

Optical Transceiver Chips Based on Co-Integration of Capacitively Coupled Proximity Interconnects and VCSELs

Xuezhe Zheng, Jon K. Lexau, Jonathan Bergey, John E. Cunningham, Ron Ho, Robert Drost, and Ashok V. Krishnamoorthy

Abstract—Combining the strengths of both proximity communication and optical communication, a new hybrid input–output (I/O) platform delivers on-chip bandwidth off-chip and over distance. We demonstrate, for the first time, a four-channel hybrid I/O interface by integrating proximity communication and vertical-cavity surface-emitting-laser-based parallel optical interconnects on the same commercial 90-nm complementary metal–oxide–semiconductor platform. The optical I/O can operate at 5 Gb/s per channel, and the complete hybrid I/O interface achieved 2.5 Gb/s per channel. We characterize the I/O link performance for various data rates and chip separations, and show 10- μm chip separation tolerance for proximity communication.

Index Terms—Communication, interconnects, input–output (I/O), optical, proximity, scalable, vertical-cavity surface-emitting laser (VCSEL).

I. INTRODUCTION

COPPER-BASED electrical links have traditionally dominated the ultrashort reach interconnects. However, optical interconnects offer a promising alternative link technology not limited by wire aspect ratios [1]. However, while much progress has been made in the direct integration of optoelectronic (OE) devices with complementary metal–oxide–semiconductor (CMOS) very large scale integration (VLSI) chips [2], there remain challenges in packaging, density, power, cost, and reliability for practical applications [3]. Our goal is to provide this integration for computing chips, whose power budgets are very tight, and space allocated for interconnects is limited.

To bridge these two solutions, we have developed a short-range low-cost electronic input–output (I/O) technology with unprecedented bandwidth density. Proximity communication is based on the observation that fast and low-cost communication is possible over very short distances [4]. It relies on capacitive coupling between chips that are placed face-to-face, such as chip1 and chip2 shown in Fig. 1. The two chips communicate through the capacitive coupling of aligned metal plates




Fig. 1 Integration of VCSEL-based parallel optical interconnects with proximity communication.

on each chip that are connected to the transmitter (TX) and receiver (RX) circuits, respectively. The two chips can be pushed together such that their surfaces touch or nearly touch, allowing the capacitively coupled plates to be in very close proximity, and thus enabling small TX and RX structures that have small parasitic capacitances and consume little power. The small plates also allow signal densities two orders of magnitude greater than traditional off-chip communication that uses wire-bonding or ball-bonding. In addition, the TX circuits drive only a high-impedance capacitive pad, very much akin to the gate of a transistor. This removes the need for the high-to-low impedance conversion that has traditionally prevented the power dissipation of off-chip drivers from scaling down with feature sizes. High-speed (up to 1.8 Gb/s) proximity communications with low power (less than 3.6 mW per channel), and very high bandwidth density (430 Gb/s/mm²) have been demonstrated with 36- μm pads in a 180-nm CMOS technology [5].

Historically, the optical I/O capability from a transceiver chip has typically been limited by the electrical I/O [6]. The use of proximity communication helps restore the electrical I/O capability and balance the bandwidth to the desired ratios. It enables low power scalable I/O for VLSI chips over short distances. It also provides a way to inexpensively attach OE devices to CMOS chips, by enabling full-bandwidth communication between a dedicated OE chip and a high-performance CMOS chip. To create this I/O bridge, we integrated proximity and optical communications together on the same CMOS platform, and call it the light-out-proximity-in (LOPI) chip. Chip1 in Fig. 1 is such a chip, and it integrates proximity I/O circuits, vertical-cavity surface-emitting laser (VCSEL) drivers, and RX circuits on the same silicon. VCSEL and photodetector arrays

Manuscript received November 28, 2006; revised January 12, 2007.

X. Zheng, J. E. Cunningham, and A. V. Krishnamoorthy are with Sun Microsystems Inc., San Diego, CA 92121 USA (e-mail: xuezhe.zheng@sun.com).

J. K. Lexau, R. Ho, and R. Drost are with Sun Microsystems Inc., Menlo Park, CA 94025 USA.

J. Bergey is with Forza Silicon Corporation, Pasadena, CA 91101 USA.

Color versions of one or more of the figures in this letter are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LPT.2007.892908




Fig. 2. Silicon chip integrates both proximity and optical I/O circuits. (a) Chip layout; (b) chip functional diagram.

are bonded on the chip via flip-chip bonding, or wire bonding as shown in the figure. Fiber arrays with matching pitch are applied to butt-couple with the VCSELs and photodetectors for optical I/O. The chips are bridged through aligned proximity TXs and RXs as shown in the inset diagram.

