256-Channel Bidirectional Optical Interconnect Using VCSELs and Photodiodes on CMOS

David V. Plant, *Member, IEEE,* Michael B. Venditti, *Associate Member, IEEE,* Emmanuelle Laprise, Julien Faucher, *Member, IEEE,* K. Razavi, Marc Châteauneuf, Andrew G. Kirk, *Member, IEEE,* and J. S. Ahearn

***Abstract*—Two-dimensional parallel optical interconnects (2-D-POIs) are capable of providing large connectivity between elements in computing and switching systems. Using this technology we have demonstrated a bidirectional optical interconnect between two printed circuit boards containing optoelectronic (OE) very large scale integration (VLSI) circuits. The OE-VLSI circuits were constructed using vertical cavity surface emitting lasers (VCSELs) and photodiodes (PDs) flip-chip bump-bonded to a 0.35-μm complementary metal–oxide–semiconductor (CMOS) chip. The CMOS was comprised of 256 laser driver circuits, 256 receiver circuits, and the corresponding buffering and control circuits required to operate the large transceiver array. This is the first system, to our knowledge, to send bidirectional data optically between OE-VLSI chips that have both VCSELs and photodiodes cointegrated on the same substrate.**

***Index Terms*—Application-specific integrated circuits (ASICs), integrated optoelectronics (OEs), optical interconnects.**

I. Introduction

TWO-DIMENSIONAL parallel optical interconnects (2-D-POIs) are capable of providing large connectivity between elements in computing and switching systems [1]–[3]. This interconnect technology is inherently scalable due to its 2-D format. Optoelectronic very-large-scale integration (OE-VLSI) circuits that combine the processing power of silicon with the efficiency of GaAs-based emitters and detectors represent an enabling technology [4], [5]. Specifically, using heterogeneous integration techniques, large 2-D arrays of vertical cavity surface emitting lasers (VCSELs) and photodiodes (PDs) can be flip-chip bonded to silicon electronics to provide optical input–output (I/O) to OE-VLSI application specific integrated circuits (ASICs) in addition to existing electrical I/O.

Using this technology, we have demonstrated a 256-channel bidirectional optical interconnect between two printed circuit boards (PCBs). This paper, which describes these results, is organized as follows. In Section II, we describe the VCSEL and PD properties and the heterogeneous integration approach. Section III describes the transmitter and receiver circuits. In Section IV, we describe the architecture of the chip that was designed, fabricated, and successfully operated. Section V is a description of the PCBs that were used not only for chip testing but also for demonstrating the optical interconnect. In Section VI, we describe the optics and optomechanics used to construct the optical interconnect, and in Section VII, we discuss the experimental results obtained for both the chip and the overall interconnect. Section VIII is a conclusion section.

Manuscript received June 12, 2000; revised March 20, 2001. This work was supported in part by BAE Systems under a contract via the DARPA/ARL VLSI-Photonics program, DAAL01-98-C-0074, and supported in part by NSERC and FCAR postgraduate fellowships and the Canadian Institute for Telecommunications Research under the NCE Program of Canada.

D. V. Plant, M. B. Venditti, E. Laprise, J. Faucher, K. Razavi, M. Châteauneuf, and A. G. Kirk are with the Department of Electrical and Computer Engineering, McGill University, Montréal, QC H3A 2A7, Canada.

J. S. Ahearn is with Information and Electronic Systems Integration Inc., BAE Systems, NHQ 4-0196, Nashua, NH 03061-0868 USA.

Publisher Item Identifier S 0733-8724(01)05313-0.

II. OE Devices: VCSELS, PDs, and Heterogeneous Integration

To achieve the OE-VLSI ASIC described in this paper, 2-D arrays of VCSELs and PDs were fabricated on separate substrates and subsequently integrated onto the silicon complementary metal–oxide–semiconductor (CMOS) die. In order to support a compact high-density microoptical interconnect, the VCSELs and PDs were interleaved [6]. We describe in this section: the design and target operating properties of the VCSELs and PDs, the OE device layout geometries, and heterogeneous integration techniques including flip-chip bonding and substrate removal of the interdigitated OE devices. This is the first report, to our knowledge, of the heterogeneous integration of interleaved VCSELs and PDs on a CMOS substrate.

A. VCSEL and PD Design and Specifications

The VCSELs were proton implanted devices designed to operate at 850 nm with threshold currents of $\sim$4.5 mA and slope efficiencies of 0.25 mW/mA [7]. The devices were also designed to be backside-emitting because of the desire to flip-chip bond them to CMOS driver circuits as described in Section II-C. This necessitated removal of the GaAs substrate to minimize absorption of light. To achieve these objectives, VCSELs were fabricated with both the n-contact and the p-contact located on the top surface of the wafer to facilitate electrical contact to the CMOS circuits. Fig. 1 shows a schematic of the VCSEL geometry indicating emission direction after substrate removal and integration to the CMOS. The p-contact was formed above the top distributed Bragg mirror (DBR) and the n-contact was brought to the substrate surface through mesa isolation and ion implantation. Fig. 2 is a photomicrograph of four isolated VCSELs prior to flip-chip bonding and substrate removal; the p-contacts and n-contacts are indicated. The VCSELs in the photograph are on a 125-μm $\times$ 125-μm pitch. Once bonded to the CMOS as per the description given below, the n-contact and DBR became the top (emitting) surface of the VCSEL.

The PDs were p-i-n structures designed to operate with a responsivity of 0.5 A/W. Fig. 3 is a photomicrograph of four isolated PDs on a 125-μm $\times$ 125-μm pitch prior to flip-chip

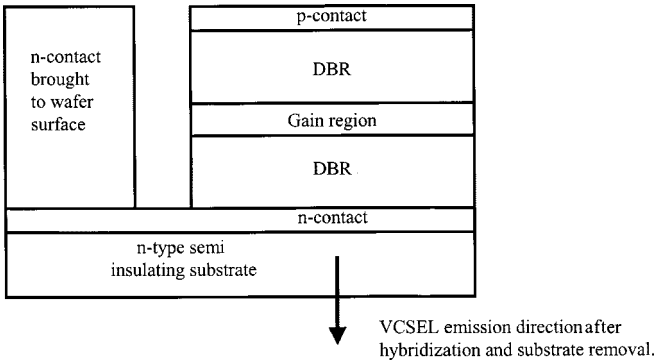

Fig. 1. A schematic of the VCSEL geometry indicating the emission direction after substrate removal and integration to the CMOS chip. Based on the partial transmittivity of the p-contact/DBR, the VCSELs could be probed prior to hybridization with the CMOS to verify functionality at the wafer level.

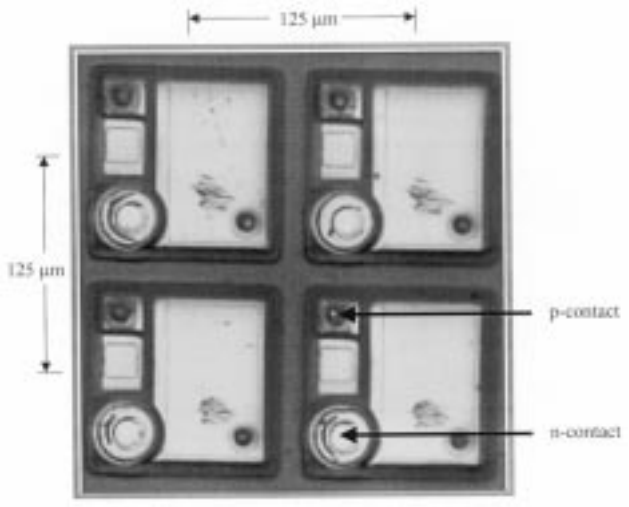

Fig. 2. Photomicrograph of four isolated VCSELs prior to flip-chip bonding. The n-contact and the p-contact are indicated. The scratch marks are a result of wafer-probing prior to flip-chipping.

