pitch of the perimeter pads dictated the overall size of the chip: the active circuitry occupies less than 20% of the total real estate. This decision to conform to typical IC flip-chip design rules for pad placement, at the expense of an expanded chip footprint, was made to allow the transceiver chip to be further packaged using common surface-mount packaging tools to minimize cost. In addition, the unused area could be used to incorporate the control, monitoring, and interface circuitry commonly found in commercial transceivers. A further important attribute of the chosen pad geometry is that all channels can be fully characterized at the Optochip level using 200-μm pitch GSSG coplanar microwave probes. Connections to the OE arrays are made through solder "microbumps" with individual anode and cathode connections to the OE diodes present at each of the array elements. The diameter of the microbumps is 30 μm for the photodiode and 45 μm for the VCSEL attachment sites, respectively. Connection to the VCSEL chip is further reinforced through a row of unconnected mechanical pads that run down the center of the transmitter block.

III. CMOS Circuits

The IBM CMOS8RF process that was employed to build the transceiver chip supports the inclusion of integrated passive devices that are employed in both the LDD and RX circuits. In addition, the TX and RX circuits both share a common footprint of 250 μm $\times$ 350 μm stepped in a 4 $\times$ 4, 2-D array. Local power supply decoupling occupies a significant portion of the real estate of both the driver and receiver cells. One major advantage of designing the transceiver using 985-nm substrate emission/detection OE devices is that the PD and VCSEL arrays are directly flip-chip bonded to the CMOS transceiver chip. This feature provides an essentially zero-length interconnection between the OE devices and analog amplification circuits which minimizes parasitics associated with this critical connection.

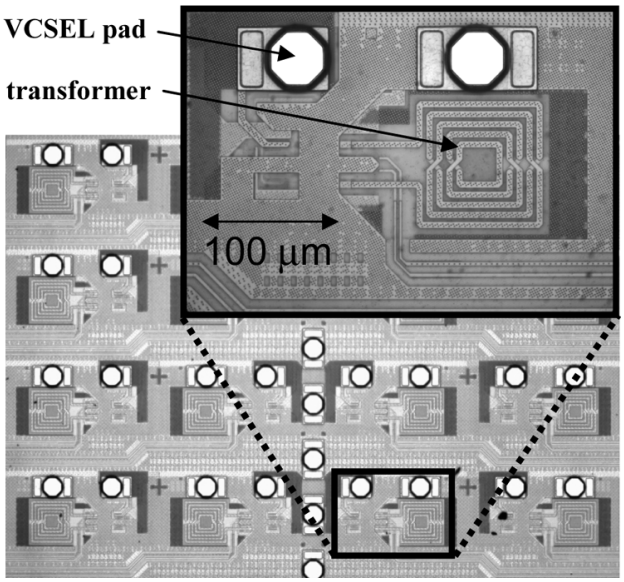

Fig. 3. Photomicrograph of the 4 $\times$ 4 transmitter array with an expanded view of a single channel illustrating the VCSEL attachment pads and multi-metal layer transformer.

A micrograph of the transmitter block, along with a close up of one of the array elements, is presented in Fig. 3. Clearly visible in Fig. 3 are the VCSEL microbump connections along with the multi-metal layer transformer. Although the lasers are fully isolated on the VCSEL chip, the cathodes of all of the diodes share a common ground connection on the CMOS chip.

The LDD circuits consist of a differential predriver stage coupled to a single-ended transconductance amplifier (TCA) output stage, as shown in Fig. 4. The input of the predriver is dc coupled and features a floating, 100-Ω differential termination. The dc-coupled design of the LDD circuits imposes no data coding restrictions, while the floating input impedance allows for either ac or dc-coupled differential data inputs. The predriver was designed to accept nominal 250 $\mathrm{mV_{pp}}$ differential input signals. The power supply for the LDD circuits is split into a 1.8-V supply for the predriver and a 2.7-V supply for the output TCA stage. This dual-supply design minimizes the power consumption of the drivers by allowing the predriver to be powered by a lower supply voltage while using a higher supply voltage for the output stage to accommodate the forward operating voltage of the VCSEL. Transformer peaking is utilized in the output stage of the predrivers to improve the transition times of the large peak to peak output voltages. The TCA output stage has a controllable preemphasis that is applied only to the falling edges of the VCSEL modulation current to improve the optical fall times of the transmitter. This circuit, dubbed fall time compensation (FTC), yields more symmetrical transmitter eye-diagrams, especially at data rates above 10 Gb/s. Details of a similar, earlier generation LDD circuit are presented in [17]. The area occupied by the driver circuits accounts for only $\sim$35% of the 250 μm $\times$ 350 μm cell footprint.

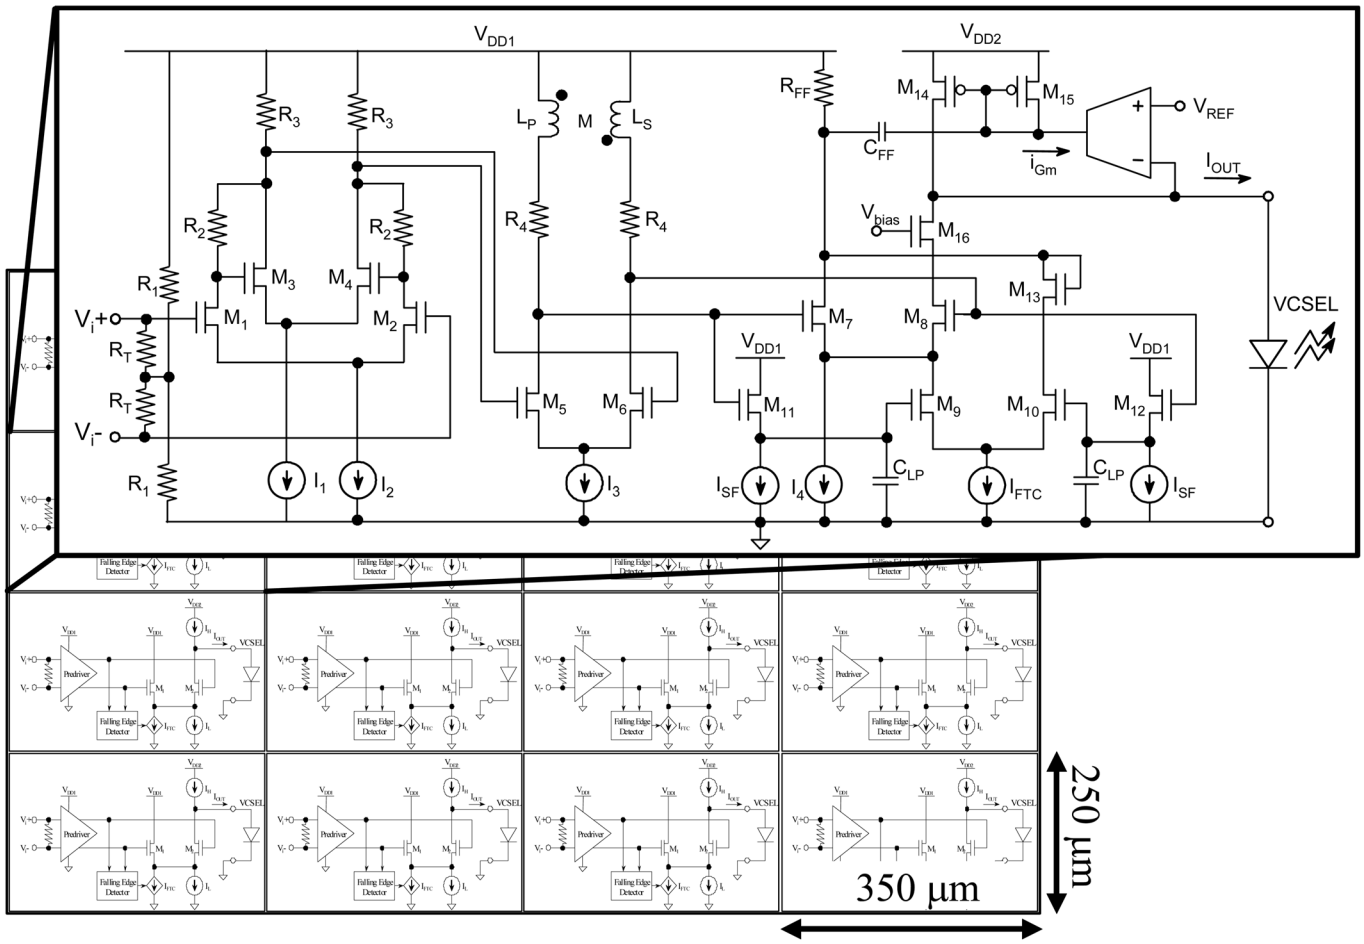

Fig. 4. 4×4 array of transmitters with a circuit schematic of one of the array elements.

The core receiver circuits consist of a transimpedance amplifier followed by a five-stage limiting amplifier. The RX cores are arrayed in the same $4 \times 4,250 \times 350\mu$m floor plan as the LDDs. The output buffers for the RX channels are located at the edges of the receiver block. Differential striplines, further shielded on their sides using dense via fences to connect the top and bottom transmission line ground planes, interconnect the core and buffer circuits. The geometry and design of these enclosed stripline structures are presented in [18]. Due to the physical layout of the chip, the core-to-buffer transmission lines for the inner channels are ~ 350 μm longer than those of the outer channels, a difference that has a noticeable impact on performance that will be further discussed in Section VI. In order to minimize switching-induced power supply noise for the sensitive receiver core circuits, the cores and output buffers are powered by separate power supplies.