II. TEST CHIP

In order to demonstrate this integration of proximity communication and optical I/O, we implemented a LOPI test chip in a 90-nm commercial CMOS process. Fig. 2(a) shows the test chip's floor plan. In addition to the proximity communication interface and optical I/O interface, the chip also has an electrical I/O interface for testing. Each interface supports four channels. The electrical I/O uses differential current mode logic (CML), allowing us to simply ac-couple the CML input buffer to a pattern generator, and eliminate the difficulty of insuring a proper input common-mode. The CML output buffer circuits can drive 5-Gb/s data off-chip and employs three digital waveform-shaping signals to control the edge rate, vary the output voltage swing, and turn off the output buffer.

The proximity I/O interface contains TX plates and RX plates both for data path and for chip alignment. With our emphasis on integration demonstration, we used relatively large proximity pads instead of pushing limits for high density and throughput. In the data path, each TX plate is subdivided into a 4×4 array of 16 independent microplates, each $26.5\text{-}\mu\text{m}$ square on a $31.25\text{-}\mu\text{m}$ pitch. While normally we send the same transmit data to all 16 microplates centered on the array, in case of chip

misalignment, we can steer the transmit data up to two micropads distance in any of the four directions. Each RX uses a single plate, which at $105.5\text{-}\mu\text{m}$ square and spaced $125\text{ }\mu\text{m}$ to its neighbors, is slightly smaller than one 4×4 group of microplates. Multiple sets of TX/RX plates, grouped into arrays that we call "where" blocks, detect the relative positions of two chips electrically. Each "where" block is a 20×20 array of small transmit and receive plates on $25\text{-}\mu\text{m}$ centers without any microsteering capability. Sweeping different patterns across the transmit "where" array and using different sets of RXs allows us to infer chip position and potentially correct any misalignment.

The optical circuits for the LOPI chip consist of a four-channel 5-Gb/s per channel VCSEL driver TX and a four-channel 5-Gb/s per channel optical RX. The optical TX, which converts a CML electrical signal into a current that can drive a VCSEL load at high speeds, can be configured by the user for optimization of various VCSEL and package configurations. These tunable parameters include VCSEL bias, modulation current, and waveform-shaping signals to control the edge rate and crossing of the current driving the VCSEL. The optical RX, which amplifies a small photodiode current into a differential CML output, consists of a transimpedance amplifier and a five-stage limiting amplifier with offset correction. Like the optical TX, the optical RX also has several user-configurable features to optimize the link performance for different components and packages.

The three different I/O interfaces interconnect via a network of multiplexer (MUX) and demultiplexer (DeMUX) circuits, allowing all I/O configurations to be implemented. Fig. 2(b) shows the test chip functional block diagram. As shown by the arrowed lines in the figure, data paths can be configured to loop back at the same I/O interface or between any two I/O interfaces by the on-chip MUX/DeMUXs.

III. EXPERIMENTAL SETUP AND TEST RESULTS

We built a printed circuit board (PCB) to support and test the LOPI chips. A set of three PCBs, each containing a single LOPI chip, allows for comprehensive testing of all I/Os. All three PCBs support proximity communication. In addition, PCB1 is configured with CML electrical in-out; PCB2 is configured with CML electrical in and optical out by attaching a VCSEL array next to the LOPI chip; and PCB3 is configured with optical in and CML electrical out by attaching a photodetector array next to the LOPI chip. The OE arrays are bonded to the LOPI chips via wire-bonding. Two LOPI chips are arranged face-to-face by mounting PCB1 face-down on a six-axis manipulator, and mounting PCB2 facing up on a fixture. With feedback from the on-chip "where" blocks, we use the six-axis manipulator to correct the relative board positions until the two LOPI chips are aligned. For practical applications, the chips will be positioned using either traditional or improved methods, and then held into place using a retention mechanism (such as a spring). The "where" blocks will only be used to analyze the resulting chip position, and guide the data steering to correct the residual misalignment.

A fiber with lens collimator couples the optical output from the VCSEL on PCB2 with 1-dB optical loss. The other end of the fiber is butt-coupled to the photodetector on PCB3 for a complete optical link. The VCSEL and photodetector array implemented are both 10-G 1 \times 4 array devices with a pitch of $250\text{ }\mu\text{m}$, from Avalon and Emcore, respectively. The VCSEL

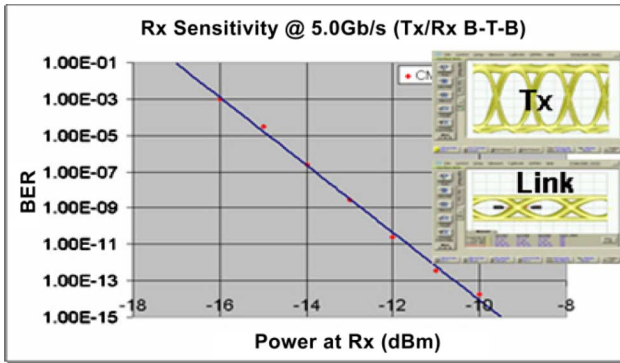


Fig. 3. Optical RX sensitivity at a data rate of 5 Gb/s with optical TX and RX connected back-to-back. Inset shows the eye diagram of the optical TX and link at a 5-Gb/s data rate.