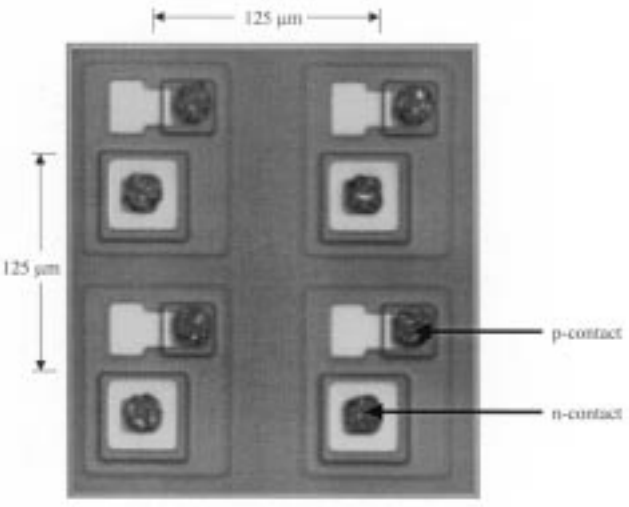

Fig. 3. Photomicrograph of four isolated p-i-n's prior to flip-chip bonding. The n-contact and the p-contact are indicated.

bonding and substrate removal; the p-contacts and n-contacts are indicated. The 2-D PD arrays were fabricated at the wafer level on a 125-μm $\times$ 125-μm pitch and were designed to be flip-chip bonded to the CMOS driver chip described below in Section II-C.

B. Optoelectronic Device Layout Geometry

In order to support a compact point-to-point optical interconnect system, the VCSELs and PDs were interleaved and arranged in a clustered geometry [6]. Fig. 4(a) is a schematic of the free-space microoptic link and Fig. 4(b) shows a schematic of the VCSEL and PD placement requirements. As is indicated in Fig. 4(b), VCSELs and PDs were grouped together in clusters, and, within each cluster, rows of VCSELs and PDs were interleaved. Specifically, a cluster consisted of eight VCSELs and eight PDs arranged in four rows. The pitch of the optoelectronic devices was 125 μm in both the horizontal and vertical directions; therefore, the VCSELs and PDs were on 125-μm horizontal by 250-μm vertical pitch. The complete 256-VCSEL and 256-PD array consisted of 32 clusters arranged in eight rows and four columns, as is shown in Fig. 5. The center-to-center spacing of clusters was 750 μm horizontally and 750 μm vertically. The CMOS was designed to accommodate this OE device pitch and placement.

It is worth noting that both the VCSELs and PDs were fabricated on a 125-μm$\times$ 125-μm pitch at the wafer level in order to be compatible with typical OE device pitches for $1 \times N$ products manufactured for the telecommunications industry. The OE devices, in principle, could be fabricated on a smaller pitch. During heterogeneous integration, OE devices that did not have corresponding contact points on the CMOS chip were removed during the flip-chip bonding and substrate removal processes described in the next section.

C. Heterogeneous Integration and Substrate Removal

The VCSELs and PDs were integrated onto the CMOS driver using flip-chip bonding and substrate removal techniques. The VCSEL flip-chip contact area was 15 μm $\times$ 15 μm and the PD flip-chip contact area was 10 μm $\times$10 μm. The contact areas on the CMOS die for the VCSEL driver and PD receiver were identical to those on the optoelectronic devices.

Heterogeneous integration was accomplished by employing relatively conventional photolithographic processes to deposit and lift off contact metals on the wafers followed by a precision assembly process using a flip-chip bonding tool. In the photolithographic step, a photoresist polymer is first spun out on the wafer and printed with the contact metal pattern and then developed. Indium is then evaporated onto the wafer and the photoresist is lifted off, leaving metal on the contact pads. This process was used for the VCSEL and PD wafers and the CMOS dies. The individual OE device dies were separated by mechanical dicing and then integrated onto a CMOS die using the precision alignment hybridization tool. The VCSEL die was first attached to the CMOS chip followed by dry etching to remove the substrate; integration of the PD die was accomplished next followed by substrate removal. The bonding of the indium metal contacts on the CMOS chip and on the OE devices was accomplished through a combination of force and controlled temperature. The process resulted in electrical isolation of individual OE devices and allowed the interleaving of the VCSEL and PD

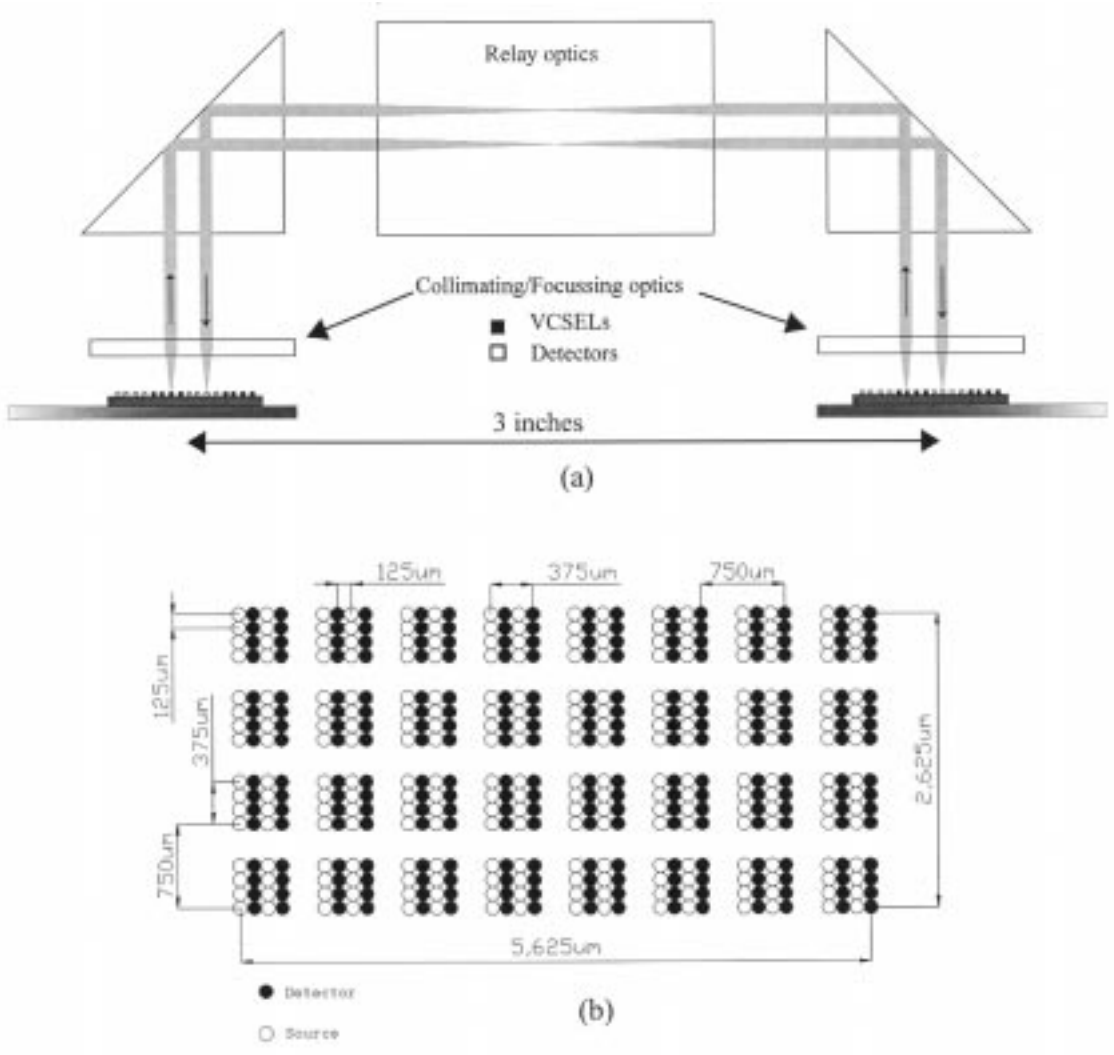

Fig. 4. (a) Schematic of the free-space microoptic link. (b) Schematic of the VCSEL and PD placement requirements.

devices onto a single CMOS die. Although individual dies were used to assemble this generation of OE-VLSI chips, migration of the process to the wafer level is relatively straightforward.

Fig. 6 shows a photomicrograph of four clusters after heterogeneous integration and substrate removal, each cluster consisting of eight VCSELs and eight PDs hybridized to the underlying CMOS chip. This layout geometry proved effective when designing the architecture of the CMOS circuit discussed in Section IV. Fig. 7 is a photograph of the complete OE-VLSI chip after VCSEL and PD integration. Fig. 8(a) shows a group of four clusters with 32 VCSELs biased below threshold, and Fig. 8(b) shows the entire VCSEL array biased above threshold. Using continuous wave measurements, the VCSEL yield after heterogeneous integration was $>$98%. In the following two sections, we describe the transceiver circuits and the CMOS chip architecture which were implemented in this interconnect system.

