End-to-End Multicore Multimode Fiber Optic Link Operating up to 120 Gb/s

Benjamin G. Lee, *Member, IEEE, Member, OSA*, Daniel M. Kuchta, *Senior Member, IEEE*, Fuad E. Doany, Clint L. Schow, *Senior Member, IEEE*, Petar Pepeljugoski, Christian Baks, Thierry F. Taunay, *Member, OSA*, Benyuan Zhu, Man F. Yan, George E. Oulundsen, Durgesh S. Vaidya, Wenlin Luo, and Neinyi Li

***Abstract*—A full multicore fiber optic link is demonstrated, transmitting greater than 100 Gb/s through a single strand of multimode fiber for the first time. The fiber, which consists of seven graded-index multimode cores, is used to transmit up to 120 Gb/s over 100 m using a custom multicore-fiber interfacing transmitter and receiver. 2-D arrays of vertical-cavity surface-emitting lasers (VCSELs) and vertically illuminated photodiodes (PDs) are fabricated with a geometry corresponding to the outer six cores of the seven-core fiber, which is arranged in a hexagonal pattern. Both flip-chip and wire-bonding technologies are used to package the VCSEL and PD chips with multichannel transmitter and receiver integrated circuits. Amplitude and timing margins of the end-to-end signals are analyzed through bit-error-rate (BER) measurements. The effects of electrical and optical crosstalk are shown to result in negligible degradation to the BER performance.**

***Index Terms*—Integrated optoelectronics, optical fiber communication, optical interconnections, optical receivers, optical transmitters.**

I. Introduction

OPTICAL interconnects have become prevalent within the top-ranked high-performance computing (HPC) systems due in large part to the superior bandwidth-distance product that they provide compared to electrical interconnects. Initially, fiber displaced only the lengthiest electrical links, as in ASC Purple [1], resulting in a relatively small volume of fiber within a system. Since then, the threshold distance at which the benefits of fiber exceed the cost has continually decreased, resulting in an exponentially increasing number of optical links, and generating a growing volume of fiber cabling. For example, the IBM Roadrunner supercomputer [2] transmits all interrack communication through 5000 active optical cables totaling 55 mi in length. Moreover, IBM's Power 775 system [3] employs fibers for all rack-to-rack links and all drawer-to-drawer links as well. Here, arrays of passive optical connectors at the back of each drawer provide connectivity for any incoming or outgoing signal, replacing the electrical backplane by as many as 5376 fibers per drawer.

This growth in fiber volume has led to new and challenging design constraints for future systems. The communication infrastructure in the top-ranked supercomputers, which is now predominately fiber based, must meet ever-growing bandwidth requirements, while maintaining feasible system-wide power and cost targets. Notably, this must be achieved in a manner that does not result in an unmanageable and expensive volume of fiber cabling. Consequently, increasing bandwidth per fiber, while minimizing further increases in link cost and power, is an important focus for the design of future optical interconnects. Not only will this be vital to next-generation supercomputing, but to future datacenters as well, where the reduced cost (in units of $/Gb/s) may be very attractive.

Fiber bandwidth density can be improved using a variety of approaches. Single-mode fiber employing dense wavelength-division multiplexing (DWDM) has become widely implemented in long-haul telecommunications systems, but suffers from tighter alignment tolerances compared to multimode fiber (MMF) links, leading to increased cost. Coarse wavelength-division multiplexing (CWDM) is an analogous MMF-compatible approach, but has, thus far, been unable to compete with the cost of single-wavelength MMF solutions, despite impressive research results [4], [5].

In contrast to the wavelength-domain approach, the spatial domain may also be employed to increase fiber density, and, thus, bandwidth density. The most straight-forward approach is to reduce the fiber-to-fiber pitch in parallel ribbon arrays. Ultimately, this requires reducing the fiber cladding diameter so that the fiber cores may be spaced more closely. Bandwidth density improves as the square of the reduction in pitch, when pitch is the same in the two lateral directions. However, if the outer diameter is greatly reduced, the fiber's mechanical stability may be compromised, adding cost during assembly. Alternatively, mode-division multiplexing using multiple-input multiple-output configurations has been investigated [6], but inherently suffers from instabilities due to unpredictable and nonstatic mode-mixing behavior.

Multicore fiber (MCF), studied now for many decades [7], presents an alternate spatial-domain path toward improving bandwidth density while maintaining standard cladding

Manuscript received September 27, 2011; revised December 20, 2011; accepted December 21, 2011. Date of publication January 11, 2012; date of current version February 15, 2012.

B. G. Lee, D. M. Kuchta, F. E. Doany, C. L. Schow, P. Pepeljugoski, and C. Baks are with IBM Research, Yorktown Heights, NY 10598 USA (e-mail: bglee@us.ibm.com; kuchta@us.ibm.com; doany@us.ibm.com; cschow@us.ibm.com; petarp@us.ibm.com; cbaks@us.ibm.com).

T. F. Taunay and B. Zhu are with OFS Laboratories, Somerset, NJ 08873 USA (e-mail: ttaunay@ofsoptics.com; bzhu@ofsoptics.com).