The full receiver amplifier circuit appears in Fig. 5. The input transimpedance stage is a fully differential, ac-coupled version of the modified common-gate amplifier described in [19]. A single-channel test circuit of a similar TIA successfully operated at data rates as high as 25 Gb/s [20]. The input ac coupling capacitors enable the front-end to be fully differential by isolating the dc bias point of the input stage from the PD bias circuit. The capacitors are sized to provide a low-frequency cutoff of ~ 30 MHz, a reasonable compromise balancing capacitor area with power penalty due to baseline wander, assuming 8b/10b data encoding. The capacitors are interdigitated vertical parallel plate structures in which the vertical "plates" are formed by fingers of metal on four of the 8 available metal layers connected by dense 1-D via arrays, similar to [21]. These capacitors offer a relatively high capacitance per unit area with a low parasitic capacitance to the substrate, making them ideal for ac-coupling high-speed signals. The RF/analog-tailored back-end metal of the 8RF process is also exploited in the TIA through the inclusion of multi-metal-layer peaking inductors applied at the input (series) and load (shunt) of the input stage. Fig. 6 is a photo micrograph of the receiver block with an expanded view of one of the receiver cores that shows the coupling capacitors and peaking inductors.

The receiver limiting amplifier (LA) consists of five cascaded gain stages as shown in Fig. 5. The last four stages are identical and are based on a modified Cherry-Hooper design [22]. The first stage of the LA shares a similar architecture, but also includes an active circuit that completes the offset cancellation (OC) feedback loop surrounding the limiting amplifier. The OC circuit is required to correct for transistor threshold variations in the differential gain stages that can result in a small imbalance in the input stages being converted to a large output offset through the high gain of the receiver amplifier chain. The OC loop samples and feeds back the average voltages of the lim-

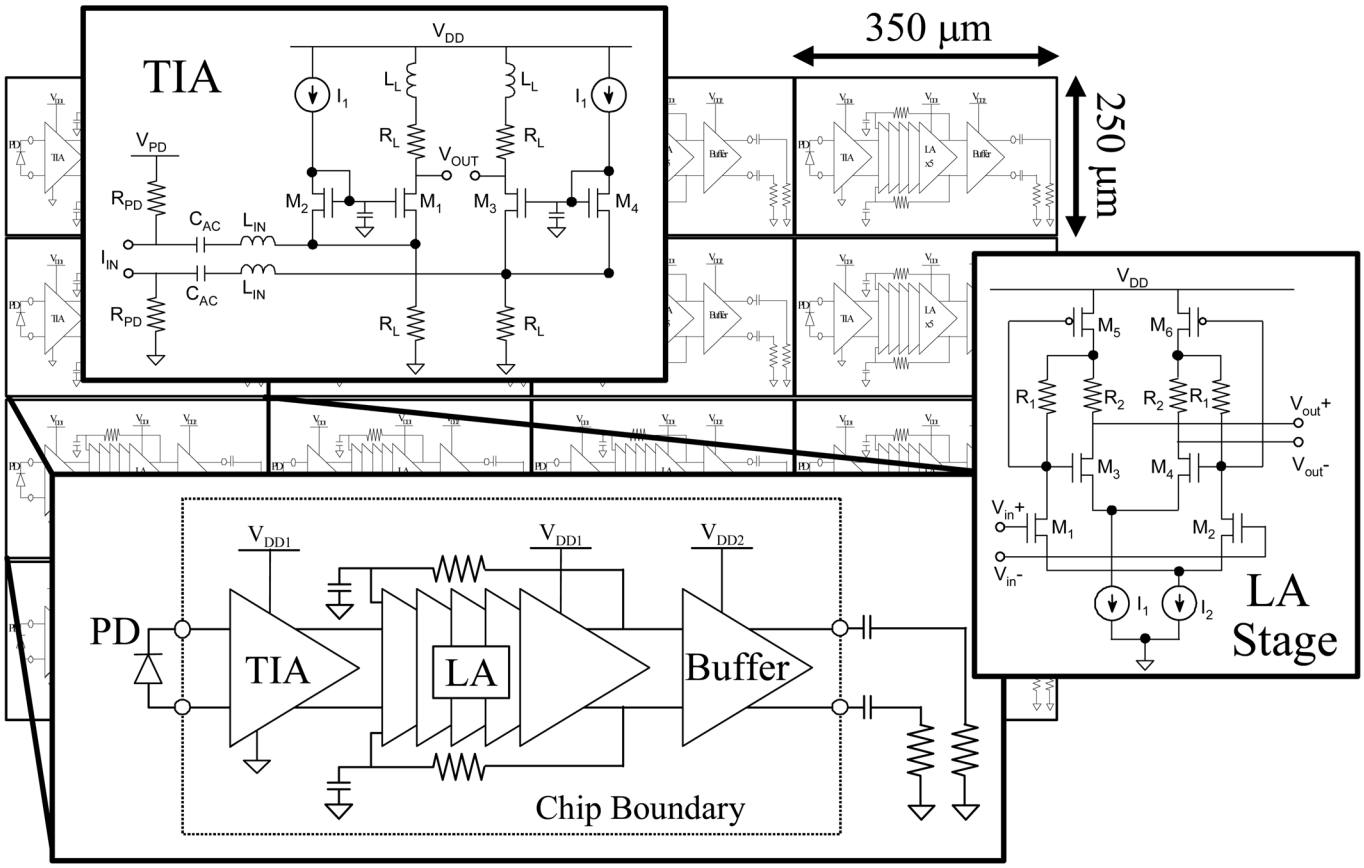

Fig. 5. Circuit schematic of the complete receiver amplifier with insets providing details of the TIA and LA circuits.

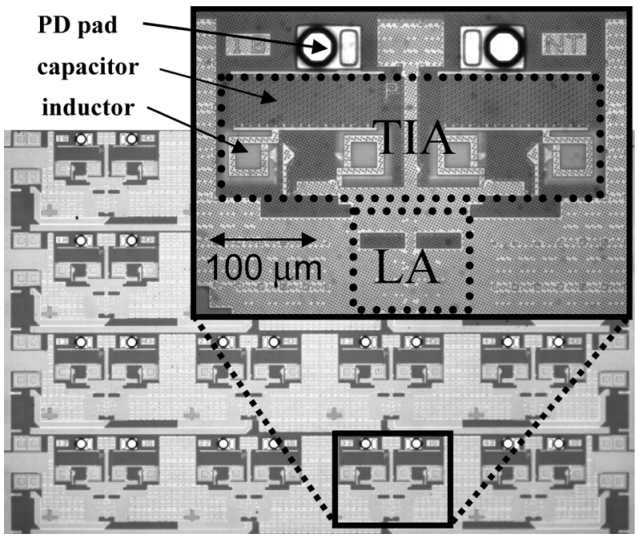

Fig. 6. Photomicrograph of the 4×4 RX array with an expanded view of one of the channels illustrating the PD attachment pads, integrated coupling capacitors, and TIA peaking inductors.

iting amplifier's differential outputs to the circuit in the first LA stage that adjusts the output offset of the first stage to minimize the overall LA output offset. The RC time constant of the OC circuit was chosen such that the response-time of the loop is below the low-frequency cutoff of the receiver. Careful design of the offset cancellation circuit is essential for proper receiver operation and is particularly important to ensure uniform performance for all of the channels within an array due to the random nature of transistor threshold variation. Following the LA is an open-drain differential pair included to drive the transmission lines that connect the limiting amplifiers to the output buffers located on the perimeter of the receiver block. Peaking inductors are again utilized at the input and load of the output buffers to help compensate for the capacitance of the large devices in this stage that are required to drive an off-chip, ac-coupled 50 Ω load.

The TIA and LA circuits account for only ∼35% of the cell area. The remaining cell area is primarily devoted to local decoupling capacitors for the core power supply and the photodiode bias supply. Metal-oxide-semiconductor (MOS) capacitors provide decoupling capacitances of 7.8 pF for the PD bias and 104 pF for the power supply at each receiver core. An additional 250 pF of decoupling capacitance is shared between the four cores comprising each row of the receiver block. Finally, the power supplies of the output buffer circuits are decoupled with 21 pF at each buffer.

IV. Optoelectronic Device Arrays

The VCSEL and PD arrays, which were designed and fabricated at Agilent Laboratories, both operate through substrate emission/detection at a nominal wavelength of 985 nm. This feature allows each array element to be equipped with an integrated lens etched into the backside of the semiconductor

substrate with a lens-to-device alignment accuracy better than ± 2 μm. The OE devices and associated integrated substrate lenses are arranged on the same 250-μm $\times$ 350- μm pitch as the IC channels, but a 62.5-μm staircase offset is applied to each row in the 350-μm direction. This offset allows the 2-D array of OE devices to be coupled to a one-dimensional array of waveguides on a 62.5-μm pitch [13]. The 4×4 OE arrays were produced by subdicing larger 4×12 arrays originally fabricated to support a different packaging concept for separate transmitter and receiver assemblies [2]. The base material of the oxide-confined VCSELs is GaAs, while the mesa-structure PDs are formed on an InP substrate. Both OE substrates are semi-insulating, and the individual array elements of both the laser and PD arrays are fully isolated from each other. Further details of the VCSEL and PD growth and fabrication can be found in [23].

We have built transmitter assemblies with various sizes of VCSEL and PD devices to explore the associated performance and packaging tradeoffs. The range of VCSEL diameters used in transceiver assemblies spans 5 to 9 μm, whereas receivers incorporating photodiodes of three different diameters: 35, 45, and 55-μm, have been characterized. However, the data presented in Section VI is for a transceiver assembled with a uniform array of VCSELs with an oxide aperture diameter of 7 μm and a PD array with elements of two different sizes: the inner devices have a 35-μm diameter, while the outer diodes have a 45-μm diameter.

The 7-μm diameter VCSELs are optimized for operation at 70 °C at data rates up to 20 Gb/s. The lasers are typically operated at a bias current of $\sim$5.5 mA, corresponding to a current density of < 14 kA/cm^2, and achieve a small signal -3 dB bandwidth of 15 GHz under these conditions. The VCSEL series resistance is 130 Ω, with a forward operating voltage of 2.05 V. The differential QE of the VCSELs is 25%, corresponding to a slope efficiency of 0.32 mW/mA, and they have demonstrated clearly open 20 G/s eye-diagrams at 70 °C [24].