III. Transmitter and Receiver Circuits

The main objective of the transceiver circuit design was to provide enough flexibility to allow for the successful simultaneous operation of large numbers of transmitters and receivers. Although simulation results described below indicate high-speed operation was achievable, high data rate operation was not a principal design objective. It was expected that the VCSEL, PD, and CMOS characteristics would vary over a large device array; thus, the transceiver designs had to allow for statistical variations in device parameters and had to avoid dependence on parameters specific to both the silicon and the OE process [8], [9]. The transceiver circuits were designed to keep their inherent switching noise generation at a practical minimum, as well as to be immune to the expected presence of substantial amounts of aggregate switching noise generated from a large array of mixed analog and digital circuits.

Given these design objectives, the design of the laser driver was based on current-steering, as outlined in Fig. 9. Specifically, a VCSEL was dc-biased with a current I_{BIAS} to a point above the VCSEL threshold current. Modulation current was provided by the current source I_{MOD} and was steered through either the VCSEL or through an electrical dummy diode load D1, which was implemented as a diode-connected PMOS transistor. Current steering was achieved with switching transistors M1L and M1R and complementary rail-to-rail digital CMOS inputs Vin and Vinb. The polarity of the laser driver circuit was noninverting; thus, when the inputs were logically low (Vin low and

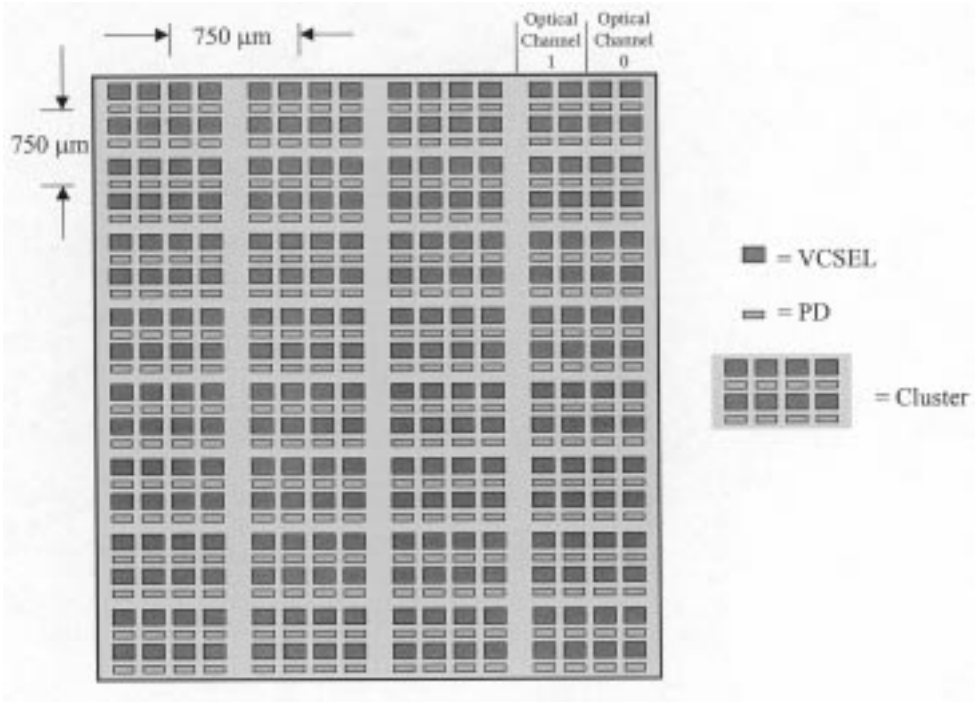

Fig. 5. A full view of the VCSEL and PD clustering design. Optical channels 0 and 1 (described in Section IV) are also indicated. Within an individual cluster and post hybridization, the VCSEL and PD pitch was 125 μm in the horizontal direction by 250 μm in the vertical direction.

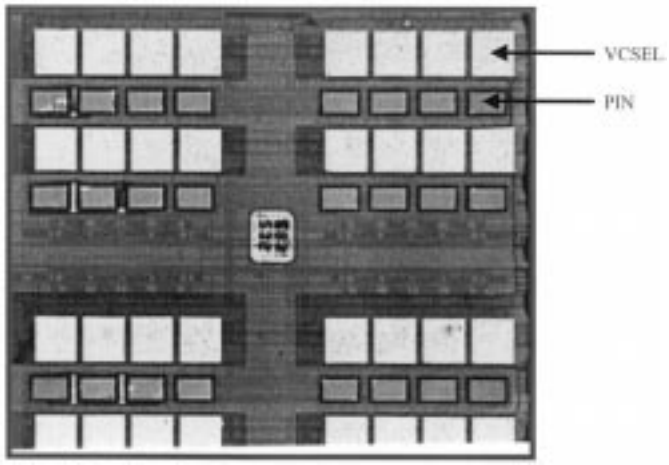

Fig. 6. Four clusters after heterogeneous integration and substrate removal. The VCSEL device was 110 μm $\times$ 112.5 μm and had a 20-μm diameter active region; the p-i-n was 125 μm $\times$ 90 μm and had a 50-μm $\times$ 50-μm active region.

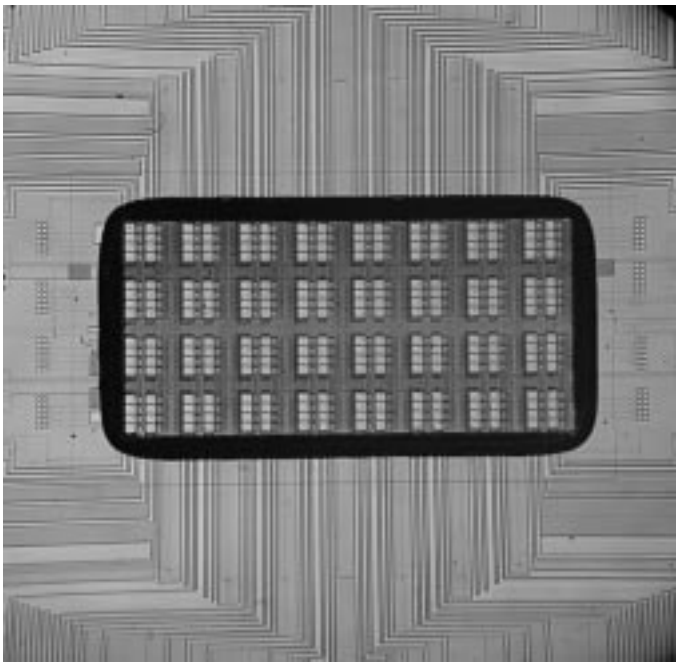

Fig. 7. A complete OE-VLSI chip. The rectangular section located in the middle of the die is the VCSEL and PD array.

Vinb high), the VCSEL was biased at $I_{\mathrm{BIAS}} + I_{\mathrm{MOD}}$ and, therefore, produced a logically high-output power. When the inputs were reversed (Vin high and Vinb low), the VCSEL was biased with only I_{BIAS} and produced a logically low output power.

The current-steering nature of the laser driver allowed the total current drawn from the power supply to remain nominally constant at $I_{\mathrm{BIAS}} + I_{\mathrm{MOD}}$ whether the VCSEL was in a high or low output power state. Power supply current transients could not be completely eliminated due to the mismatch in electrical parameters of the dummy load D1 and the VCSEL, but the approach allowed current transients (dI/dt noise) to be kept to a small fraction of I_{MOD}. The range of currents settable for I_{MOD} and I_{BIAS} was approximately 6 and 12 mA, respectively. The nominal voltage supply was 4.8 V. The power dissipation per laser driver circuit depended on the magnitudes of I_{BIAS} and I_{MOD}, and was estimated to be 86.4 mW in the worst case. An individual transmitter circuit was successfully simulated under worst-case (i.e., largest magnitude) conditions for I_{BIAS} and I_{MOD} at data rates in excess of 1 Gb/s.