M. F. Yan is with OFS Laboratories, Murray Hill, NJ 07974 USA (e-mail: mfy@ofsoptics.com).

G. Oulundsen and D. S. Vaidya are with OFS, Sturbridge, MA 01566 USA (e-mail: goulundsen@ofsoptics.com; dvaidya@ofsoptics.com).

W. Luo and N. Li are with Emcore Digital Product Division, Albuquerque, NM 87123 USA (e-mail: Wenlin_Luo@emcore.com; Nelson_Li@Emcore.com).

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

Digital Object Identifier 10.1109/JLT.2012.2183853

dimensions. MCFs have recently been demonstrated in both solid-core [7]–[11] and holey-fiber [12], [13] variations. Among solid-core MCFs, a four-core fiber made from 10 μm diameter step-index cores has been demonstrated, achieving 1-Gb/s transmission at 850 nm [8]. Two cores were coupled to a linear vertical-cavity surface-emitting laser (VCSEL) array, and no receiver was demonstrated. Additionally, a graded-index multimode seven-core fiber has been reported, achieving 70-Gb/s transmission by employing custom-fabricated breakout cables constructed to inject/collect light into/from each core [11]. Again, no transceiver was demonstrated.

In this paper, we report a 100 m transmission demonstration using a transmitter and receiver custom-designed to interface with a cleaved facet of the same multimode MCF geometry reported in [11]. In this manner, the cumbersome breakout cable is entirely avoided, as the cores are butt-coupled directly to a matched pattern of VCSELs and photodiodes (PD) at each end facet of the fiber. Previously, our transmitter-only demonstration [14] exhibited a threefold increase over the then-current MMF transmission record of 40 Gb/s, obtained using four CWDM channels [4], [5]. However, constraints in the MCF-interfacing receiver implementation [15] required design modifications—reported here—in order to reach the targeted speed. We now show an end-to-end record transmission of 120 Gb/s by transmitting and receiving six 20 Gb/s signals through the six outer cores of the seven-core fiber. Crosstalk measurements with all channels enabled confirm the viability of this approach.

The remainder of this paper is organized as follows. Section II describes the geometry, fabrication, and performance of the multimode MCF. Section III reviews the assembly and performance of the VCSEL-array transmitter. And, Section IV describes the MCF-interfacing receiver with full-link performance data. Conclusions are drawn in Section V.

II. Multicore Graded-Index MMF

A. Structure

The multimode MCF was made from seven graded-index OFS LaserWave fiber core rods [16] arranged in a hexagonal array. The fiber core diameters are reduced from the standard dimension of 50 to 26 μm in order to maintain a standard cladding diameter of 125 μm. Adjacent cores are spaced by 39 μm, measured from center to center. Fig. 1 includes an image of the fiber cross section.

B. Characterization

Intercore crosstalk was measured on the 100 m fiber sample by scanning a fiber in two dimensions across the input facet of the MCF. The injection fiber carries 850 nm light from a continuous-wave laser source, and has a 4 μm core diameter to ensure single-mode operation. Within the first meter of MCF, tight bends ($\sim$ 1 cm diameter) are implemented in order to instigate strong modal mixing and to strip away cladding modes which may be excited when the light is launched outside the core region. The light is then extracted from the central core at the output facet using a single-core MMF with a 26-μm core diameter. Finally, the power is detected and recorded as a function of the injection fiber position (see Fig. 1). The worst case

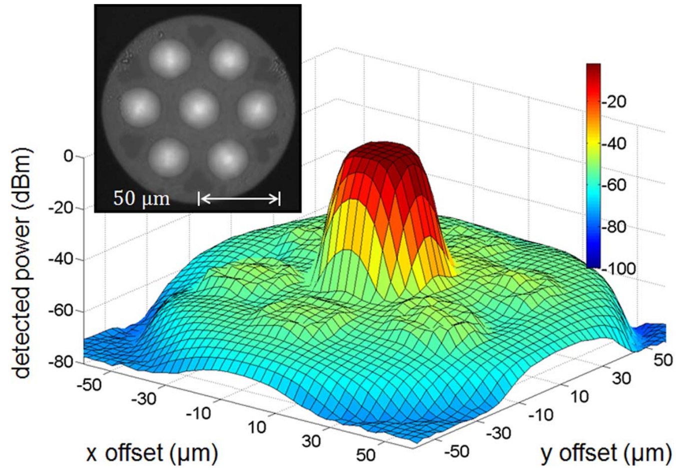

Fig. 1. Core-to-core optical crosstalk for the 100 m MCF sample plotted as the detected power in the central core versus the offset in the two lateral dimensions of the single-mode-fiber launch location relative to the facet center. The inset shows a cross-sectional image of the seven-core fiber.

crosstalk observed from a single-input launch is below -40 dB. Cladding modes, when they are not stripped by the tight bends, may induce additional crosstalk (still less than -20 dB); however, these modes are excited only when light is launched outside of the core region, as would only be observed under extreme misalignment.