The PDs have a measured responsivity of 0.67 A/W, and are typically operated at a bias voltage of -1.5 to -2.5 V. With a -2 V bias, the capacitance is 110 fF and 150 fF for the 35 and 45-μm diameter devices, respectively. Of this total capacitance, the pads account for $\sim$35 fF. The series resistance of the photodiodes, extracted from the forward *I–V* characteristic, is less than 20 Ω. The bandwidth of the photodiodes was obtained from impulse response measurements performed with a mode-locked Ti:Sapphire laser emitting 2-ps pulses at a wavelength of 985 nm [2]. The measured bandwidth was *RC*-limited even for devices with diameters as small as 30 μm, and is greater than 20 GHz for both of the device sizes used in the transceiver assembly reported here.

V. Transceiver Assembly

Transceiver Optochips are assembled by flip-chip soldering 4×4 VCSEL and PD arrays to the CMOS ICs. The gold-plated contact pads of the OE arrays are bonded to the IC pads using eutectic AuSn solder (80% Au, 20% Sn). The AuSn solder is predeposited on the IC bond pads using a two step wafer-level process. The completed 8" CMOS wafers first undergo a NiAu plating step, followed by the deposition of a thick layer of AuSn solder ($\sim$5 μm). The melting temperature of eutectic AuSn is 278 °C during reflow and > 400°C after reflow, enabling the sequential bonding of the two OE device arrays to the transceiver IC. To assemble a complete Optomodule, a transceiver Optochip is subsequently flip–chip attached to a high-density and high-speed surface laminar circuit (SLC) carrier [24]. The solder used in this step is SnPb (63% Sn, 37% Pb) with a melting temperature of 183 °C. Additional details of the Optochip, Optomodule, and Optocard assembly processes can be found in [13] and [14].

VI. Performance Characterization

More than 40 transceiver Optomodules have been assembled using the process described in Section V. After assembly, Optomodules are attached to test cards and wirebond connections made for all of the power and control signals. The transceiver Optomodules exhibit uniform performance, with the most significant variations arising from using OE arrays with devices of different sizes. As aforedescribed in Section IV, the data in this paper were obtained for a transceiver Optomodule with 7-μm diameter VCSELs and 35/45 μm diameter photodiodes. In general, the performance of the transmitter block is similar for assemblies using VCSELs with active-area diameters ranging from 6 to 8 μm. All testing was performed using a $2^7 - 1$ pseudorandom bit stream (PRBS), chosen since it is readily available on all test equipment and is a reasonable substitute test pattern for the 8b/10b data encoding that the Optochip was designed to support. Finally, the voltage sources used to power the Optochip were adjusted to overcome the series resistance in the IC power distribution networks to yield the desired voltages at the core circuitry. The resulting elevated supplies would add an additional $\sim$20% to the measured power consumption, but this contribution has been discounted since the parasitic voltage drop is not an inherent effect and could be eliminated by optimizing the IC power network. In the sections that follow, the voltages and powers reported are for the core circuitry.

The Optochip was designed to be further packaged into an Optomodule, including attaching a heat sink to the backside of the transceiver chip. Using reasonable assumptions for expected ambient temperature and airflow, thermal modeling of the complete package predicts that the IC temperature can be held to < 50 °C. Measurements on fully powered and operating Optomodules have validated this assumption for IC temperature. However, unless otherwise noted, the characterization data that follows was taken with the Optochip operating at $\sim$35 °C (measured at the backside of the transceiver IC). Further details of the thermal performance of complete Optomodule transceiver packages are addressed in [25].

A. Transmitter Section

The transmitter block was powered with a 1.8-V supply for the predrivers and a 2.7-V supply for the output stage. Under these operational conditions, each transmitter channel consumes 73 mW, with the power dissipation split roughly equally between the predriver and output stage. Although its effect at data rates below 10 Gb/s is relatively minor, the FTC circuit was fully enabled for all channels in the data that follows. In addition, the two modulation control voltages, that set the high current level

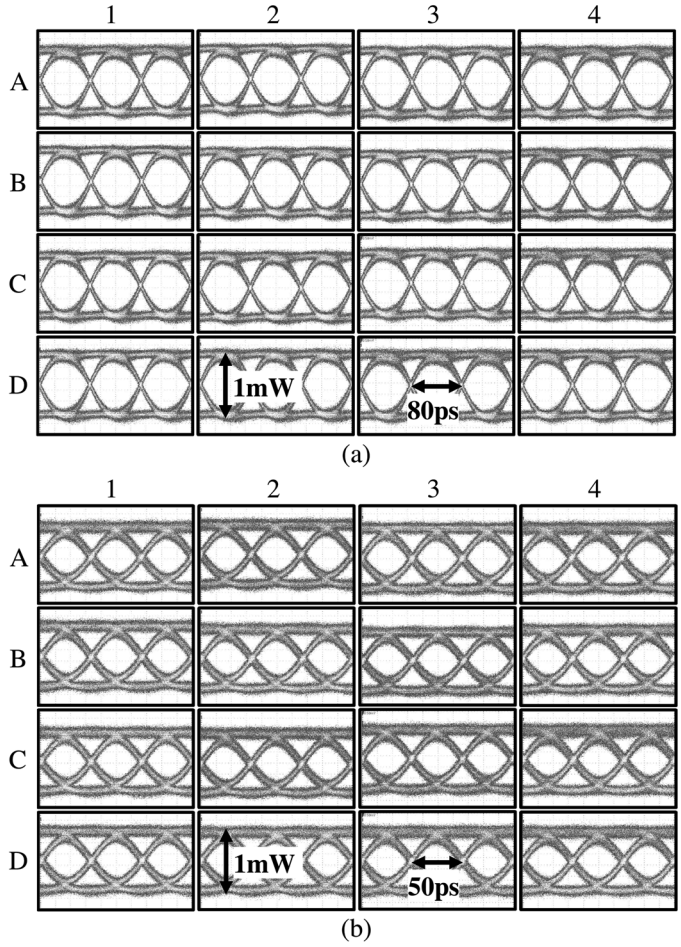

Fig. 7. Eye-diagrams at (a) 12.5 Gb/s and (b) 20 Gb/s for all 16 transmitter channels.

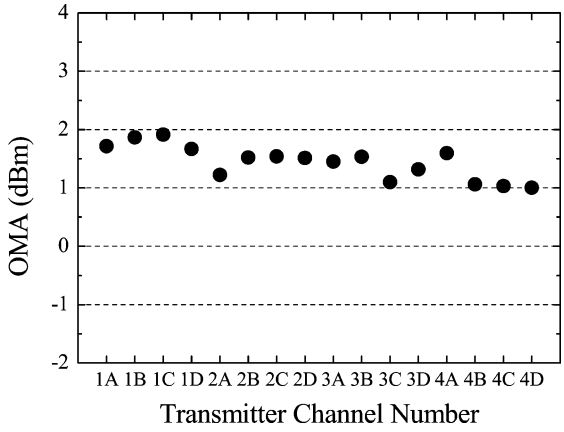

Fig. 8. Measured OMA demonstrating 1-dB power uniformity for all 16 transmitter channels.

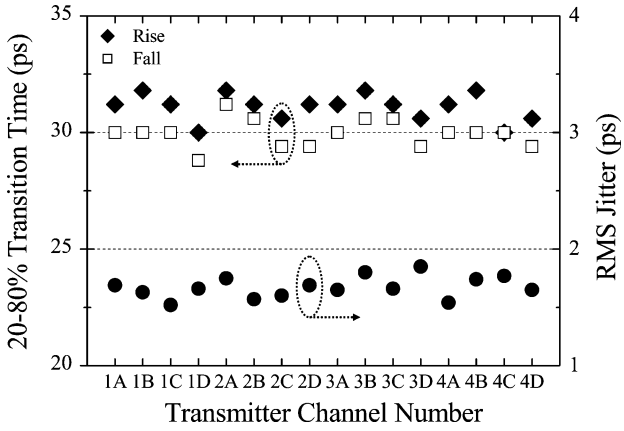

Fig. 9. Measured rise/fall times and RMS jitter for all 16 TX channels.

and depth of modulation, were set to the same levels for all channels: no per-channel optimization of control or supply voltages was necessary or performed.

High-speed characterization of each transmitter channel was undertaken using a differential coplanar microwave probe to apply the input signals and a cleaved 62.5 μm multimode fiber to collect the optical outputs. This fiber type was chosen to maximize the amount of light that was collected while still maintaining compatibility with the high-speed photodiode that was used, a Newport D-25 XR. The -3 dB bandwidth of this unamplified detector is specified to be 17 GHz, and a 20 GHz sampling head was used to display the signals for eye-diagram measurements. The single-ended input data amplitude was set to 500 mVpp, and Fig. 7 presents the eye-diagrams of all 16 transmitter channels at data rates of 12.5 and 20 Gb/s. The extinction ratio (ER) of the transmitters was targeted to be > 4 dB for these data sets.

As described in Section I, the VCSEL lenses were designed for efficient coupling to optical waveguides using a two-lens optical system. As a result, direct butt-coupling to lensed or cleaved 50- or 62.5-μm diameter multimode fiber results in excess coupling loss. Therefore, in order to more accurately measure the output power of the transmitter, a cleaved 100- μm core fiber was used to collect as much of the emitted light as possible. The average transmitted power is bounded between 0 and $+1$ dBm for all channels. The extinction ratio, measured from the transmitter eye diagrams, shows a similar tight distribution between 4 and 5 dB. Current versions of data communication standards such as Fibre Channel [26] and Ethernet [27] have moved away from specifying link budgets in terms of average power and ER, and instead use optical modulation amplitude (OMA). OMA, defined as the difference in power between the logical "1" and "0" levels, has gained favor because in an ac-coupled system, peak-to-peak modulated power is the critical parameter. Additionally, by moving away from an ER specification, there is more freedom for component manufacturers to set the laser bias and modulation to simultaneously meet output power, jitter, and laser eye safety requirements. Fig. 8 is a plot of the measured OMA for each of the transmitter channels. Consistent with the tight distributions of average power and ER, the transmitter OMA ranges from $+1$ to $+2$ dBm.

Fig. 9 shows the rise/fall times and RMS jitter extracted from 10-Gb/s eye-diagrams using the measurement setup described earlier. As expected from the uniform and symmetric eye diagrams of Fig. 7, the rise and fall times are nearly equal and the jitter values are similar for all of the 16 channels. No significant degradation of the transmitters: in output power, rise/fall times, or jitter, was observed when the backside IC temperature was raised from 35 to $\sim$60 °C.