The receiver circuit is shown schematically in Fig. 10. It was optically and electrically single-ended and was based on a

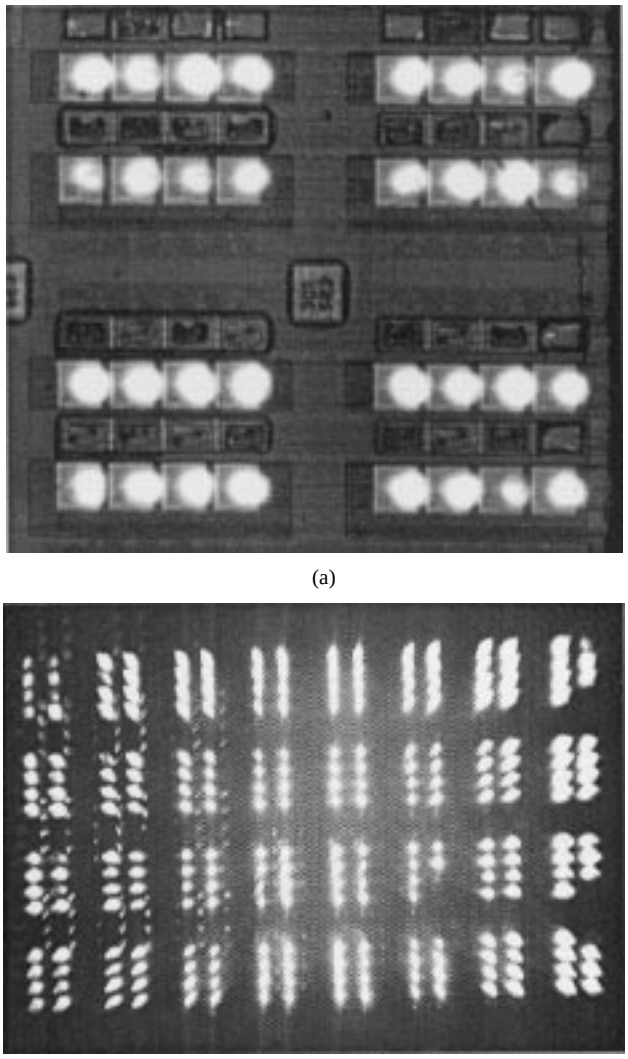

(a)

(b)

Fig. 8. (a) Four clusters with 32 VCSEL biased below threshold. (b) The entire VCSEL array biased above threshold. The distortion is caused by the optical system used to image the 3-mm × 6-mm array simultaneously.

common source transimpedance amplifier (TIA) front end [10], [11]. An offset-control stage was included to compensate for both the dc-coupled nature of working with CMOS amplifier stages and the dc-coupled nature of the optical input. This allowed properties of the receiver such as sensitivity (preamplifier feedback resistance) and the accommodation of various average optical power levels (offset control) to be dealt with independently, providing greater operational flexibility. The final stage of the receiver consisted of a Schmitt trigger that served as a final gain stage for decision-making and provided some hysteresis in its transfer function to help reduce the effects of power-supply switching noise in an array environment [12]. As was the case with the laser-driver circuit, the operation of the receiver was noninverting. If the optical input was logically high, then the receiver output was also logically high, and *vice versa*. Power dissipation for the receiver circuit was dependent on a multitude of operating conditions such as optical input power levels and the bias condition of the offset correction stage. Under typical operating conditions, the power dissipation per receiver was estimated to be 5 mW. An individual receiver

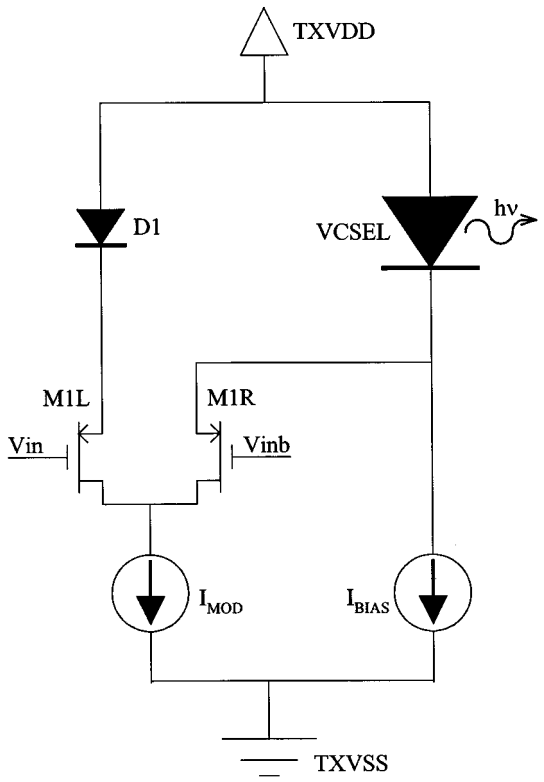

Fig. 9. Laser driver design.

was successfully simulated at 500 Mb/s with input optical power levels of 20 and 42 μW in the maximum-gain setting and with 200 and 800 μW in the minimum-gain setting.

Via the processes described in Section II-B above, each driver circuit was integrated with a VCSEL and each receiver was integrated with a PD; this resulted in a 2-D array of 256 transmitters and receivers. We describe, in the next section, the digital design and control architecture of the CMOS chip. We also provide details on the die size and the fabrication technology used to construct the chip.

IV. CMOS Chip Architecture

The CMOS chip was designed to act as a network interface chip for nodes in a computing network [13], [14]. A 32-bit-wide data bus was a key design goal. To meet these objectives, the chip was designed using 256 pixels, each pixel being comprised of a VCSEL hybridized to a laser driver (a transmitter), a photodiode hybridized to an amplifier (a receiver), and control logic to define the operating mode of the pixel. A schematic of a pixel is shown in Fig. 11.

Referring to Fig. 11, there were four modes of operation for each pixel: mode 1) electrical data could be transmitted optically [E-to-O conversion]; mode 2) optical data could be received electrically [O-to-E conversion]; mode 3) optical data could be received and retransmitted optically via combinational logic which resided between receiver and transmitter circuits [O-to-E-to-O]; and mode 4) received optical data could be latched in a flip-flop before being resent optically or electrically. In this mode, data could be stored to allow for protocols which required synchronization.

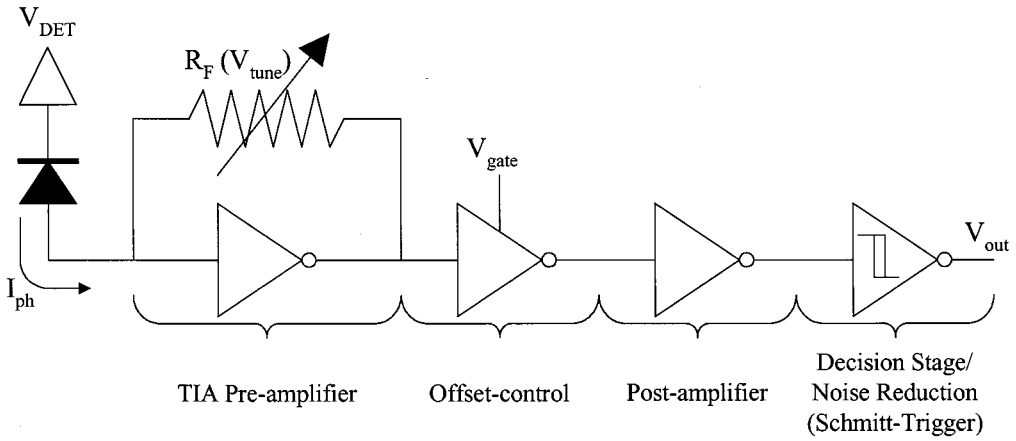

Fig. 10. Receiver circuit design.

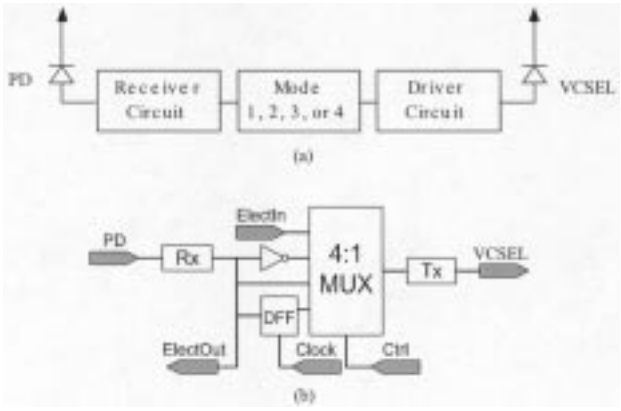

Fig. 11. (a) Simplified schematic of an individual pixel. (b) Logical description of an individual pixel indicating optical I/Os, the electrical I/Os, and the mode control multiplexer.

Thirty-two-bit-wide optical channels were defined by grouping 32 pixels together. Data fed into a control register was used to define the operating mode of each of the 32 pixels in the channel. All 32 bits in a given optical channel operated in the same mode. The chip had eight channels for a total of 256 bits.