Furthermore, differential mode delay (DMD) [17] was measured on a 300 m length of fiber made using a single LaserWave 26 μm diameter core. Results indicate a DMD value of 0.12 ps/m within an 8 μm radial mask width. The measurements were taken using a Ti:Sapphire laser that produces picosecond-scale optical pulses at a wavelength near 850 nm. These pulses were launched into the fiber's input facet from a single-mode (4 μm core diameter) fiber, while the output was collected in a short length of MMF and received using a fast detector (21 GHz bandwidth) followed by a sampling oscilloscope (70 GHz bandwidth). Additional details of the MCF can be found in [11].

III. Flip-Chip Packaged 24-Channel Transceiver

A. Structure

Four custom arrays of VCSELs, arranged in a hexagonal pattern, have been designed and fabricated in order to interface the outer six cores of the MCF [see Fig. 2(a)]. Likewise, a corresponding PD chip has been fabricated [see Fig. 2(b)]. Limitations in wiring density precluded the implementation of fully isolated devices. As a result, the VCSEL and PD arrays are arranged in common cathode configurations. Furthermore, a central device within the array for interfacing the central core in the fiber is omitted due to the same density limitations. The locations of the device arrays and bond pads are designed to be consistent with a previous 24-channel transceiver package [18], [19], so that the high-density carrier and ICs could be reused. The properties of the individual VCSELs and PDs themselves have also been previously reported [20]. However, in order to position the six PDs in the allotted array size, it was necessary to reduce the PD diameters from those previously reported down to 21 μm.

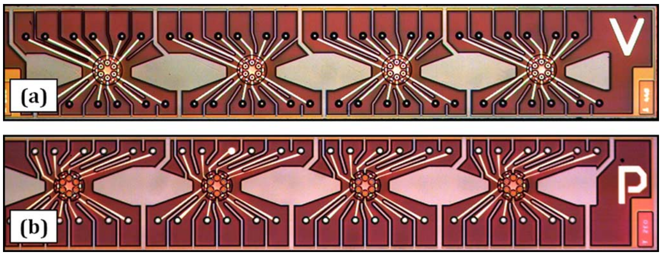

Fig. 2. Die images of custom (a) VCSEL and (b) PD chips for interfacing four strands of the seven-core MCF.

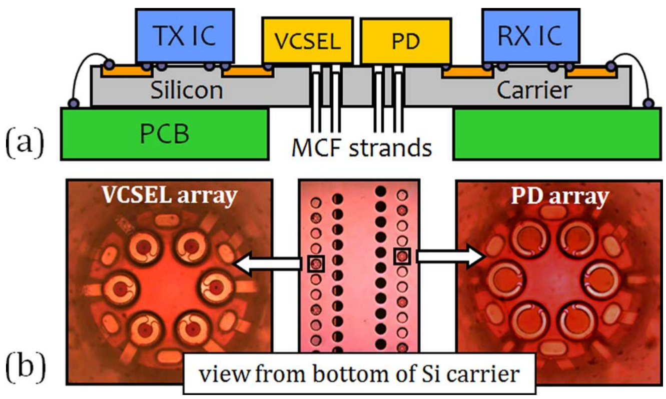

Fig. 3. (a) Transceiver schematic and (b) selected images showing the VCSEL and PD arrays viewed through the silicon carrier's optical vias from the underside of the package.

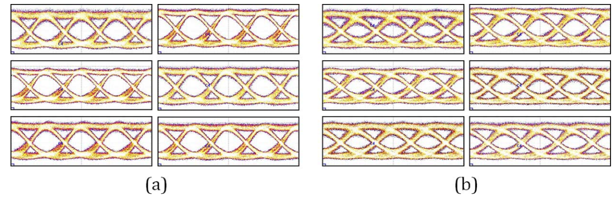

Fig. 4. Eye diagrams recorded one channel at a time at data rates of (a) 15 Gb/s and (b) 20 Gb/s for one of the four transmitter banks. The eyes are recorded following 100 m of fiber using a high-speed detector and sampling oscilloscope. Each image has identical amplitude scales and time scales of 26.7 and 20 ps/div at 15 and 20 Gb/s, respectively.