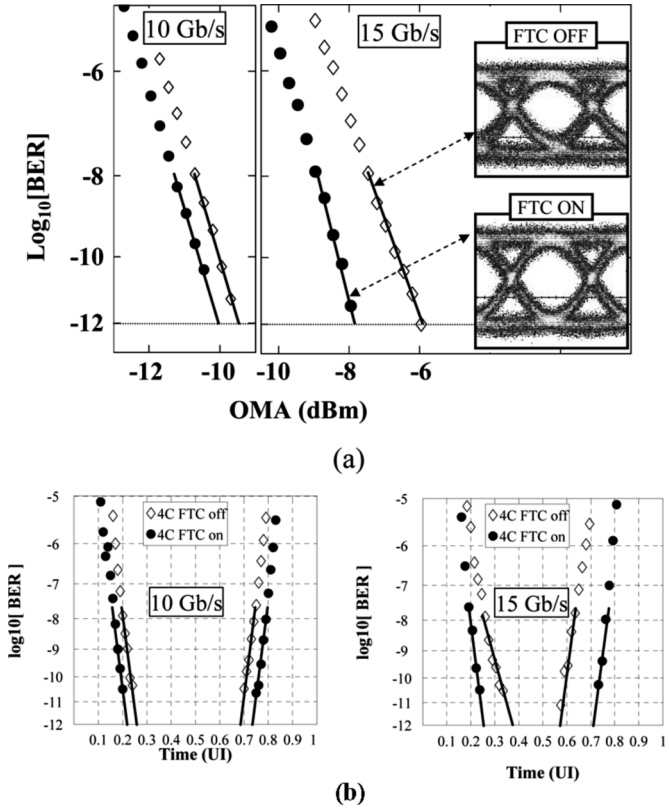

Fig. 10. Effect of fall-time compensation (FTC) at 10 and 15 Gb/s: (a) reference receiver sensitivity measurements; (b) timing margin measurements. Solid symbols: FTC enabled, open symbols: FTC disabled.

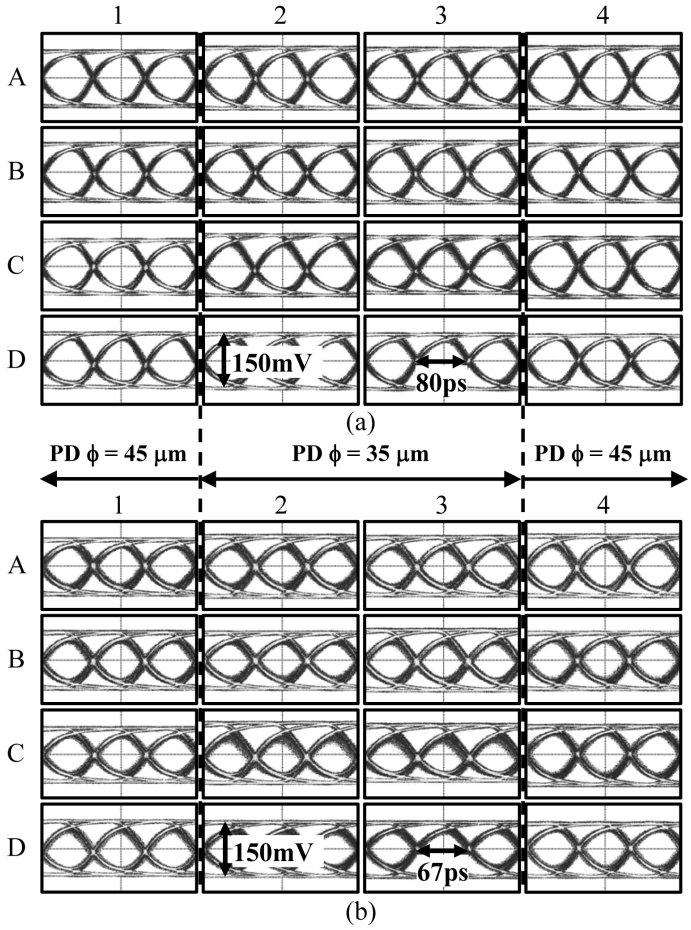

Fig. 11. Single-ended receiver electrical output eye diagrams (full Optochip links) for all 16 RX channels at (a) 12.5 Gb/s and (b) 15 Gb/s. Columns 2 and 3 are comprised of photodiodes with a 35 μm diameter, while columns 1 and 4 have 45 μm diameter devices.

The effect of the FTC circuit was quantified by measuring sensitivity and timing margin for a typical transmitter channel with FTC enabled and disabled and comparing the results. The reference receiver consisted of a Discovery Semiconductor 401HG receiver front-end (PD with TIA) followed by a low-noise linear post amplifier. The sensitivity curves obtained at 10 and 15 Gb/s are presented in Fig. 10. Enabling the FTC circuit results in a sensitivity improvement of approximately 2 dB at 15 Gb/s, with a corresponding increase of temporal eye opening of 0.27 unit interval (UI), or 18 ps. The effect of the FTC circuit at 10 Gb/s is less pronounced, with improvements of $\sim$0.5 dB and 0.1 UI (10 ps) in the sensitivity and timing characteristics, respectively. The inset eye-diagrams of Fig. 10 were obtained with the Newport D-25 XR photodetector described above and not with the reference receiver used for the sensitivity measurements.

B. Receiver Section

As detailed in Section III, the receiver is powered with three separate supplies: the core power supply, the photodiode bias, and the output buffer supply. Supplying 1.8-V to the core and buffer supplies yields a reasonable tradeoff between performance and power dissipation. Under these operational conditions, the per-channel power consumption is 44 mW for the core circuits (TIA and LA), and 18 mW for the output buffers, for a total of 62 mW/channel. The PD bias supply was set to 2.5 V for all measurements, but could be decreased to 2 V at the cost of a slight reduction in sensitivity at the highest data rates (12.5 and 15 Gb/s). Based upon previous measurements on single-channel test versions of the receiver circuits, the receiver gain is 87 dB Ω with a -3 dB bandwidth of 6.6 GHz [28].

One of the transmitter channels in a separate Terabus Optochip was used as the optical source for receiver characterization. Therefore, all of the data collected on the receiver sections represent the performance of complete transceiver links. Eye diagrams obtained for all channels of a transceiver with 35/45 μm photodiodes at data rates of 12.5 and 15 Gb/s are presented in Fig. 11. The 12.5-Gb/s eye diagrams were taken at an input OMA of -4.5 dBm, whereas for the 15-Gb/s data the OMA was increased to -3.5 dBm. All 16 channels in the receiver section produce clearly open eyes at data rates up to 15 Gb/s. The outputs of the receivers are fully limiting for incident OMA above -14 dBm, and the differential output amplitude ranges from 275–350 $\mathrm{mV_{pp}}$ for all 16 channels.

The rise and fall times measured on eye-diagrams obtained at 10 Gb/s are presented in Fig. 12. The rise/fall time data reveals a systematic difference between the inner (2 and 3) and outer rows (1 and 4). Based upon the diameter of the PDs used in the receiver block, it would be expected that the inner rows, comprised of 35- μm diameter devices, would exhibit faster rise

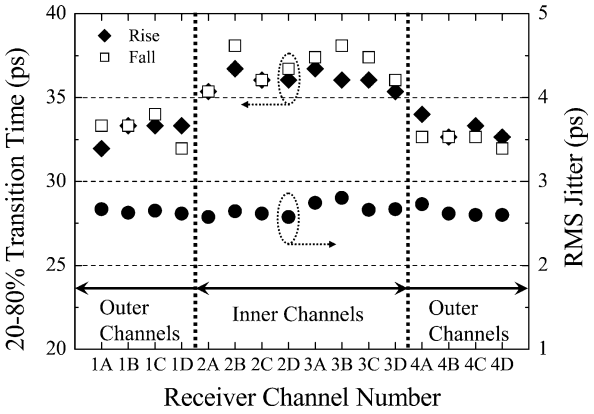

Fig. 12. Measured rise/fall times and RMS jitter for the 16 RX channels.

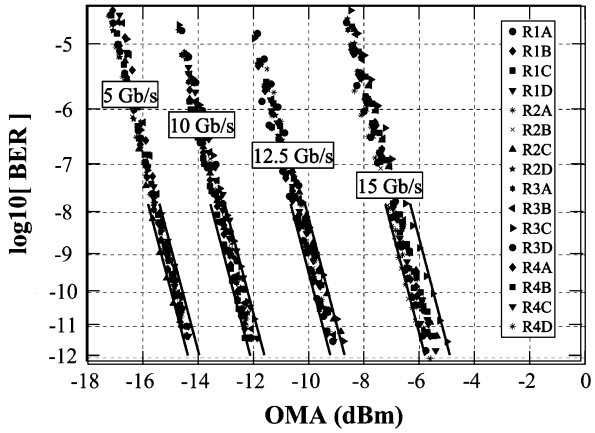

Fig. 13. Receiver sensitivity curves obtained on all channels at data rates of 5, 10, 12.5, and 15 Gb/s.

and fall times compared to the outer rows that are made up of 45- μm diameter devices. The measured data of Fig. 12 is in direct contradiction to this expectation: the outer channels have shorter transition times compared to the inner channels.

The difference between the inner and outer channel performance can be traced to the physical layout of the receiver chip. The transmission lines interconnecting the core and buffer circuits are 350 μm longer for the inner channels compared to the outer channels (cf., Fig. 19). As discussed in Section III, the transmission lines were designed to be fully shielded stripline structures to minimize interchannel crosstalk. This design decision resulted in the signals being routed on one of the middle metal layers, with a nominal thickness of 0.32 μm [18]. In addition, due to the dense channel packing of the core circuits and the fact that the transmission line design occupies all of the metal layers, forming a complete routing blockade, the width of the complete transmission line structure was constrained to be 16 μm. To make the differential impedance of the transmission lines $\sim 100\Omega$, the receivers in first-generation transceiver chips incorporated 2 μm wide signal traces [18]. However, the first-generation receiver arrays only operated at data rates up to 12 Gb/s [29], in contrast to single-channel receiver test chips (with a direct connection between core and buffer: no transmission line) that were shown to operate up to 17 Gb/s [30]. Because the single-channel receivers performed significantly faster than the array versions, with the only difference between them being the transmission lines, the decision was made to increase the width of the transmission lines to reduce their resistance and correspondingly their high-frequency loss. In the second-generation transceivers reported here, the width of the signal traces in the differential striplines was increased by 50%, to 3 μm, in an attempt to increase the maximum receiver operating bit-rate. This approach was successful, as evidenced by the open 15 Gb/s eye diagrams of Fig. 11. However the difference in rise and fall times between the inner and outer channels indicates that the transmission lines still impact the receiver performance, particularly as they traverse longer distances.

