20-Gb/s Power-Efficient CMOS-Driven Multimode Links

Caroline P. Lai*, Clint L. Schow, Alexander V. Rylyakov, Benjamin G. Lee, Fuad E. Doany, Richard A. John, and Jeffrey A. Kash

IBM Research, T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, New York 10598
**Current Address: Department of Electrical Engineering, Columbia University, 500 West 120th St., New York, New York 10027*
caroline@ee.columbia.edu

Abstract: We report several firsts for CMOS-driven, VCSEL-based, multimode transmitters and receivers: serial links up to 20 Gb/s; unprecedented full-link power efficiencies from 8 pJ/bit (15 Gb/s) to 17 pJ/bit (20 Gb/s); and >25-Gb/s transmitter operation.

OCIS codes: (060.2360) Fiber optics links and subsystems; (060.2380) Fiber optics sources and detectors; (200.4650) Optical interconnects.

1. Introduction

Contemporary high-performance computing systems are poised to leverage optical technologies for achieving high interconnect bandwidth requirements with sufficiently low power consumption. IBM's Roadrunner supercomputer in 2008 utilized approximately 50 thousand optical links for 5-Gb/s fiber connections at the rack level [1], while IBM Blue Waters, slated for 2011, is estimated to have more than two million 10-Gb/s fiber-to-the-module optical interconnections [2]. It is clear parallel module-to-module optical links comprise a promising solution to optimize numerous interconnect metrics, such as bandwidth, power, and area, while simultaneously increasing link distance.

The Terabus program [3] endeavors to develop high-throughput, low-power optical interconnects between chip-scale multimode optical transceivers communicating through polymer waveguides integrated on printed circuit boards. This approach was successfully demonstrated using transceivers with 130-nm complementary metal-oxide-semiconductor (CMOS) chips with flip-chip bonded 985-nm optical devices [3]. In later generations, the Terabus transceivers have been re-designed to incorporate conventional 850-nm vertical cavity surface emitting lasers (VCSELs) and photodiodes (PDs) packaged with 130-nm integrated circuits (ICs) on a common silicon carrier (Fig. 1) [4]. Migrating the CMOS circuits forward to the 90-nm generation affords improvements in power consumption, bandwidth, and area performance, as demonstrated with initial transmitter (TX) and receiver (RX) prototypes [5]. We now report results on second-generation 90-nm TX and RX components including full-link characterization over a range of data rates, power consumptions, and link distances. The single-channel test sites (Fig. 1) incorporate core analog circuitry and optoelectronic components designed for future multichannel parallel transceivers.

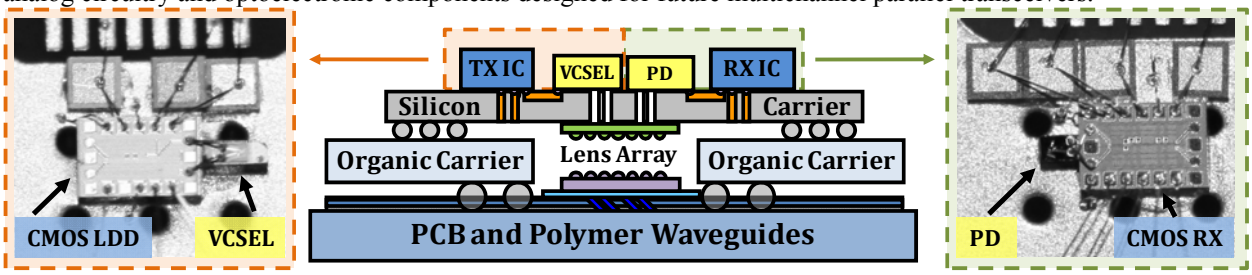

Fig. 1. Block diagram of envisioned Terabus assembly, with pictures of 90-nm CMOS single-channel transmitter (left) and receiver (right) test sites with wire-bonded VCSELs and PDs.