Electrical data was brought onto the chip and extracted from the chip via a 32-bit-wide electrical input bus and a 32-bit-wide electrical output bus, respectively. Using this arrangement, the chip supported a one-to-one mapping of the optical channel width to the electrical bus width. As can be seen in Fig. 12, there was one electrical input bus, one electrical output bus, and eight optical channels. In the transmit mode (mode 1 above), electrical data was fed onto the chip and transmitted optically from one of the eight channels. In receive mode (mode 2 above), optical data was received on one of the eight channels, and fed off-chip via the electrical output bus. In the retransmit mode (mode 3 or mode 4 above), optical data was received and retransmitted optically without the use of the electrical buses. An additional feature of the chip was the ability to optically broadcast data. In this mode, the electrical input bus could be connected to multiple optical channels simultaneously.

Referring to Section II, the relationship between the eight optical channels and the 32 clusters was as follows. A 32-bit-wide optical channel was comprised of four pixels per cluster grouped together across eight clusters. This grouping was done along the long axis, as is shown in Fig. 5, where two of the eight optical channels are indicated.

The 1-cm $\times$ 1-cm CMOS chip was fabricated in 0.35-μm foundry silicon and was manufactured by the Taiwanese Semiconductor Manufacturing Corporation (TSMC). Access to the technology was made available through the Canadian Microelectronics Corporation [15]. The chip had 250 electrical I/O and power connections. These connections included pads for electrical I/O buses, analog and digital control signals, PD bias, and power and ground pads for the digital circuitry, the laser driver circuits, and the receiver circuits. Fig. 13 shows a photograph of the hybridized CMOS chip wirebonded into a 256-pin grid array (PGA) package. The packaged chip was subsequently inserted into a PCB, described in the next section, and successfully tested in each of the four modes described above. The results of the testing are presented in Section VII.

V. PCB Architecture and Performance

The PCBs were designed to support not only testing but also integration of the chip into the free-space microoptic interconnect. Using a field programmable gate array (FPGA) and a computer interface, the boards provided the necessary functionality to operate the CMOS chip in the modes described in Section IV. In addition, the chip temperature was control by a thermoelectric cooler (TEC) that was attached to the back of the PGA. Fig. 14(a) shows a picture of an assembled board. The packaged chip was inserted into the socket on the right-hand side of the board, and the TEC was interfaced to the back of the PGA via the large through hole in the PCB.

The architecture of the PCB and control system is shown in Fig. 14(b). The board was connected to a PC via a digital interface card capable of handling 96 I/Os. A graphical user interface (GUI), written in Visual C++, served as an interface for controlling the chip. The computer was used to generate the control signals and either the computer or the FPGA could generate data vectors depending on the settable PCB operating mode. The rate at which the computer could generate data through the interface card was in the kHz range. Thus, when MHz operation of the

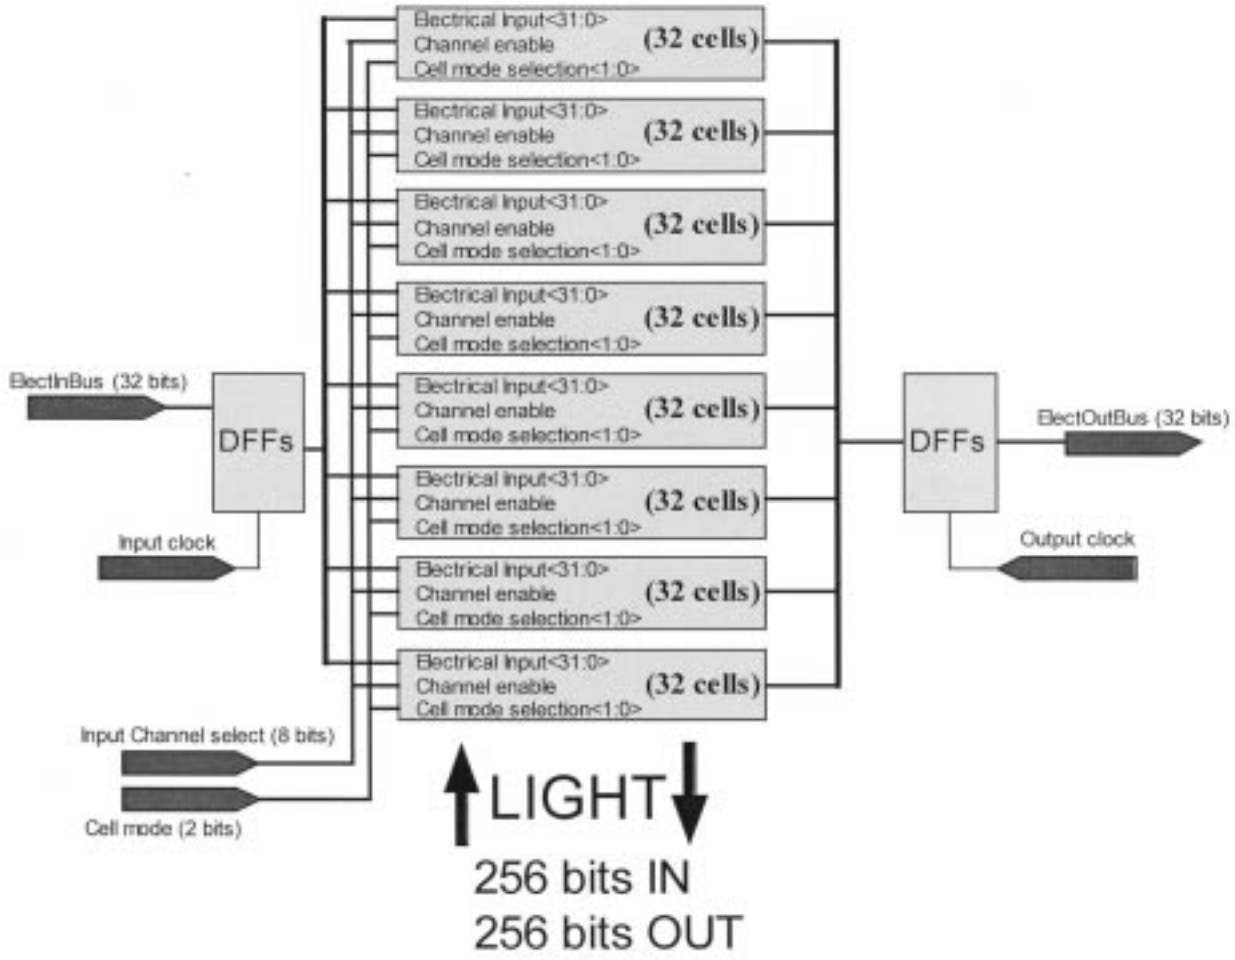

Fig. 12. Architecture of OE-VLSI ASIC. The "ElectInBus (32 bits)" and the "ElectOutBus (32 bits)" represent the 32-bit-wide electrical input and output buses, respectively. Each bus is registered by D-flip-flops (DFFs) and controlled by the "Input clock" and the "Output clock" signals, respectively.

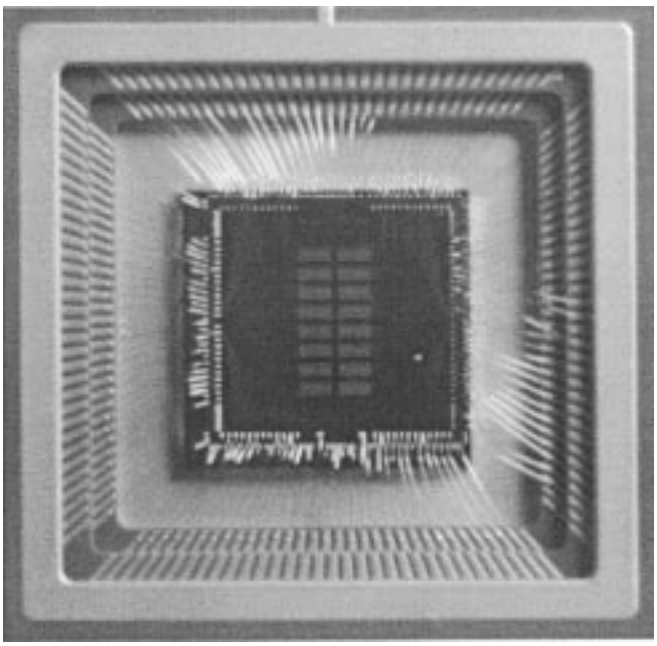

Fig. 13. 1 cm × 1 cm 0.35-μm OE-VLSI ASIC packaged in a 256-pin PGA.

chip was required, it was necessary to incorporate the FPGA between the interface card and the chip.