The receiver sensitivity characteristics measured at 5, 10, 12.5, and 15 Gb/s for all 16 receiver channels are shown in Fig. 13. The spread in measured sensitivities is very low: < 0.5 dB for rates up to 12.5 Gb/s and < 0.8 dB at 15 Gb/s, indicating uniform performance for all receiver channels and no dependence on photodiode diameter. At a $\mathrm{BER} = 10^{-12}$, the worst-case receiver sensitivity, expressed in OMA, is $-11.5, -8.5$, and -4.7 dBm for bit rates of 10, 12.5, and 15 Gb/s, respectively. As with the transmitters, the two power domains in the receiver block were operated under the same conditions with no per-channel optimization. However, some degradation in receiver performance is incurred if the receivers are operated at elevated temperatures. At a backside IC temperature of 70 °C, the receiver sensitivity is reduced by 1 dB at 5 Gb/s, 2 dB at 10 Gb/s, and 2.6 dB at 12.5 Gb/s compared to measurements at 35 °C (such as those of Fig. 13). At 70°C and 12.5 Gb/s, the temporal eye opening is also reduced by 0.1 UI (8 ps). Fortunately, at the expected operating chip temperature for fully packaged transceiver modules, 50 °C, the degradation in sensitivity compared to 35 °C operation is less pronounced: 0.2 dB at 5 Gb/s, and 0.7 dB at both 10 and 12.5 Gb/s. These characteristics are consistent with prior high-temperature characterization of single-channel receiver circuits that indicated both the gain and bandwidth of the receiver amplifiers are reduced at elevated operating temperatures [31].

The slightly wider distribution of measured sensitivity at 15 Gb/s compared to the lower data rates can be attributed to one of the receiver channels (2C) that exhibits higher output offset compared to the other channels. This characteristic is noticeable in the eye-diagrams produced by this channel, particularly at 15 Gb/s (Fig. 11), through a low crossing level of $\sim 37\%$. However, it should be noted that the spread in receiver performance has been substantially improved in the second generation chips compared to the first-generation [29]. The first-generation receivers showed a spread in sensitivity of ~ 1 dB for measurements taken at 10 Gb/s, and a greater number of channels showed eye-diagram crossings deviating from 50% compared to the second-generation circuits. The improvement in the later generation is attributed to a strengthened offset cancellation circuit that incorporates 50% more gain in the active feedback loop. The higher-gain offset cancellation circuit is able to correct for higher levels of input offset, resulting in a greater number of balanced receiver channels. In addition, as evidenced by the tight distributions, the transmission line effect mentioned above does not impact the measured receiver

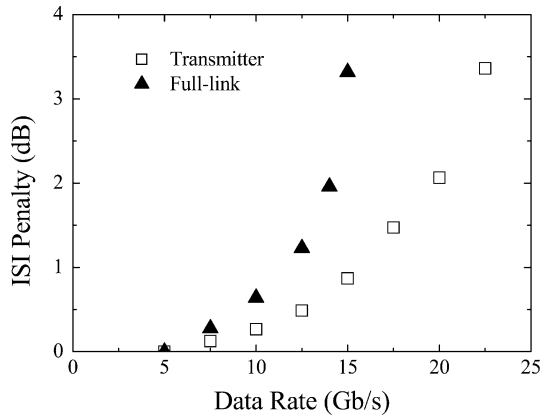

Fig. 14. Measured ISI penalty as a function of data rate for typical channels: TX and full-link.

sensitivity. This is as expected since the output of the limiting amplifier is a digital signal: the logical "1"/"0" decision has already been made earlier in the LA chain.

C. Transceiver Performance Summary

In order to better illustrate the speed capability of the transceiver Optochip, the intersymbol interference (ISI) power penalty is plotted as a function of data rate for typical channels in Fig. 14. The ISI penalty is defined as the amount of additional optical power that would be required to overcome deterministic eye closure due to limited bandwidth. The transmitter ISI penalty was characterized using the reference photodiode described earlier, while the full-link ISI penalty was obtained from the receiver output with input provided by an Optochip transmitter channel. As seen in Fig. 14, the transmitters show very little eye closure over the tested data rates, with less than 1 dB of ISI penalty at 15 Gb/s. On the other hand, the full Optochip links show good performance up to 12.5 Gb/s, with less than 1.5 dB of ISI penalty. At 15 Gb/s, the link ISI penalty exceeds 3 dB, indicating that complete links are not viable at higher rates. The difference between the transmitter and full-link ISI penalties indicate that the bandwidth of the receivers is the primary factor that limits the maximum operational speed of the Optochips.

It is important to note that the ISI penalty increases less rapidly as a function of data rate than the measured decrease in receiver sensitivity due to the limiting function of the receivers. The ISI penalty is a measure of the eye closure of the fully limiting digital receiver output, whereas receiver sensitivity is determined by the effective eye opening at the point in the receiver amplifier chain where the logical decision is made. Fig. 14 directly illustrates the primary challenge in building multimode optical links operating at $>$ 10 Gb/s: producing receiver amplifiers (either limiting or linear) that simultaneously provide high gain and high bandwidth. Extending the speed capability of LDDs and VCSELs appears to be a more tractable problem, with recent reports of 985-nm devices operating at 35 Gb/s [32] and 850-nm VCSELs capable of 30 Gb/s [33].

TABLE I
SUMMARY OF TRANSMITTER AND RECEIVER PERFORMANCE

Parameter	TX	RX	Unit
Power dissipation/channel	73	62	mW
Data rate, typ.	12.5	12.5	Gb/s
Data rate, max.	20	15	Gb/s
t_r/t_f, 20-80%, max.	32	38	Ps
TX OMA, min.	1		dBm
TX Avg. power, min.	0		dBm
TX extinction ratio, min.	4		dB
RX OMA Sensitivity, min. 12.5 Gb/s, BER=10^{-12}		-8.5	dBm
RX output amplitude, min.		270	mV_{pp-d}

TABLE II
SUMMARY OF TRANSCEIVER PERFORMANCE

	Channel Data Rate (Gb/s/ch)			
Parameter	**10**	**12.5**	**15**	**Unit**
Aggregate bit rate	160	200	240	Gb/s
Power/channel	135	135	135	mW
Power consumption	13.5	10.8	9	mW/Gb/s
Area efficiency	9.8	11.8	14	$Gb/s/mm^2$
ISI penalty, max.	0.75	1.5	3.5	dB
Supported Link Budget	12.5	9.5	5.7	dB

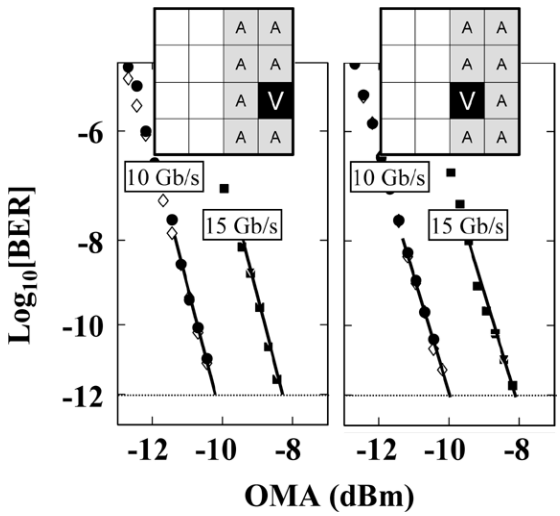

Fig. 15. Measured transmitter crosstalk characteristics at 10 and 15 Gb/s for a typical (a) outer and (b) inner victim channel, when operating alone (solid symbols) and surrounded by seven neighboring aggressor channels (open symbols).

Table I provides a summary of the key transmitter and receiver performance parameters. Table II summarizes the transceiver Optochip (full TX-RX link) performance. The parameters in Tables I and II are for operation at 35 °C. As discussed in Section VI-B, the receiver sensitivity is 0.7 dB lower at 50 °C, resulting in a corresponding reduction of the supported link budget at this expected operational temperature.

VII. CROSSTALK CHARACTERIZATION

Interchannel crosstalk is an obvious concern for dense transceiver chips that incorporate 2-D arrays of driver circuits and high-gain receiver circuits. Characterization of three types of channel-to-channel crosstalk was undertaken: transmitter-to-receiver, transmitter-to-transmitter, and receiver-to-receiver.

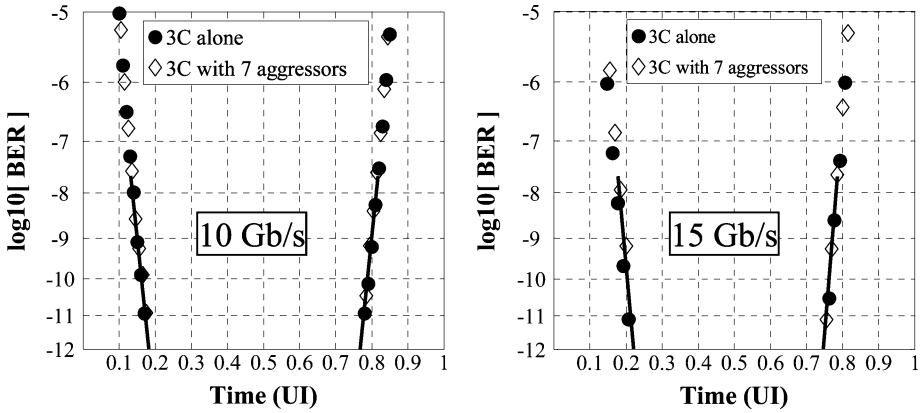

Fig. 16. Timing margin measured on a typical transmitter channel at 10 and 15 Gb/s when operating alone (solid circles) and in the presence of seven surrounding aggressor channels (open diamonds).