2. Transmitter and Receiver Performance

The TX and RX ICs were fabricated in a standard digital bulk 90-nm IBM CMOS process. Each chip is wire-bonded to a circuit board that provides power and bias connections. Short wire bonds connect the ICs to the optoelectronic devices. Lensed fiber probes couple light out of/into the VCSELs and PDs, and coaxial microwave probes provide access to high-speed electrical I/O. All testing was performed with 2^7-1 pseudo-random binary sequence (PRBS).

The TX uses a custom-designed laser diode driver (LDD), occupying 125 × 500 μm^2, and is based on a two-stage Cherry-Hooper pre-amplifier followed by an inductively peaked single-ended output stage. The input of the LDD is 50-Ω single-ended terminated. The VCSEL, fabricated by Emcore Corporation, has a ~5-μm aperture diameter, sub-mA threshold current, and 3-dB electrical bandwidth of 14 GHz (when biased at ~6 mA) [6].

The transmitter was operated over a wide range of data rates and voltages. It was tested using a reference receiver constructed with the same PD in the receiver assemblies described below, bonded to a commercial 40-Gb/s transimpedance amplifier (TIA). Total TX power consumption, including the laser and driver, is as low as 14 mW

while maintaining error-free performance (bit-error rate (BER) < 10^{-12}) at 20 Gb/s. To the authors' knowledge, the resulting energy efficiency of 0.7 pJ/bit is the best reported to date for a complete optical transmitter. Eye diagrams at 20 Gb/s for three different power settings are provided (Fig. 2). Operating the TX with higher supply voltages (56 mW power consumption) allows for operation up to 28 Gb/s with BERs < 10^{-12}. The temporal eye openings (at a BER of 10^{-8}) are 0.45 UI at 20 Gb/s and 0.4 UI at 25 Gb/s. The transmitter measured energy-per-bit values are plotted with respect to data rate (Fig. 2); at a given bit rate, the power efficiency was obtained by lowering the TX power consumption to a minimum value while maintaining a BER less than 10^{-12}. The best power efficiency is achieved at a data rate of 17.5 Gb/s. Above this speed, the higher TX power consumption (mW) required to maintain performance outpaces the increased data rates (Gb/s), degrading the efficiency in mW/Gb/s, or equivalently, pJ/bit. The power efficiency curve rises at low data rates since the TX must produce enough modulated power to maintain a minimum receiver output swing; once the TX operates at its minimum power for a given receiver output, the efficiency degrades at lower rates as the TX power is amortized over fewer transmitted bits.

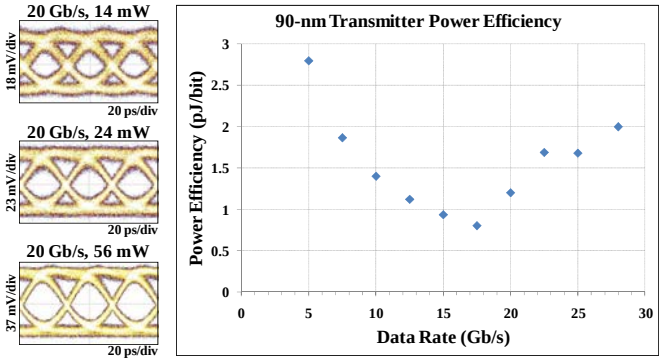

Fig. 2. Transmitter performance with reference photodiode: eye diagrams at 20 Gb/s for three power settings, and power efficiency curve.

The RX is comprised of a TIA and a six-stage limiting amplifier (LA) with a 125 × 600 μm^2 footprint. A 25-μm aperture diameter Emcore PD is wire-bonded to the CMOS chip. The PD has a measured responsivity of 0.55 A/W, dark current of 0.1 nA, a capacitance of 80 fF, and a 3-dB electrical bandwidth greater than 20 GHz [6]. The receiver was tested under various operating conditions using a reference VCSEL as the optical source. The RX operates at 15 Gb/s with a total power consumption of 43 mW, resulting in 2.9 pJ/bit. At the maximum speed of 20 Gb/s, the RX consumes 135 mW (Fig. 3). The timing margin (at a BER of 10^{-8}) is 0.45 UI at 10 Gb/s and 0.32 UI at 17.5 Gb/s. The receiver measured energy-per-bit values are plotted with respect to data rate (Fig. 3). Similar to the TX, at a given data rate, the best power efficiency was obtained by minimizing the RX power consumption while ensuring that the BERs remained less than 10^{-12} and the differential output voltage swing was greater than 200 mV. The RX power efficiency curve indicates an optimal energy-per-bit performance at a data rate of 15 Gb/s. At higher and lower data rates, the efficiency is degraded for the same reasons presented above considering the TX.