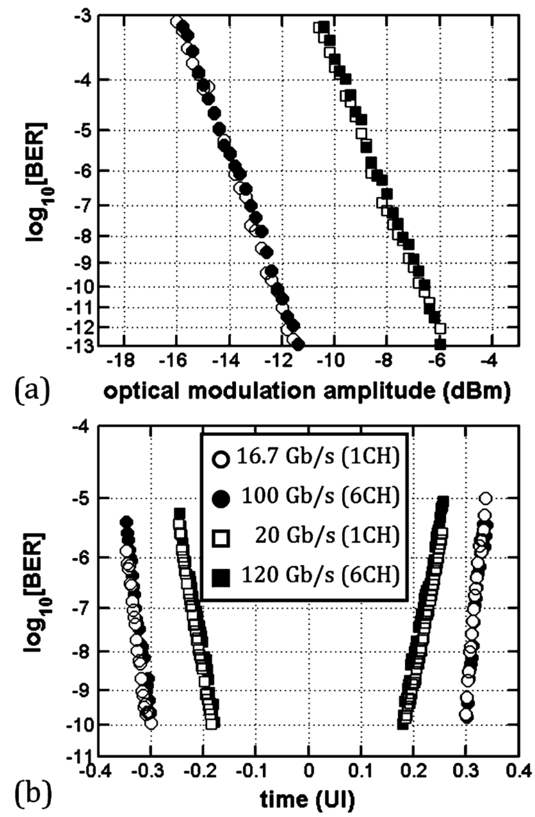

Fig. 5. (a) Amplitude and (b) timing margin BER curves for a channel at 16.7 and 20 Gb/s in the presence of 0 (1CH) and 5 (6CH) aggressors.

The VCSEL and PD array chips were integrated with laser diode driver (LDD) and receiver amplifier ICs in a 24-channel transceiver package [see Fig. 3(a)]. The 24-channel analog ICs were fabricated in an IBM 130 nm CMOS process. These ICs, along with the VCSEL and PD chips, were flip-chip bonded to a silicon carrier. On the carrier surface, the high-speed electrical signals are routed to probe pads, while the bias and control signals are wire-bonded from the carrier's perimeter to a printed circuit board (PCB). Each 24-channel IC is segmented into four six-channel banks. Each bank interfaces with a six-channel VCSEL or PD array, and shares biases internally, reducing the number of control signals required, but eliminating the possibility of enabling and tuning devices independently.

Optical vias or holes are etched through the carrier substrate to allow the MCF to be directly interfaced with the optoelectronic arrays [see Fig. 3(b)]. It is worthwhile to note that the previous transceiver utilized 24 optical vias per direction (i.e., an optical via for each VCSEL and PD). Because we reuse the same carrier, the 24 VCSELs are aggregated into only four optical vias, as are the PDs. This leaves 20 vias unoccupied on the transmit side and 20 unoccupied on the receive side. Further reductions in package area can, thus, be easily envisioned. Additional details about the circuits and assembly processes can be found in previous reports [18], [19].

B. Transmitter Performance

Each of the 24 transmitter channels operated successfully with similar high-speed performance. Eye diagrams are shown at 15 and 20 Gb/s for one bank representing a typical response (see Fig. 4). The eyes remain open up to 20 Gb/s, although intersymbol interference (ISI) can be observed above 15 Gb/s. Each bank of six VCSELs produces a total output power of 13 mW at 3 V supply, indicating more than 2 mW of average power per VCSEL assuming identical performance. Each bank of six transmitter channels dissipates 550 mW, resulting in about 92 mW/channel or 4.6 pJ/bit at 20 Gb/s/channel.

In order to assess optical and electrical crosstalk generated from the MCF transmitter, all channels within one bank are simultaneously modulated. Meanwhile, a cleaved MCF is inserted into the optical via and aligned to the VCSEL array. Crosstalk is measured on one asynchronous channel by toggling the modulation to the other five channels. Data rates of 16.7 and 20 Gb/s/channel (100 and 120 Gb/s aggregate) are investigated. Negligible crosstalk is observed in both the timing- and amplitude-margin bit-error-rate (BER) curves (see Fig. 5). Some of

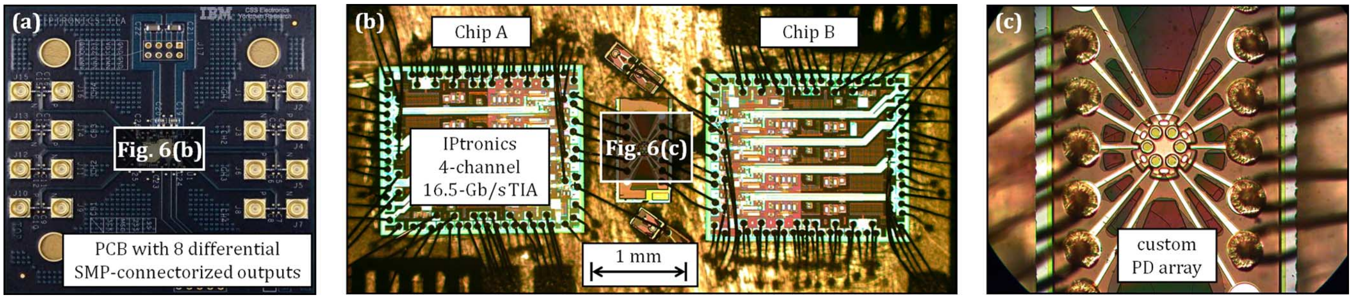

Fig. 6. Images of the wire-bonded MCF receiver implementation including (a) PCB for testing, (b) assembled PD chip and two TIA chips, and (c) magnified view of custom PD array.

the degradation in sensitivity—measured in optical modulation amplitude (OMA)—at increasing data rates can be attributed to ISI generated at the transmitter. However, bandwidth limitations within the high-speed reference receiver used during the characterization also contribute to the loss of sensitivity. As a result, the sensitivity at the maximum data rate of 120 Gb/s is greatly improved in the redesigned six-channel receiver described in Section IV. For the 100 Gb/s and 120 Gb/s aggregate rates, the eye openings at a BER of 10^{-9} are 0.6 and 0.4 unit intervals (UIs). Disabling the five aggressor channels improves the eye opening by only 0.01 UI (0.6 ps) at 16.7 Gb/s and 0.02 UI (1 ps) at 20 Gb/s, indicating that transmitter crosstalk is negligible.