Degradation in receiver sensitivity and timing margin were used to quantify crosstalk effects: these two parameters directly capture reductions in both vertical and horizontal eye opening.

A. Transmitter-to-Receiver Crosstalk Characterization

Transmitter to receiver crosstalk was evaluated by performing several loop-back sensitivity tests. The sensitivity at 10 Gb/s of one of the receiver channels located closest to the transmitter section was measured repeatedly with its input provided by different transmitter channels on the same transceiver Optochip. The measured sensitivity characteristics were identical for all configurations, and were indistinguishable from measurements of the receiver channel taken with the transmitter section completely turned off. These measurements indicate that the transmitter to receiver crosstalk is negligible, a conclusion that is not surprising based upon the relatively large distance between the transmitter and receiver sections ($\sim$500 μm).

B. Transmitter Interchannel Crosstalk Characterization

Transmitter interchannel crosstalk was quantified at the Optomodule level by measuring the sensitivity characteristics of a reference receiver with its optical input provided by the transmitter channel under test: the "victim" channel. The baseline receiver characteristics, obtained with only the victim channel being driven with high-speed differential electrical inputs, were then compared to the receiver sensitivity curves measured with up to seven surrounding aggressor channels being driven with independent $2^7 - 1$ PRBS electrical data. The reference receiver used for this measurement is described in Section VI-A. Transceiver crosstalk was evaluated at the Optomodule level to allow the simultaneous probing of eight transmitter channels, which is not feasible on Optochips. The seven independent data streams for the aggressor channels were supplied by seven independent test boards incorporating PRBS generator chips that operate at data rates up to 15 Gb/s. The aggressor data generators all shared a clock that was separate from the clock for the victim channel. The crosstalk penalties measured in this asynchronous fashion should represent worst-case values, especially compared to typical applications where all channels are operated synchronously with a common clock. Transmitter crosstalk characterization was performed on multiple channels at data rates of 10 and 15 Gb/s. Typical data obtained on two transmitter channels, one an inner channel, the other an outer channel, is presented in Fig. 15 for both data rates. The measured reference receiver characteristics are identical for the victim channel operating alone or in the presence of seven surrounding aggressor channels, indicating a negligible transmitter interchannel crosstalk penalty.

Transmitter timing margin was also measured for victim channels operating alone and in the presence of seven surrounding aggressor channels. Fig. 16 presents the results of these measurements on a typical channel at data rates of 10

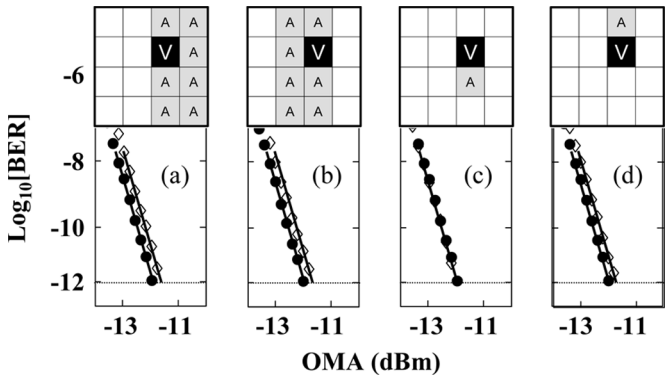

Fig. 17. Crosstalk penalty measurements at 10 Gb/s of an inner receiver channel for varying configurations of aggressor channels.

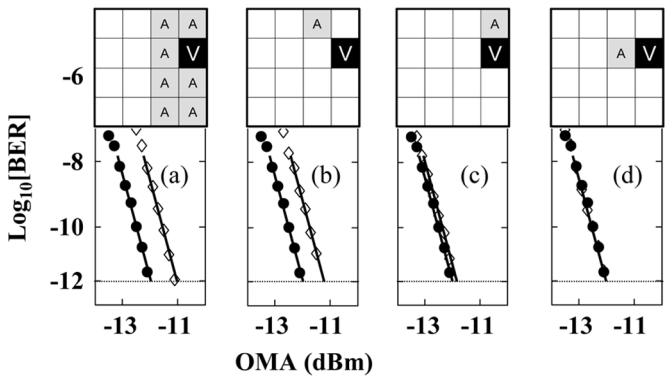

Fig. 18. Measured crosstalk characteristics at 10 Gb/s of an outer receiver channel for various aggressor channel configurations.

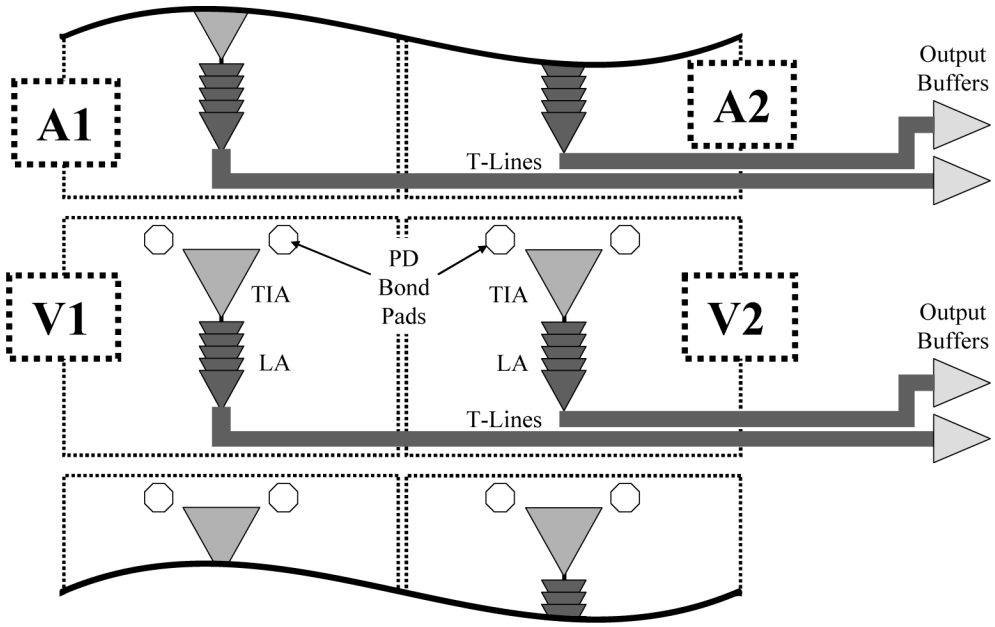

Fig. 19. Simplified block diagram of the RX layout illustrating the signal routing in relation to victim and primary aggressor channels.

and 15 Gb/s. As with the sensitivity measurements, the timing margin of victim channels was unaffected by the operation of neighboring channels, providing further confirmation that transmitter channel-to-channel crosstalk is insignificant in the transceiver Optochips.

C. Receiver Interchannel Crosstalk Characterization

Interchannel receiver crosstalk was investigated at the Optochip level through sensitivity and timing margin characterization of a channel under investigation (the victim channel), surrounded by up to seven neighboring aggressor channels. For these crosstalk studies, control measurements were first performed in which the victim channel was the only element of the array receiving an optical signal. To assess the impact of channel-to-channel crosstalk, the original measurements were then compared to the results of subsequent measurements on the victim channel when differing combinations of its neighboring channels were excited with optical data.

The optical sources for the aggressor channels were provided by a custom-built 12-channel, 985-nm transmitter with data inputs supplied by the 12-channel test station described in [4]. The aggressor channels were operated at 10.3125 Gb/s with a $2^{31} - 1$ PRBS pattern. Although the low-frequency cutoff of the receivers is too high to support the long run lengths of this pattern, it was acceptable for driving the aggressor channels since for this measurement they were only required to generate digital outputs, not necessarily run error-free. The aggressor channels all shared a clock that was separate from the clock for the victim channel. As with the transmitter crosstalk characterization, the receiver crosstalk penalties measured in this asynchronous manner should represent worst-case values compared to expected operating conditions. In addition, the OMA presented to the aggressors was set to approximately -4 dBm: more than 7 dB above the receiver sensitivity of the victim channel at 10 Gb/s. This represents a beyond worst-case configuration for point-to-point parallel data links in which all channels likely share similar routing and therefore should experience similar loss. However, since the receivers provide fully limited digital outputs for incident OMA greater than -14 dBm, approximately 3 dB below the 10-Gb/s receiver sensitivity $(\mathrm{BER} = 10^{-12})$, the crosstalk results were found to be independent of the aggressor incident optical power level.

The optical signals for crosstalk testing were introduced to the receivers using two cleaved ribbon fibers comprised of 12 fibers each, spaced at a 250-μm pitch. Four fibers from each ribbon were concurrently coupled into four receiver channels so that a total of up to eight channels could simultaneously receive optical data. The victim channel within this 2×4-element array was supplied data from a transmitter of a separate transceiver Optochip, while the aggressor channels surrounding the victim received data from up to seven channels of the 12-channel transmitter described above.

Crosstalk testing identified a consistent pattern of power and timing margin penalties based upon the location of the victim channel within the RX block. Typical crosstalk results for inner RX channels are presented in Fig. 17. The inner channels exhibit very low crosstalk penalties, less than 0.3 dB, when operated with seven surrounding aggressor channels. As shown in Fig. 17(a) and (b), similar sensitivity characteristics are obtained with the neighboring aggressors located in either the inner or outer columns.

The outer receiver channels exhibit markedly different crosstalk penalties compared to the inner channels. The lower three channels on each side of the receiver block have crosstalk penalties slightly less than 1 dB when operated with seven surrounding aggressor channels. Typical results obtained on one of these outer channels are shown in Fig. 18. A similar crosstalk penalty is incurred for seven surrounding aggressor channels as for only one particular aggressor, as illustrated in Fig. 18(a) and (b). In other measurements for a single aggressor located in all other possible locations surrounding the victim channel, a negligible crosstalk penalty was measured as evidenced in Fig. 18(c) and (d).