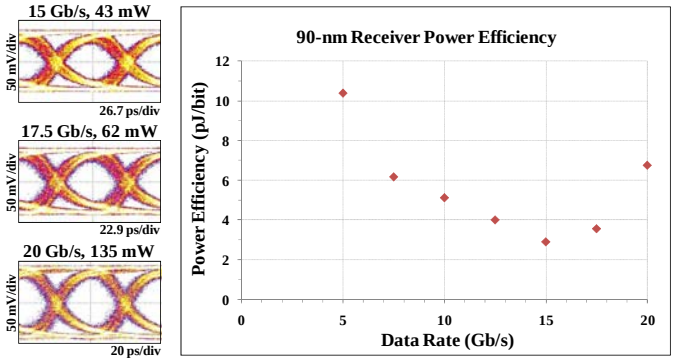

Fig. 3. Receiver performance with reference VCSEL: eye diagrams at several data rates and bias conditions, and power efficiency curve.

3. Link Experimental Results

We further explore the performance of a full link created by connecting the single-channel transmitter and receiver components described above through standard multimode fiber. At a total power consumption of 114 mW, BERs less than 10^{-12} are obtained at 15 Gb/s, exhibiting a power efficiency of 7.6 pJ/bit. Error-free performance is achieved at 20 Gb/s at an increased power consumption of 340 mW, yielding an energy-per-bit of 17 pJ/bit. Fig. 4 shows the link eye diagrams at several data rates for various total link power dissipation. The power efficiency for

the full link with a very short ("0 m") fiber is shown versus data rate in Fig. 4; the link BER is less than 10^{-12} for all these measurements. It is clear that the achieved energy-per-bit performance for the link is optimal at 15 Gb/s. Fig. 4 also shows a family of RX sensitivity curves at several data rates.

BER measurements were also conducted to evaluate the link transmission penalty performance at several distances. Error-free performance is maintained at 18 Gb/s after 50 m of OM3 multimode fiber and at 17.5 Gb/s after 100 m. BER curves are recorded (Fig. 5) while varying the link length, for data rates of 15 Gb/s and 17.5 Gb/s. Increasing the link length to 100 m adds a minimal transmission penalty of ~1 dB at both of these data rates.

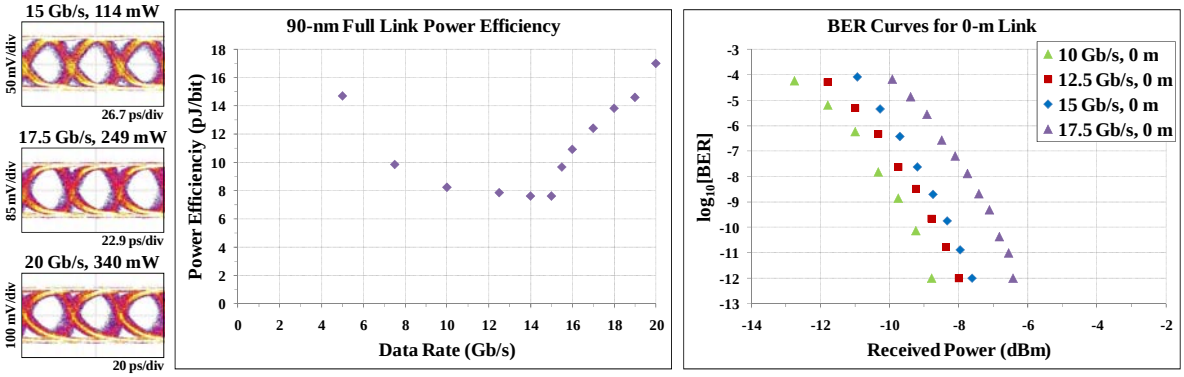

Fig. 4. Full link: eye diagrams at several data rates and bias conditions, power efficiency curve, and BER curves for the link at 0 m.