C. Receiver Performance

The transceiver assembly yields 100% of the 24 receiver channels as well. Yet, the receiver IC provides significantly lower bandwidth compared to the transmitter with wide-open eyes up to 10 Gb/s [15]. Furthermore, it suffers from significant interchannel crosstalk, since the receiver IC was not specifically designed to interface with common cathode PD arrays. Consequently, the receiver was redesigned using commercial ICs and a custom wire-bonded assembly described in the following section for the purpose of testing a single end-to-end MCF link.

IV. Redesigned Six-Channel Receiver

The common cathode photodetector array—arranged to optically interface with the MCF—was wire-bonded to two commercial four-channel 16.5-Gb/s amplifier chips available from IPtronics, labeled "Chip A" and "Chip B" in Fig. 6. These three chips have been mounted on a PCB. Three of the inputs to the receiver chips are connected to each side of the PD array, while the fourth receiver channel is wire-bonded to a reference photodetector. The outputs of the receiver are connected to four differential traces that route the digitized signals to SMP connectors.

Initially, the six-channel receiver was tested one channel at a time. To do so, the setup shown in Fig. 7(a) was employed. A single channel of the previously reported MCF transmitter was driven and coupled to a 100 m MCF sample. The activated core was coupled to a single-core 50 μm fiber terminating in a lensed facet at the receiver under test. Each channel was subsequently evaluated on an error detector and sampling oscilloscope. Since

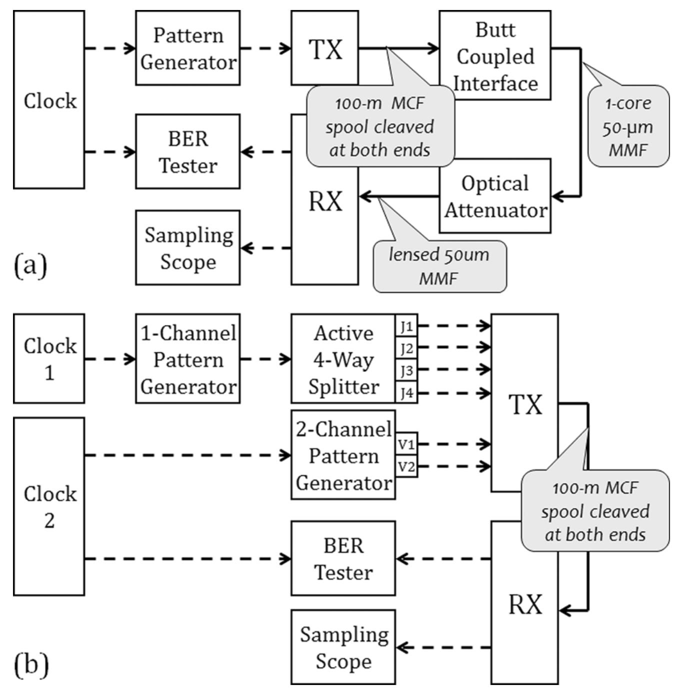

Fig. 7. Full-link test setup during (a) single-channel and (b) multichannel measurements.

the receiver is tested using the MCF transmitter, effectively a single-channel link is demonstrated.

The eye diagrams (see Fig. 8) illustrate that although the receiver is specified for up to 16.5 Gb/s operation only, open eye diagrams may be obtained on all channels at 20 Gb/s. Each channel's sensitivity was measured at a BER of 10^{-12} (see Fig. 9). At 16 Gb/s, the sensitivities in OMA of all six channels fall between -12.3 and -10.8 dBm. At 20 Gb/s, the sensitivities degrade by about 2 dB (-10.5 to -8.8 dBm), due to transmitter and receiver bandwidth limitations.

Finally, the receiver was tested under multichannel operation utilizing the setup shown in Fig. 7(b). Again the MCF receiver measurements represent a full-link result. Here, all channels of the multicore transmitter were driven simultaneously and coupled to the MCF sample. At the end of the 100 m fiber sample, the cores were coupled to the six receiver channels. All receiver output channels were enabled, but only one or two channels were monitored on test equipment at any given time. Two VCSELs—located at opposite ends of the array—were driven

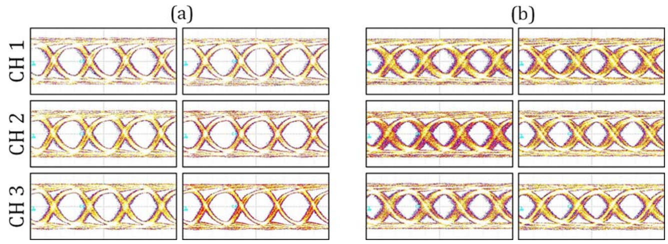

Fig. 8. Eye diagrams for the six receiver channels recorded one at a time at data rates of (a) 16 Gb/s and (b) 20 Gb/s using constant amplitude scales and time scales of 25 and 20 ps/div, respectively.