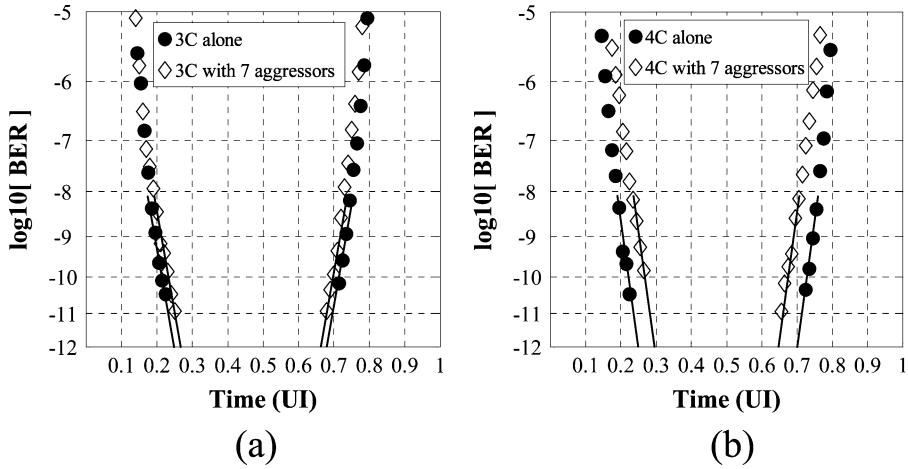

Fig. 20. Receiver timing margin curves at 10 Gb/s for a typical (a) inner and (b) outer victim channel, when operating alone (solid symbols) and surrounded by seven neighboring aggressor channels (open symbols).

The layout of the receiver section, part of which is shown as a block diagram in Fig. 19, may explain the measured crosstalk penalties. Fig. 19 depicts several channels on the right side of the receiver chip; the transceiver layout is completely symmetrical so the left-hand channels have identical circuitry and routing and therefore crosstalk characteristics. The data in Fig. 17 was taken on a victim channel corresponding to the array element labeled V1 in Fig. 19. The sensitivity curves obtained on V1 indicated that one particular aggressor channel, dubbed the primary aggressor, and corresponding to A1 in Fig. 19, contributed nearly the entire observed sensitivity penalty.

The higher crosstalk penalties measured for outer channels are also dominated by a primary aggressor channel. The data of Fig. 18 were taken on a victim channel corresponding to V2 in Fig. 19. Once again, nearly the entire sensitivity penalty measured for V1 when completely surrounded by operating aggressor channels can be attributed to the same primary aggressor: A1. The larger penalty measured for V2 compared to V1 is attributed to the longer interaction length between the transmission lines of A1 and the input stage of V2. A1's transmission lines run in close proximity to the inputs of V2 along the entire 350-μm length of the cell.

The receiver interchannel crosstalk penalties are believed to arise due to coupling between the sensitive input of the victim TIA and the large signals ($\sim$1 V_{pp-d}) traveling on the core-to-buffer transmission lines of the primary aggressor. Switching-induced noise introduced at the input of the victim TIA due to this transmission line coupling directly degrades the sensitivity of the victim RX. Another observation that supports this explanation is that the topmost outer receiver channels incurred no sensitivity penalty when surrounded by seven aggressors. These channels do not have primary aggressor channels: no aggressor transmission lines are located near their TIAs.

For completeness, receiver sensitivity measurements were also performed at 5 Gb/s on typical inner and outer channels to ensure that the intrinsically lower sensitivity of the receivers at 10 Gb/s, due to their limited bandwidth, was not masking additional penalties. The measured sensitivity penalties at 5 Gb/s were in agreement with the 10-Gb/s data, indicating no crosstalk dependence on bit-rate.

The effect of interchannel crosstalk on receiver timing margin was also assessed with the same measurement setup used for the sensitivity measurements. However, for the timing measurements, the input optical power to the victim receiver channel was held constant and the BER measured as a function of the temporal sampling point of the receiver output signal. The results are shown in Fig. 20 for inner and outer channels at an input power level 1 dB above the receiver sensitivity point at 10 Gb/s ($\mathrm{BER} = 10^{-12}$). As with the sensitivity measurements, the input optical power for each of the aggressor channels was set 7 dB above the receiver sensitivity. Consistent with the sensitivity measurements on the inner channels that exhibited minimal crosstalk penalties, the timing margin of the inner channels is not significantly impacted by the presence of aggressors. For the outer channels, crosstalk causes a horizontal eye-closure of $\sim$0.1 UI (equal to 10 ps at a bit rate of 10 Gb/s). At higher input powers the temporal eye-closure is less pronounced because the receiver output is more strongly limited.

Due to physical constraints, only the victim channel could be electrically probed during the crosstalk measurements reported here. Therefore, the electrical outputs of the aggressor channels were unterminated: the open circuit at the probe pads completely reflects each aggressor's output back into its output buffer circuit. This effect may lead to an overestimation of the crosstalk penalty compared to actual operation in which all channels are properly terminated. Additional measurements at the Optomodule level, where the aggressor channels can be probed and terminated, will be conducted in the future to quantify the effect of output termination on receiver crosstalk penalty.

In summary, our investigation into receiver crosstalk at the Optochip level demonstrated a maximum crosstalk penalty of 1 dB for outer channels with a concurrent reduction in horizontal eye opening of $\sim$10 ps at 10 Gb/s. Inner channels were shown to incur a maximum sensitivity penalty of 0.3 dB with no significant reduction in temporal eye opening.

VIII. Conclusion

We have demonstrated a single-chip, CMOS-based optical transceiver that supports aggregate bidirectional data rates up to 240 Gb/s through 16 receiver and 16 transmitter channels. The

complete transceiver assembly utilizes the hybrid integration of dense, 2-D, substrate emitting/detecting OE arrays with low-power CMOS amplifier circuits and occupies an area of 17 mm^2. All 16 TX channels exhibited clearly open eyes at data rates up to 20 Gb/s with modulated optical powers greater than 1 mW. The 16 receivers were fully operational for data rates up to 15 Gb/s with measured worst-case OMA sensitivities of $-11.5, -8.5$, and -4.7 dBm $(\mathrm{BER} = 10^{-12})$ for bit rates of 10, 12.5, and 15 Gb/s, respectively. Receiver interchannel crosstalk was identified as the dominant source of crosstalk, resulting in a worst-case sensitivity penalty less than 1.0 dB. The total core power consumption for all of the transmitters (73 mW/ch) and receivers (62 mW/ch) in the Optochip is 2.15 W, corresponding to 9 mW/Gb/s/link at 15 Gb/s/ch. The transceivers achieve an area efficiency of 14 Gb/s/mm^2 that is unprecedented for parallel optics and are the first single-chip transceivers that are capable of channel rates above 10 Gb/s. The compact chip-like packaging approach readily supports scaling to higher numbers of transmitter and receiver channels to enable future wide and fast board-level optical data buses through integrated waveguides.

ACKNOWLEDGMENT

The authors would like to acknowledge the outstanding technical work performed by individuals formerly with Agilent Laboratories in producing the photodiode and VCSEL arrays during the first phase of the Terabus program, in particular G. Trott, M. Tan, and C. K. Lin. The efforts of D. Kucharski and D. Guckenberger, now with Luxtera, Inc., in producing the circuit designs for the first phase of the project are also appreciated. The authors also would like to thank A. Rylyakov, P. Pepeljugoski, L. Schares, and M. Taubenblatt for useful discussions in support of the project. Finally, the authors would like to thank T. LaMothe at IBM Burlington, VT, for foundry support.

REFERENCES

[1] A. Benner, M. Ignatowski, J. Kash, D. Kuchta, and M. Ritter, "Exploitation of optical interconnects in future server architectures," *IBM J. Res. Dev.*, vol. 49, pp. 755–775, 2005.

[2] L. Schares *et al.*, "Terabus: Tb/s-class card-level optical interconnect technologies," *IEEE J. Sel. Topics Quantum Electron., Special Issue on Optoelectron. Pack.*, vol. 12, pp. 1032–1044, Sept./Oct. 2006.

[3] SNAP12 Multi Source Agreement Group May 2002 [Online]. Available: http://www.avagotech.com, SNAP12 12 Channel Pluggable Optical Module MSA Specifications,'' Rev. 1.1, Accessed on 1/08/2008

[4] D. M. Kuchta, Y. H. Kwark, C. Schuster, C. Baks, C. Haymes, J. Schaub, P. Pepeljugoski, L. Shan, R. John, D. Kucharski, D. Rogers, M. Ritter, J. Jewell, L. A. Graham, K. Schrödinger, A. Schild, and H.-M. Rein, "120-Gb/s VCSEL-based parallel-optical interconnect and custom 120-Gb/s testing station," *J. Lightw. Technol.*, vol. 22, pp. 2200–2212, Sep. 2004.

[5] D. M. Kuchta, D. Kucharski, Y. H. Kwark, and R. John, "120 Gb/s VCSEL-based parallel optical links," in *CPT 2005 Eighth Int. Symp. Contemp. Photon. Technol.*, Tokyo, Japan, Jan. 2004, vol. 08 G-4, pp. 169–172.

[6] L. A. Buckman-Windover, J. N. Simon, S. A. Rosenau, K. S. Giboney, G. M. Flower, L. W. Mirkarimi, A. Grot, B. Law, C.-K. Lin, A. Tandon, R. W. Gruhlke, H. Xia, G. Rankin, M. R. T. Tan, and D. W. Dolfi, "Parallel optical interconnects > 100 Gb/s," *J. Lightw. Technol.*, vol. 22, pp. 2055–2063, Sep. 2004.

[7] B. E. Lemoff, M. E. Ali, G. Panotopoulos, G. M. Flower, B. Madhavan, A. F. J. Levi, and D. W. Dolfi, "MAUI: Enabling fiber-to-the-processor with parallel multiwavelength optical interconnects," *J. Lightw. Technol.*, vol. 22, pp. 2043–2054, Sep. 2004.

[8] B. E. Lemoff, M. E. Ali, G. Panotopoulos, E. de Groot, G. M. Flower, G. H. Rankin, A. J. Schmit, K. D. Djordjev, M. R. T. Tan, A. Tandon, W. Gong, R. P. Tella, B. Law, and D. W. Dolfi, "500-Gbps parallel WDM optical interconnect," in *Proc. 55th Electron. Comput Technol. Conf.*, Orlando, FL, 2005.