In order to support MHz data rates, the effects of transmission lines on signal integrity had to be minimized. This was achieved by placing the FPGAs as close as possible to the chip and by using short traces on the PCB. At the time that the board was designed and manufactured, the FPGA was one of the fastest available. For simple designs such as counters or linear feedback shift registers (LFSRs), the FPGA could be clocked at 166 MHz. Data vectors could, in theory, have been generated at 83 MHz which was half the frequency of the clock. In order to minimize clock skew, the chip received its clock from the FPGA thus reducing the true data generation rate by a factor of two. As a result, the chip could be exercised at 41.5 MHz, or 83 Mb/s. Fig. 15 shows a trace obtained at 100 Mb/s indicating that the PCB and FPGA outperformed specifications. The FPGA could generate patterns such as binary codes, 16-bit patterns, and pseudo-random bit sequence (PRBS) sequences. After describing a bulk optics-based interconnect in Section VI, experimental results on the performance of the chip obtained while being driven by this PCB will be described in Section VII.

VI. Optical System

The optical system used was a bulk lens optical system that was based upon a double Petzval design, as is shown in Fig. 16.

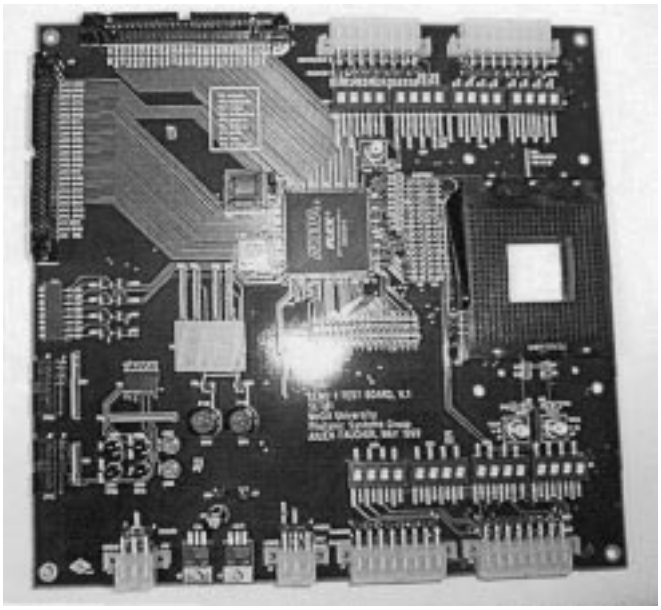

(a)

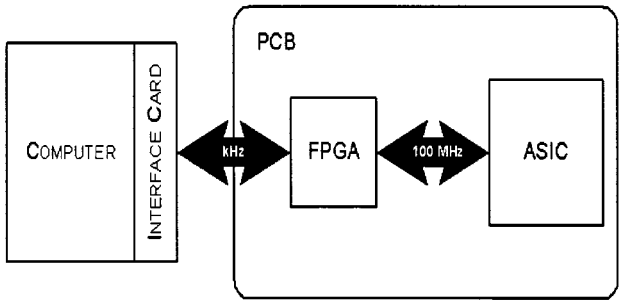

(b)

Fig. 14. (a) Photograph of assembled PCB. The PGA packaged chip was inserted into the ZIF socket (shown empty in this photograph) on the right-hand side of the board. (b) Schematic of the control system.

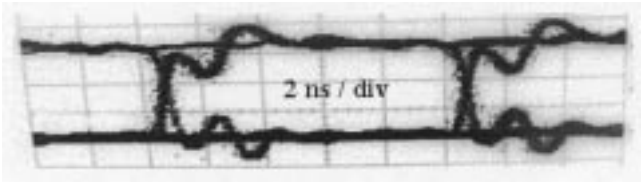

Fig. 15. 100-Mb/s electrical PRBS on the bus connecting the FPGA to the chip.

The lens types that were employed are given in the table shown in Fig. 16. Both lenses were achromatic doublets. This optical system was used primarily to validate the operation of the chip and did not represent the final optical configuration.

The system was designed to operate with a field of view that matched the chip dimensions. It had low distortion and was optimized to give acceptable aberrations across the field. It was also designed to allow a pellicle beam splitter to be inserted into the center of the optical system in order to allow simultaneous observation of both chips. The optical system was successfully attached to the PCBs described above, and board-to-board optical

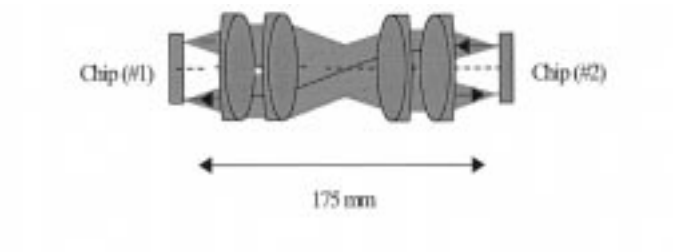

Lens #	Focal length (mm)	Diameter (mm)	Thickness (mm)
1	80.011	31.0	9.50
2	60.00	25.00	11.40

Fig. 16. Optical interconnect design and specifications on the achromatic doublets utilized in the interconnect.

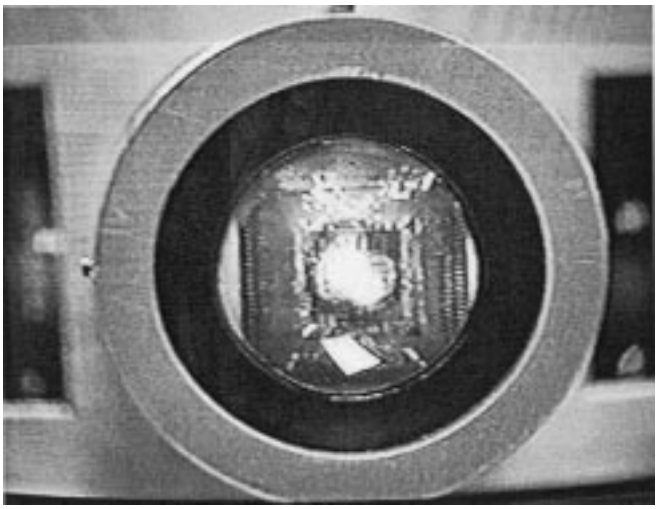

(a)

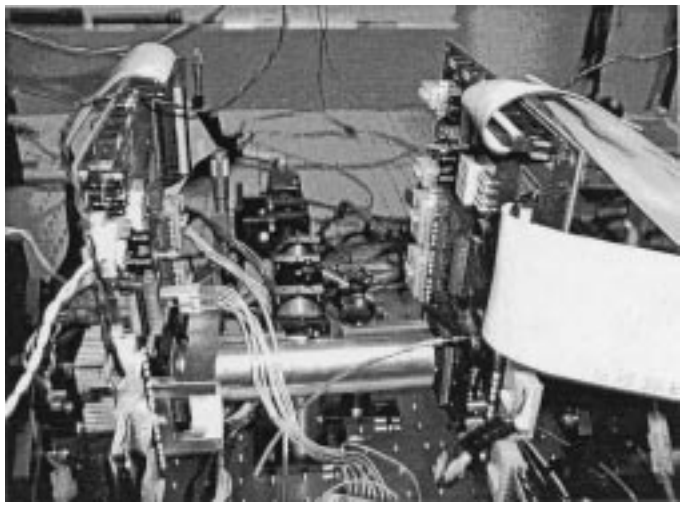

(b)

Fig. 17. (a) A view of the chip seen through the optical system. (b) A 256-channel bidirectional PCB-to-PCB optical interconnect system.

communication was achieved. Fig. 17(a) shows a chip imaged through the optical system, and Fig. 17(b) shows the assembled board-to-board system.