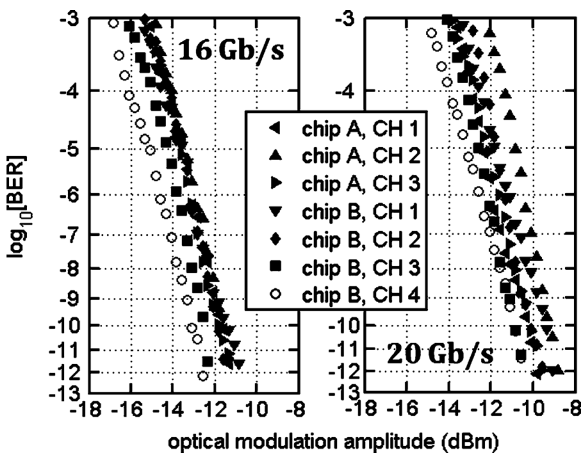

Fig. 9. Receiver sensitivity curves for the six receiver channels recorded one-at-a-time at (a) 16 Gb/s and (b) 20 Gb/s. Chip B, channel 4 is bonded to a 25-μm-diameter reference PD.

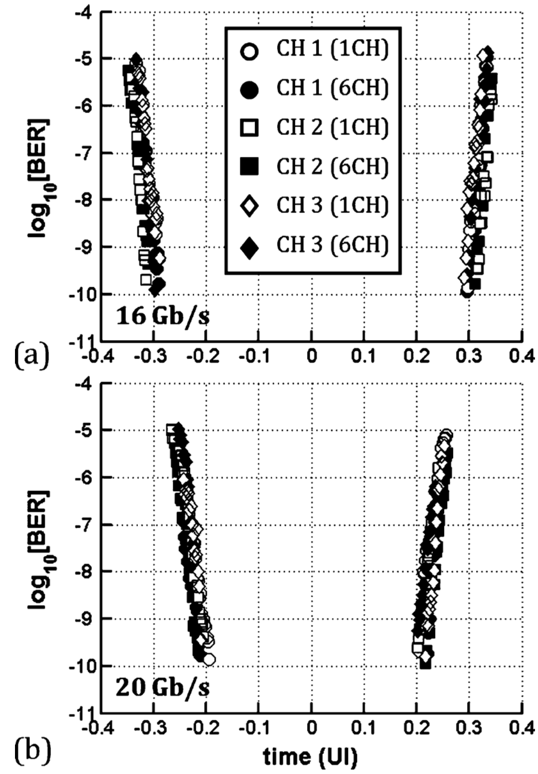

Fig. 10. Timing margin BER curves for the six receiver channels at (a) 16 Gb/s and (b) 20 Gb/s with and without aggressor channels enabled.

synchronously, while the four remaining neighbors were driven asynchronously with a separate clock and pattern generator.

Due to nonuniformity in the VCSEL output power and the transmitter fiber coupling, the OMA exiting each core was recorded at the output facet of the MCF. The per-channel (i.e., per-core) OMA at the fiber end facet varied between -3.0 and -0.8 dBm. One channel having OMA of -2.1 dBm was chosen as the victim channel. The six cores were then coupled simultaneously to the six receiver channels. BER timing margin curves were taken systematically for coupling occurring between the victim channel and each of the three receiver channels for one of the receiver chips (i.e., chip A, channels 1, 2, and 3). In this manner, the data obtained on each receiver channel were obtained using the same transmitter channel making comparisons more meaningful. As in the transmitter characterization, the modulation to the five aggressor channels was toggled in order to evaluate the effect of optical and electrical crosstalk at the receiver.

The measured performance is shown in Fig. 10. The three channels demonstrate uniform performance at 16 Gb/s with wide-open margin ($\sim$ 0.6 UI at 10^{-9} BER) and no apparent degradation due to crosstalk. At 20 Gb/s, the jitter increases; however, all three channels demonstrate margin greater than 0.4 UI at 10^{-9} BER. The effect of the five aggressor channels on the eye opening of the victim channel results in a degradation of only 0.01 UI in the worst case.

The transmitter and receiver dissipate 550 and 650 mW of power, respectively, resulting in 10 pJ/bit of full-link energy efficiency when operating at 120 Gb/s.