[9] D. V. 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. Lightw. Technol.*, vol. 19, pp. 1093–1103, Aug. 2001.

[10] M. B. Venditti, E. Laprise, J. Faucher, P.-O. Laprise, J. E. A. Lugo, and D. V. Plant, "Design and test of an optoelectronic-vlsi chip with 540-element receiver-transmitter arrays using differential optical signaling," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, pp. 361–379, Mar./Apr. 2003.

[11] A. V. Krishnamoorthy, K. W. Goosen, L. M. F. Chirovsky, R. G. Rozier, P. Chandramani, W. S. Hobson, S. P. Hui, J. Lopata, J. A. Walker, and L. A. D'Asaro, "16×16 VCSEL array flip-chip bonded to CMOS VLSI circuit," *IEEE Photon. Technol. Lett.*, vol. 12, pp. 1073–1075, Aug. 2000.

[12] J. Trezza, H. Hamster, J. Iamartino, H. Bagheri, and C. DeCustatis, "Parallel optical interconnects for enterprise class server clusters: needs and technology solutions," *IEEE Commun. Mag.*, vol. 41, pp. S36–S42, Feb. 2003.

[13] F. E. Doany, C. L. Schow, C. Baks, R. Budd, Y.-J. Chang, P. Pepeljugoski, L. Schares, D. Kuchta, R. John, J. A. Kash, and F. Libsch, "160-Gb/s bidirectional parallel optical transceiver module for board level interconnects using a single-chip CMOS IC," in *Proc. 57th ECTC*, Reno, NV, May/Jun. 2007, pp. 1256–1261.

[14] F. E. Doany, C. L. Schow, J. A. Kash, C. Baks, R. Budd, P. Pepeljugoski, D. Kuchta, R. Dangel, F. Horst, and B. Offrein, "Terabus: A 160-Gb/s bidirectional board level optical data bus," in *Proc. IEEE LEOS 21st Ann. Meet.*, Lake Buena Vista, FL, Oct. 2007.

[15] R. Dangel, U. Bapst, C. Berger, R. Beyeler, L. Dellmann, F. Horst, B. Offrein, and G.-L. Bona, "Development of a low-cost low-loss polymer waveguide technology for parallel optical interconnect applications," in *Tech. Dig. IEEE/LEOS Summer Top. Meet.*, San Diego, CA, Jun. 2004.

[16] L. F. Miller, "Controlled collapse reflow chip joining," *IBM J. Res. Dev.*, vol. 13, no. 3, pp. 239–250, 1969.

[17] D. Kucharski, Y. H. Kwark, D. M. Kuchta, K. T. Kornegay, M. R. T. Tan, C.-K. Lynn, and A. Tandon, "A 20 Gb/s VCSEL driver with pre-emphasis and regulated output impedance in 0.13 μm CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, Feb. 2005.

[18] D. Guckenberger, C. Schuster, Y. H. Kwark, and K. Kornegay, "On-chip crosstalk mitigation for densely packed differential striplines using via fence enclosures," *Electron. Lett.*, vol. 41, pp. 412–414, Mar. 2005.

[19] D. Guckenberger, J. D. Schaub, and K. T. Kornegay, "A dc-coupled low-power transimpedance amplifier architecture for Gb/s communication system applications," in *Dig. RFIC Symp.*, June 2004, pp. 515–518.

[20] C. L. Schow, L. Schares, R. A. John, L. S. Fischer, and D. Guckenberger, "A 25-Gb/s transimpedance amplifier in 0.13 μm CMOS," *Electron. Lett.*, vol. 42, pp. 1240–1241, Oct. 2006.

[21] R. Aparicio and A. Hajimiri, "Capacity limits and matching properties of integrated capacitors," *IEEE J. Solid-State Circuits*, vol. 37, pp. 384–393, Mar. 2002.

[22] E. Cherry and D. Hooper, "The design of wide-band transistor feedback amplifiers," *Proc. Inst. Elect. Eng.*, vol. 110, pp. 375–389, Feb. 1963.

[23] C. K. Lin, A. Tandon, K. Djordjev, S. W. Corzine, and M. R. T. Tan, "High-speed 985 nm bottom-emitting VCSEL arrays for chip-to-chip parallel optical interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 13, no. 5, pp. 1332–1339, Sep./Oct. 2007.

[24] Kyocera SLC Technologies, Yasu, Shiga, Japan (2008). [Online]. Available: http://www.kyocera-slc.co.jp/english/

[25] F. E. Doany, C. L. Schow, C. W. Baks, D. M. Kuchta, P. Pepeljugoski, L. Schares, R. Budd, F. Libsch, R. Dangel, F. Horst, B. J. Offrein, and J. A. Kash, "160 Gb/s bidirectional polymer waveguide board-level optical interconnects using CMOS-based transceivers," *," IEEE Trans. Adv. Packag.*, to be published.

[26] *Fibre Channel Physical Interfaces (FC-PI-2) Standard* 2005, FC-PI-2 INCITS 404.

[27] *IEEE Standard 802.3ae* Jun. 2002.
[28] C. L. Schow, L. Schares, S. J. Koester, G. Dehlinger, R. John, and F. E. Doany, "A 15-Gb/s, 2.4 V optical receiver using a Ge-on-SOI photodiode and a CMOS IC," *IEEE Photon. Technol. Lett.*, vol. 18, pp. 1981–1983, Oct. 2006.
[29] C. L. Schow, F. E. Doany, O. Liboiron-Ladouceur, C. Baks, D. M. Kuchta, L. Schares, R. John, and J. A. Kash, "160-Gb/s, 16-channel full-duplex, single-chip CMOS optical transceiver," in *Proc. Opt. Fiber Commun. Conf. OFC 2007*, Mar. 25–29, 2007, pp. 1–3, paper OThG4.
[30] O. Liboiron-Ladouceur, C. L. Schow, P. K. Pepeljogoski, F. E. Doany, R. A. John, and J. A. Kash, "A 17-Gb/s, 200-m multimode optical fiber link using CMOS analog ICs and silicon carrier packaging," in *Proc. IEEE LEOS 20th Ann. Meet.*, Montreal, Canada, Oct. 29–Nov. 2 2006, pp. 573–574.
[31] S. J. Koester, C. L. Schow, L. Schares, G. Dehlinger, J. D. Schaub, F. E. Doany, and R. A. John, "Ge-on-SOI-detector/CMOS-amplifier receivers for high-performance optical-communication applications," *J. Lightw. Technol.*, vol. 25, pp. 2046–2058, Jan. 2007.
[32] Y.-C. Chang, C. S. Wang, and L. A. Coldren, "High-efficiency, high-speed VCSELs with 35 Gbit/s error-free operation," *Elecron. Lett.*, vol. 43, pp. 1022–1023, Sept. 2007.
[33] R. H. Johnson and D. Kuchta, "30 Gb/s directly modulated 850 nm datacom VCSELs," in *Postdeadline Paper CLEO/QELS 2008*, San Jose, CA, May 4–9, 2008.

Clint L. Schow received the B.S. degree in electrical engineering in 1994, and the M.S. and Ph.D. degrees in electrical engineering from the University of Texas at Austin in 1997 and 1999, respectively.

In 1999, he joined IBM, Rochester, MN, assuming responsibility for the optical receivers used in IBM's optical transceiver business. From 2001 to 2004, he was with Agility Communications, Santa Barbara, CA, developing high-speed optoelectronic modulators and tunable laser sources for optical communications. In 2004, he joined the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, as a Research Staff Member. He is currently working on parallel optical interconnect technologies and high-speed CMOS circuits for fiber-optic data links.

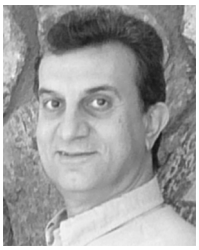

Fuad E. Doany received the Ph.D. degree in physical chemistry in 1984 from the University of Pennsylvania, Philadelphia.

He was a Postdoctoral Fellow with the California Institute of Technology, Pasadena, from 1984 to 1985. He subsequently joined IBM Thomas J. Watson Research Center, Yorktown Heights, NY, where he worked on laser spectroscopy, applied optics, projection displays, and laser material processing for electronic packaging. Since 2000, he has focused on high-speed optical link and systems design, and optoelectronic packaging. He is an author or coauthor of many technical papers and holds more than 40 U.S. patents.

Dr. Doany is a member of the Optical Society of America.

Christian W. Baks received the B.S. degree in applied physics from Fontys College of Technology, Eindhoven, in 2000 and the M.S. degree in physics from the State University of New York (SUNY), Albany, in 2001.

He joined the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, as an engineer in 2001, where he is involved in high-speed optoelectronic package and backplane interconnect design specializing in signal integrity issues.

Young H. Kwark (S'80–M'86) received the B.S.E.E. degree from the Massachusetts Institute of Technology, Cambridge, and the M.S.E.E. and Ph.D. degrees from Stanford University, Stanford, CA.

His work experience as a Research Staff Member with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, includes circuit design for optical links and wireless applications. He is currently involved in package characterization for high-performance computing platforms.

Daniel M. Kuchta (S'85-M'92-SM'97) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer science from the University of California at Berkeley, in 1986, 1988, and 1992, respectively.

He is a Research Staff Member with the Communication Technology Department, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, where he works on high-speed VCSEL characterization, multimode fiber links, and parallel fiber optic link. He is an author or coauthor of 12 patents and more than 50 technical papers.

Jeffrey A. Kash (SM'94-F'05) received the Ph.D. degree in physics from the University of California at Berkeley.

He joined the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, in 1981. Presently, he manages the Optical Link and Systems Design group, which explores advanced optical links for future servers. He is the Principal Investigator for the Terabus project. In past work, he has investigated the optical properties of hot electrons in semiconductors, including the development of Picosecond Imaging Circuit Analysis, which is used by major semiconductor manufacturers today to debug CMOS ICs.

Dr. Kash is a Fellow of the American Physical Society.