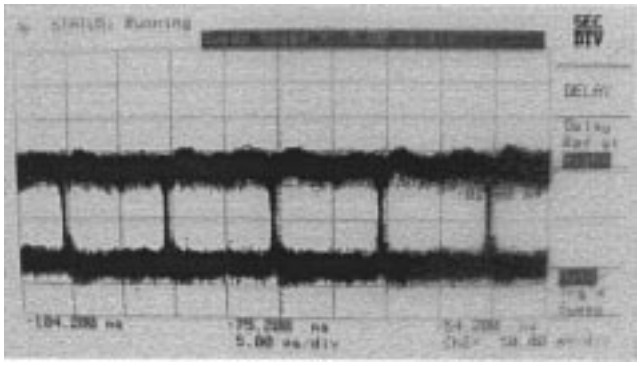

(a)

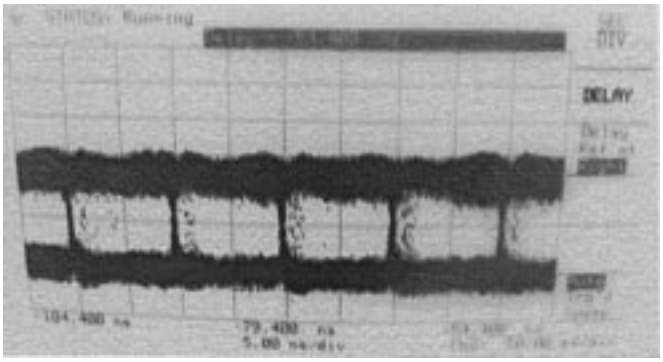

(b)

Fig. 18. (a) Mode 1 operation at 83 Mb/s. (b) Mode 1 operation with multiple adjacent VCSELs operating simultaneously at 83 Mb/s.

VII. Experimental Results

Experimental measurements were performed on both the chip and the board-to-board link. Experiments performed on the chip showed that it operated successfully in all four modes described above. We describe, in this section, experimental data obtained on individual chips and in the board-to-board link.

In operational mode 1, electrical-to-optical (E-to-O) data transfer was achieved. A pseudorandom data stream generated by the FPGA was injected onto the chip via the 32-bit electrical bus and then directed toward the appropriate channel. Optical output data was measured using an avalanche photodiode (APD), and the corresponding electrical signal was displayed on a digitizing oscilloscope. Fig. 18(a) shows the output of single transmitter (driver + VCSEL) being driven at 83 Mb/s; Fig. 18(b) shows the same transmitter with multiple adjacent transmitters being driven with pseudorandom data. Results were obtained using a bias current of 5 mA and a modulation current of 4 mA. Inspection of the data indicates there was some crosstalk injected onto the optical data. Coupling of the power-supply rails was believed to be partially responsible for this crosstalk.

The edge speed of the transmitter was measured using the APD and digitizing scope in an effort to extrapolate the bit rate that the laser driver circuits could potentially support in a stand-alone optical link. Fig. 19 shows an oscilloscope trace of the rising edge of the optical signal (the APD has a negative-polarity output, thus the laser on-state is the lower voltage portion of the trace). The horizontal scale is 200 ps/div, indicating a rising edge of less than 200 ps. This rise time suggests that the laser driver

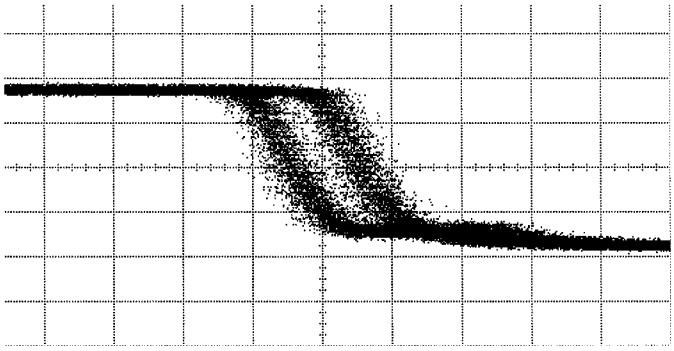

Fig. 19. High-speed VCSEL and driver measurement indicating a turn-on time of $\sim$200 ps.

circuit could potentially support data rates in excess of 1 Gb/s. These results are consistent with the simulation results obtained in Section III.

In operational mode 2 (O-to-E), an optical input was applied to a receiver circuit, and the receiver output was clocked to the output bus, where a logic analyzer could measure it. Fig. 20 is a logic analyzer trace showing a received signal at 150 Mb/s. Based on limitations imposed by the CMOS pad drivers and the PCB interface, this represented the maximum data rate obtainable in this operation mode.

We could avoid the fundamental limitations of performance found in operational modes 1 and 2 by using operational mode 3 (O-to-E-to-O). In this mode, an optical signal was input to a receiver, converted to an electrical signal, passed through combinational digital logic, and was then sent to a laser driver circuit which converted the electrical signal back to an optical signal. In doing this the combined performance of the receiver and transmitter circuits was measured. Fig. 21(a) and (b) shows measured eye diagrams at data rates of 250 and 400 Mb/s, respectively.

While obtaining measurements on a single receiver on a single chip was a relatively simple task, obtaining similar data on multiple receivers simultaneously proved to be considerably more difficult. All of the receivers in sets of four neighboring clusters (refer to Section II-A for a discussion of the clustered architecture) were commonly biased and controlled. This placed limitations on the operational freedom of these common-control receiver groups in terms of, for example, variations in incident optical power and PD responsivity. Due to these restrictions, we found the maximum number of simultaneously operable receivers in a common-control group to be limited and dependent on data rate. Of the 32 receivers in a common-control group, we were able to simultaneously operate two at data rates exceeding 50 Mb/s. This number increased to four receivers at data rates of approximately 1 Mb/s and to eight receivers at speeds less than 1 kHz. The limitations on the achievable operational parallelism for receivers in common-control groups is being further studied with the objective of devising design techniques to overcome them.

In addition to measurements on individual ASICs, we have also performed measurements on the board-to-board link shown in Fig. 17(b), successfully obtaining a bidirectional data link on four pairs of receivers at a data rate of approximately 1 Mb/s.

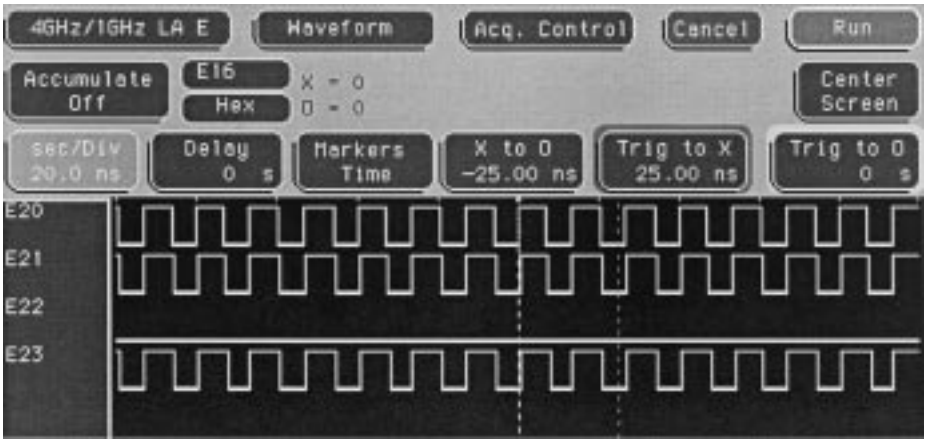

Fig. 20. Mode 2 operation at 150 Mb/s. Data taken with a logic analyzer reading signals off the PCB.

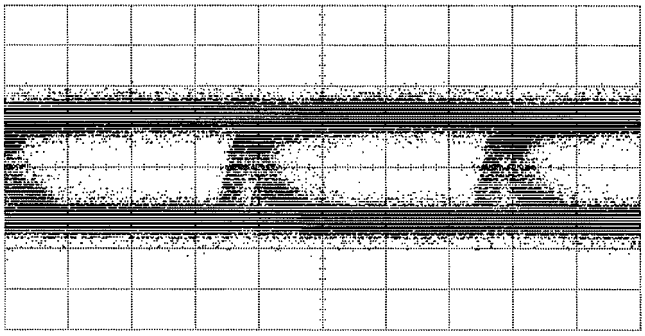

(a)

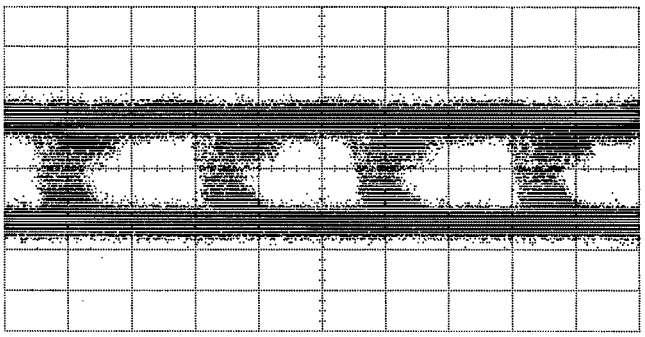

(b)

Fig. 21. Mode 3 operation at: (a) 250 Mb/s and (b) 400 Mb/s.