V. Conclusion

A 100 m MMF optic link has been demonstrated transmitting 120 Gb/s over a single fiber strand with standard 125 μm cladding while dissipating only 10 pJ/bit. Within the fiber's cladding, seven graded-index cores are arranged in a hexagonal lattice pattern, and the outer six cores of the pattern are interfaced to a custom transmitter and receiver. The transmitter includes an LDD IC and a VCSEL array chip flip-chip bonded to a silicon carrier. The receiver includes two transimpedance amplifier ICs and a PD array chip, all mounted on a custom PCB with wire-bond connections. Interchannel crosstalk due to both optical and electrical effects is shown to be insignificant for both the transmitter alone and for the full link (transmitter plus receiver). This link will provide a sixfold improvement in the bandwidth density at a midpoint passive optical connector over equivalent single-core versions. Although the reutilized silicon carrier package did not demonstrate it here, a similar improvement in bandwidth density at the link's endpoints can be envisioned as well. This density provided by the MCF link may be needed in order to realize efficient next-generation high-performance computers and datacenters, while the larger core sizes compared to single-mode fiber solutions may serve to keep packaging costs sustainable. Nonetheless, continued developments related to efficient connectorization, optical coupling, rotational alignment, and core identification will be required.

REFERENCES

[1] 100 teraFLOPS Dedicated to Capability Computing. Internet:, updated Nov 22, 2010 [accessed May 11, 2011] [Online]. Available: https://asc.llnl.gov/computing_resources/purple/

[2] D. Grice, H. Brandt, C. Wright, P. McCarthy, A. Emerich, T. Schimke, C. Archer, J. Carey, P. Sanders, J. A. Fritzjunker, S. Lewis, and P. Germann, "Breaking the petaflops barrier," *IBM J. Res. Develop.*, vol. 53, pp. 1:1–1:16, 2009.

[3] A. F. Benner, D. M. Kuchta, P. K. Pepeljugoski, R. A. Budd, G. Hougham, B. V. Fasano, K. Marston, H. Bagheri, E. J. Seminaro, H. Xu, D. Meadowcroft, M. H. Fields, L. McColloch, M. Robinson, F. W. Miller, R. Kaneshiro, R. Granger, D. Childers, and E. Childers, "Optics for high-performance servers and supercomputers," in *Proc. Opt. Fiber Commun. Conf.*, 2010, pp. 1–3, paper OTuH1.

[4] 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. IEEE Electron. Compon. Test Conf.*, 2005, pp. 1027–1031.

[5] R. Michalzik, G. Giaretta, K. W. Goossen, J. A. Walker, and M. C. Nuss, "40 Gb/s coarse WDM data transmission with 825 nm wavelength VCSEL's over 310 m of high performance multimode fiber," in *Proc. Eur. Conf. Opt. Commun.*, Munich, Germany, Sep. 2000, pp. 33–34.

[6] A. Tarighat, R. C. J. Hsu, A. Shah, A. H. Sayed, and B. Jalali, "Fundamentals and challenges of optical multiple-input multiple-output multimode fiber links," *IEEE Commun. Mag.*, vol. 45, no. 5, pp. 57–63, May 2007.

[7] S. Inao, T. Sato, S. Sentsui, T. Kuroha, and Y. Nishimura, "Multicore optical fiber," presented at the presented at the Opt. Fiber Commun. Conf., 1979, Paper WB1.

[8] B. Rosinski, J. W. D. Chi, P. Grosso, and J. Le Bihan, "Multichannel transmission of a multicore fiber coupled with vertical-cavity surface-emitting lasers," *J. Lightw. Technol.*, vol. 17, no. 5, pp. 807–810, May 1999.

[9] K. Imamura, K. Mukasa, and T. Yagi, "Investigation on multi-core fibers with large Aeff and low micro bending loss," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2010, pp. 1–3.

[10] K. Takenaga, S. Tanigawa, N. Guan, S. Matsuo, K. Saitoh, and M. Koshiba, "Reduction of crosstalk by quasi-homogeneous solid multi-core fiber," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2010, pp. 1–3.

[11] B. Zhu, T. F. Taunay, M. F. Yan, M. Fishteyn, G. Oulundsen, and D. Vaidya, "70-Gb/s multicore multimode fiber transmissions for optical data links," *IEEE Photon. Technol. Lett.*, vol. 22, no. 22, pp. 1647–1649, Nov. 2010.

[12] B. J. Mangan, J. C. Knight, T. A. Birks, P. S. J. Russell, and A. H. Greenaway, "Experimental study of dual-core photonic crystal fibre," *Electron. Lett.*, vol. 36, no. 16, pp. 1358–1359, Aug. 2000.

[13] K. Imamura, K. Mukasa, Y. Mimura, and T. Yagi, "Multi-core holey fibers for the long-distance ($>$ 100 km) ultra large capacity transmission," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2009, pp. 1–3, paper OTuC3.

[14] B. G. Lee, D. M. Kuchta, F. E. Doany, C. L. Schow, C. Baks, R. John, P. Pepeljugoski, T. F. Taunay, B. Zhu, M. F. Yan, G. E. Oulundsen, D. S. Vaidya, W. Luo, and N. Li, "120-Gb/s 100-m transmission in a single multicore multimode fiber containing six cores interfaced with a matching VCSEL array," in *Proc. IEEE Photon. Soc. Summer Top. Meet.*, Jul. 2010, pp. 223–224.