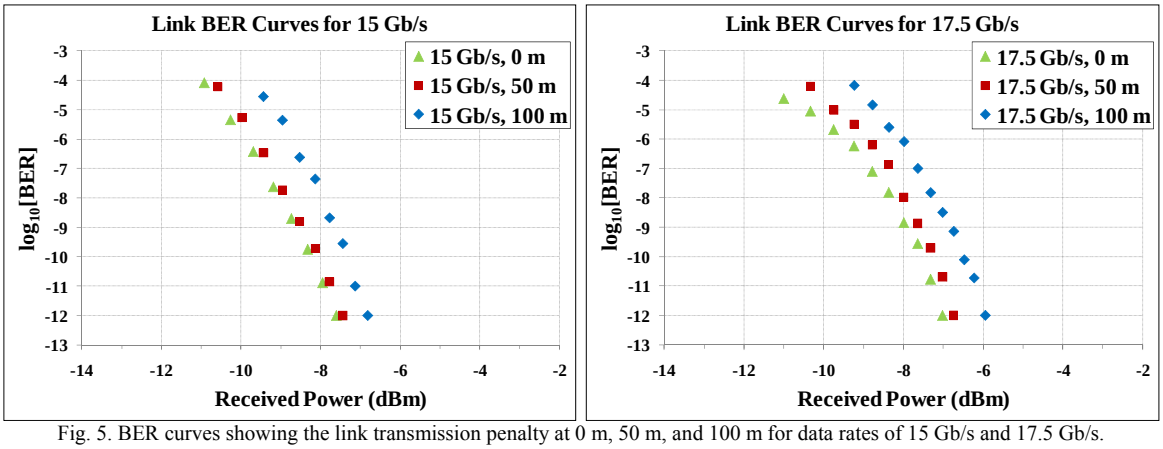

Fig. 5. BER curves showing the link transmission penalty at 0 m, 50 m, and 100 m for data rates of 15 Gb/s and 17.5 Gb/s.

4. Conclusions

Interconnects in future high-performance computing systems will require a tremendous number of high-bandwidth, low-power, densely-integrated optical interconnects. Parallel, VCSEL-based links dominate the current market for short-reach (< 100 m), high-bandwidth interconnects, and will remain a viable technology choice for future systems. The work reported here highlights the promise of CMOS-driven, VCSEL-based links, demonstrating for the first time serial links at 20 Gb/s, full-link power efficiencies less than 8 pJ/bit at 15 Gb/s, and transmitters producing wide-open eyes at 25 Gb/s.

The authors gratefully acknowledge support from DARPA under contract HR0011-08-C-0102. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of DARPA or the Department of Defense. Approved for public release; distribution unlimited.

5. References

[1] D. Grice *et al.*, "Breaking the petaflops barrier," *IBM J. Res. & Dev.*, vol. 53, no. 5, paper 1, 2009.

[2] National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana-Champaign, Urbana, IL. http://www.ncsa.illinois.edu/BlueWaters/, [Online]. Accessed 6/09/2010.

[3] F. E. Doany *et al.*, "160 Gb/s Bidirectional Polymer-Waveguide Board-Level Optical Interconnects Using CMOS-Based Transceivers," *IEEE Transactions on Advanced Packaging*, vol. 32, no. 2, pp. 345–359, May 2009.

[4] F. E. Doany *et al.*, "Terabit/s-class 24-channel bidirectional optical transceiver module based on TSV Si carrier for board-level interconnects," in *Proc. 60th Electronic Components and Technology Conference (ECTC)*, pp. 58-65, Jun 2010.

[5] B. G. Lee *et al.*, "Lower-Power CMOS-Driven Transmitters and Receivers," in *Proc. Conference Lasers and Electro-Optics (CLEO)*, paper CMB5, May 2010.

[6] N. Y. Li *et al.*, "High-performance 850 nm VCSEL and photodetector arrays for 25 Gb/s parallel optical interconnects," in *Proc. Optical Fiber Communications (OFC)*, paper OTuP2, Mar 2010.