VIII. Conclusion

In conclusion, we have demonstrated a 256-channel, bidirectional optical interconnect between two PCBs containing VCSELs and photodiodes flip-chip bump-bonded to a 0.35-μm CMOS transceiver array. Experimental measurements on the chip indicate it was fully functional as per both the digital and analog design objectives. Unidirectional and bidirectional communication between boards was successfully demonstrated. This is the first system, to our knowledge, to send bidirectional data optically between optoelectronic VLSI chips that have both VCSELs and photodiodes cointegrated on the same substrate.

Acknowledgment

The authors gratefully acknowledge the contributions of T. Maj, A. Ghanem, F. Thomas-Dupuis, P. Sehgal, and T. Chiu of the McGill Photonic System Group; J. Trezza of Sanders/Lockheed Martin and J. T. Ciemiewicz of BAE Systems (formerly Sanders/Lockheed Martin Corporation) for heterogeneous integration and PD fabrication. The authors also gratefully acknowledge the Canadian Microelectronics Corporation (CMC) for providing the CMOS, and Emcore for providing the VCSELs used in this work.

References

[1] N. M. Jokerst and A. L. Lentine, Eds., *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 3–148, Apr. 1996.
[2] N. M. Jokerst, R. Baets, and K. Kobayashi, Eds., *IEEE J. Select. Topics Quantum Electron.*, vol. 5, pp. 143–398, Mar.-Apr. 1999.
[3] Y. Li, E. Towe, and M. Haney, Eds., *Proc. IEEE*, vol. 88, pp. 723–878, June 2000.
[4] D. L. Mathine, "The integration of III–V optoelectronics with silicon circuitry," *IEEE J. Select. Topics Quantum Electron.*, vol. 3, pp. 952–959, June 1997.
[5] A. V. Krishnamoorthy and K. W. Goossen, "Optoelectronic-VLSI: Photonics integrated with VLSI circuits," *IEEE J. Select. Topics Quantum Electron.*, vol. 4, pp. 899–912, Nov./Dec. 1998.
[6] M. Chateauneuf, M. B. Venditti, E. Laprise, J. Faucher, K. Razavi, F. Thomas-Dupuis, A. G. Kirk, D. V. Plant, T. Yamamoto, J. A. Trezza, and W. Lou, "Design, implementation, and characterization of a 2-D bi-directional free-space optical link," in *Proc. Optics Computing 2000 Conf.*, vol. 4089, 2000, pp. 530–538.
[7] Emcore, Somerset, NJ, private communication.
[8] N. H. E. Weste and K. Eshraghian, *Principles of CMOS VLSI Design, A Systems Perspective*, 2nd ed. Reading, MA: Addison-Wesley, 1993.
[9] T. K. Woodward, A. L. Lentine, A. V. Krishnamoorthy, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, B. T. Tseng, S. P. Hui, and R. E. Leibenguth, "Parallel operation of 50 element two-dimensional CMOS smart-pixel receiver array," *Electron. Lett.*, vol. 34, pp. 936–937, May 1998.
[10] D. A. Van Blerkom, C. Fan, M. Blume, and S. C. Esener, "Transimpedance receiver design optimization for smart pixel arrays," *J. Lightwave Technol.*, vol. 16, pp. 119–126, Jan. 1998.
[11] T. K. Woodward, A. V. Krishnamoorthy, K. W. Goossen, J. A. Walker, B. J. Tseng, J. E. Cunningham, R. E. Leibenguth, and W. Y. Jan, "1Gb/s single-beam smart-pixel receiver/transmitter in hybrid MQW-CMOS OE-VLSI technology," in *IEEE/LEOS Summer Topicals Meeting Smart Pixels*, Monterey, CA, July 1998, ThA2.
[12] A. V. Krishnamoorthy, P. Chandramani, T. K. Woodward, K. W. Goossen, B. J. Tseng, J. A. Walker, and W. Y. Jan, "Triggered receivers for optoelectronic-VLSI," in *Proc. Optics Computing 1999 Conf.*, PD1-1. Postdeadline paper .

[13] D. R. Rolston, D. V. Plant, T. H. Szymanski, H. S. Hinton, W. S. Hsiao, M. H. Ayliffe, D. Kabal, M. B. Venditti, A. V. Krishnamoorthy, K. W. Goossen, J. A. Walker, B. Tseng, S. P. Hui, J. C. Cunningham, and W. Y. Jan, "A hybrid-SEED smart pixel array for a four-stage intelligent optical backplane demonstrator," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 97–105, Apr. 1996.

[14] T. H. Szymanski and H. S. Hinton, "Reconfigurable intelligent optical backplane for parallel computing and communications," *Appl. Opt.*, vol. 35, pp. 1253–1268, 1996.

[15] Canadian Microelectronic Corp, Somerset, NJ. [Online]. Available: www.cmc.ca.

David V. Plant (S'86–M'89) received the Ph.D.degree in electrical engineering from Brown University, Providence, RI, in 1989.

From 1989 to 1993, he was a Research Engineer in the Department of Electrical and Computer Engineering at the University of California, Los Angeles (UCLA). He has been a Professor and member of the Photonic Systems Group within the Department of Electrical and Computer Engineering at McGill University, Montreal, QC, Canada, since 1993. Since joining McGill University, he has been studying optical interconnects for very short reach (VSR) applications. He is the Major Project Leader of the Photonic Interconnection Systems major project within the Canadian Institute for Telecommunications Research (CITR) and a Project Co-Leader within the Canadian Institute for Photonic Innovation (CIPI). During the 2000–2001 academic year, he took a leave of absence from McGill University to become the Director of Optical Integration at Accelight Networks, a startup company based in Pittsburgh, PA, and Ottawa, ON, Canada, which will deploy optical internetworking equipment.

Dr. Plant received the Outstanding Departmental Teaching Award and the Faculty of Engineering Teaching Award in 1996. In recognition of his accomplishments in research and teaching, he was named a James McGill Professor in 2001.

Michael B. Venditti (S'94–A'96) was born in Montreal, QC, Canada, on November 13, 1974. He received the B.Eng. (honors) and M.Eng. degrees from McGill University, Montreal, in 1996 and 1999, respectively. He is currently pursuing the Ph.D. degree at McGill University, focusing on transmitter, receiver, and system design for OE-VLSI applications.

His academic and research interests include optoelectronic devices, mixed analog/digital ASIC design, and OE-VLSI systems.

Emmanuelle Laprise received the B.Eng. degree in computer engineering from McGill University, Montreal, QC, Canada, in 1996. She is currently pursuing the Master's degree at McGill with the Photonics Systems Group.

She works at PMC-Sierra, Montreal, QC, Canada, as a product design engineer. Her research interests include the challenges that a designer faces when designing ASICs to be used in highly parallel optical interconnect systems.

Julien Faucher (S'99–M'99) received the B.Eng. degree in computer engineering from McGill University, Montreal, QC, Canada, in 1999. He is currently pursuing the M.Eng. degree in electrical engineering at McGill.

His research interests include parallel error-correction coding schemes and optoelectronic-VLSI circuits.

K. Razavi, photograph and biography not available at the time of publication.

Marc Châteauneuf received the B.Eng. degree in electrical engineering from McGill University, Montreal, QC, Canada, in 1998. He is currently pursuing the Ph.D. degree in electrical engineering at McGill University.

His research interests include free-space optical interconnects.

Andrew G. Kirk (M'97) received the B.Sc. degree in physics from the University of Bristol, Bristol, U.K., in 1989 and the Ph.D. degree in physics from the University of London (King's College), London, U.K., in 1992.

He has been a Professor and member of the Photonic Systems Group in the Department of Electrical and Computer Engineering at McGill University, Montreal, QC, Canada, since 1986. For the last three years, he has provided industrial training to a variety of major photonics companies. His current research interests are in the area of microoptics, optical MEMS, and free-space optical interconnects. He has published a number of refereed papers in these fields and has collaborated actively with a variety of companies.

Dr. Kirk was appointed as William Dawson Scholar at McGill University in 2001. He also recently received the annual Teaching Award of the Electircal and Computer Engineering Department.

J. S. Ahearn, photograph and biography not available at the time of publication.