Fig. 4. Timing "bathtub" of LOPI I/O. Inset shows the output signal eye diagram at 1.85 and 2.5 Gb/s.

TXs achieved an average optical power of -3 dBm, and an extinction ratio better than 6 dB. The optical link demonstrated operation up to 5 Gb/s, with the bit-error rate (BER) versus the RX optical power plotted in Fig. 3. We observed no apparent noise floor to a BER of 10^{-14} , and we deduced an RX sensitivity of -11.5 dBm at a BER of 10^{-12} . The inset of Fig. 3 shows the eye diagram of the optical TX (top) and link (bottom) at 5 Gb/s.

We then tested the complete link, including proximity communication. Data into the electrical input of the LOPI chip on PCB1 was steered to the proximity interface, transferred to LOPI chip on PCB2 via proximity communication, steered to the optical interface, converted to optical output at the VCSEL, coupled into the fiber, and transmitted to the RX on PCB3. This total path, with data flowing across three chips without clocking, achieved a data rate of 2.5 Gb/s with BER better than 10^{-12} . Fig. 4 shows the timing "bathtub" plots for data rates of 1.85, 2.0, and 2.5 Gb/s, respectively. The inset shows the output signal eye diagram at 1.85 (top) and 2.5 Gb/s (bottom). As described above, in addition to the proximity and optical I/O interface, the data stream also goes through many MUX/DeMUXs for path configuration. Without retiming, the data stream accumulates too much jitter along the way, limiting the link speed to only half of what the optical link is capable of.

The proximity capacitance falls with chip separation, which in turn affects the signal magnitude coupled through, and ultimately the signal-to-noise ratio. We performed measurements to characterize the tolerance of proximity communication to chip separation while running the complete LOPI link at 1.85 Gb/s. We varied the chip separation and observed the link performance degradation. Fig. 5 shows the timing margin for different




Fig. 5. Timing margin of LOPI I/O for different chip separations.

proximity chip separations. The timing margin is measured at a BER of 10^{-12} . Typical clock data recovery circuitry requires a minimum phase margin of 40%UI to work properly. With this criterion, the plot indicates that the LOPI link can tolerate up to 10- μ m chip separation.

For all the results shown above, we used 5 m of 62.5/125 multimode fiber. We also tested the LOPI link with 100 m of fiber, but observed no noticeable performance degradation.

IV. CONCLUSION

Combining advanced optoelectronic technologies together with high-performance proximity communication can provide dense and high-speed I/O. We successfully integrated proximity communication with optical communication on a commercial 90-nm CMOS platform, demonstrating a promising I/O solution that can scale with VLSI technology. Four I/O channels, each at 2.5 Gb/s (or 10-Gb/s throughput) is achieved for the complete LOPI interface, with optical I/O capable of 5 Gb/s per channel (or 20-Gb/s throughput). This successful demonstration suggests that much higher density and throughput can be obtained seamlessly by using smaller proximity pads and integrating more optical channels for future applications.

REFERENCES

- [1] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proc. IEEE*, vol. 88, no. 6, pp. 728–749, Jun. 2000.
- [2] A. V. Krishnamoorthy, "Photonics-to-electronics integration for optical interconnects in the early 21st century," *Optoelectron. Lett.*, vol. 2, no. 3, pp. 163–168, May 2006.
- [3] D. V. Plant and A. G. Kirk, "Optical interconnects at the chip and board level: Challenges and solutions," *Proc. IEEE*, vol. 88, no. 6, pp. 806–818, Jun. 2000.
- [4] R. J. Drost, R. D. Hopkins, R. Ho, and I. E. Sutherland, "Proximity communication," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1529–1535, Sep. 2004.
- [5] D. Hopkins, A. Chow, R. Bosnyak, B. Coates, J. Ebergen, S. Fairbanks, J. Gainsley, R. Ho, J. Lexau, F. Liu, T. Ono, J. Schauer, I. Sutherland, and R. Drost, "Circuit techniques to enable 430 Gb/s/mm/mm proximity communication," in *IEEE Int. Solid-State Circuits Conf.*, Feb. 2007, pp. 368–369.
- [6] C. Cook, J. E. Cunningham, A. Hargrove, G. G. Ger, K. W. Goossen, W. Y. Jan, H. H. Kim, R. Krause, M. Manges, M. Morrissey, M. Perinpanayagam, A. Persaud, G. J. Shevchuk, V. Sinyansky, and A. V. Krishnamoorthy, "A 36-channel parallel optical interconnect module based on optoelectronics-on-VLSI technologies," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 387–399, Mar./Apr. 2003.