[15] B. G. Lee, D. M. Kuchta, F. E. Doany, C. L. Schow, C. Baks, R. John, P. Pepeljugoski, T. F. Taunay, B. Zhu, M. F. Yan, G. E. Oulundsen, D. S. Vaidya, W. Luo, and N. Li, "Multimode transceiver for interfacing to multicore graded-index fiber capable of carrying 120-Gb/s over 100-m lengths," in *Proc. IEEE Photon. Soc. Annu. Meet.*, Nov. 2010, pp. 564–565.

[16] [Online]. Available: http://www.ofsoptics.com/resources/LaserWave-550-300-web.pdfLaserWave 550/300 Multimode Fibers. Internet:, [accessed July 26, 2011]

[17] R. Olshansky and S. M. Oaks, "Differential mode delay measurement," presented at the presented at the Eur. Conf. Opt. Commun., Genoa, Italy, 1978, Paper III.1.

[18] L. Schares, J. A. Kash, F. E. Doany, C. L. Schow, C. Schuster, D. M. Kuchta, P. K. Pepeljugoski, J. M. Trewhella, C. W. Baks, R. A. John, L. Shan, Y. H. Kwark, R. A. Budd, P. Chiniwalla, F. R. Libsch, J. Rosner, C. K. Tsang, C. S. Patel, J. D. Schaub, R. Dangel, F. Horst, B. J. Offrein, D. Kucharski, D. Guckenberger, S. Hegde, H. Nyikal, C.-K. Lin, A. Tandon, G. R. Trott, M. Nystrom, D. P. Bour, M. R. T. Tan, and D. W. Dolfi, "Terabus: Terabit/second-class card-level optical interconnect technologies," *IEEE J. Sel. Topics Quantum Electron.*, vol. 12, no. 5, pp. 1032–1044, Sep./Oct. 2006.

[19] F. E. Doany, C. L. Schow, C. K. Tsang, N. Ruiz, R. Horton, D. M. Kuchta, C. S. Patel, J. U. Knickerbocker, and J. A. Kash, "300-Gb/s 24-channel bidirectional Si carrier transceiver optochip for board-level interconnects," in *Proc. Electron. Compon. Technol. Conf.*, May 2008, pp. 238–243.

[20] N. Y. Li, C. L. Schow, D. M. Kuchta, F. E. Doany, B. G. Lee, W. Luo, C. Xie, X. Sun, K. P. Jackson, and C. Lei, "High-performance 850 nm VCSEL and photodetector arrays for 25 Gb/s parallel optical interconnects," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2010, pp. 1–3.

Benjamin G. Lee (M'04) received the B.S. degree from Oklahoma State University, Stillwater, in 2004, and the M.S. and Ph.D. degrees from Columbia University, New York, in 2006 and 2009, respectively, all in electrical engineering.

In 2009, he became a Postdoctoral Researcher at IBM Thomas J. Watson Research Center, Yorktown Heights, NY, where he is currently a Research Staff Member. He is also an Assistant Adjunct Professor of electrical engineering at Columbia University. His research interests include silicon photonic devices, integrated optical switches and networks for high-performance computing systems and datacenters, and highly parallel multimode transceivers.

Dr. Lee is a member of the IEEE Photonics Society and the Optical Society of America. He has served on the technical program committee for the Fourth, Fifth, and Sixth ACM/IEEE International Symposium on Networks-on-Chip.

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

He subsequently joined IBM at the Thomas J. Watson Research Center, Yorktown Heights, NY, where he is currently a Research Staff Member in the Communication Technology Department, and involved in high-speed vertical-cavity surface-emitting lasers characterization, multimode fiber links, and parallel fiber optic link research. He is the author or coauthor of more than 50 technical papers and inventor/coinventor of more than 10 patents.

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

From 1984 to 1985, he was a Postdoctoral Fellow at the California Institute of Technology. He subsequently joined IBM at the T.J. Watson Research Center, Yorktown Heights, NY, where he was involved in laser spectroscopy, applied optics, projection displays, and laser material processing for electronic packaging. Since 2000, his research has been focused on optoelectronic packaging and high-speed optical link and systems design. He is the author or coauthor of many technical papers and holds over 40 U.S. patents.

Clint L. Schow (SM'10) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Texas at Austin, Austin, in 1994, 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, where he was involved in the development of high-speed optoelectronic modulators and tunable laser sources for optical communications. In 2004, he joined the IBM T.J. Watson Research Center as a Research Staff Member, where he is currently involved in research on parallel optical interconnect technologies and high-speed CMOS circuits for fiber-optic data links.

Petar Pepeljugoski, biography not available at the time of publication.

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

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

Thierry F. Taunay received the Ph.D. degree in physics from the University of Lille, Lille, France, in 1997.

He subsequently joined the Naval Research Laboratory, Washington, DC, where he was involved in fiber photosensitivity, fiber-optic sensors, radiation effects on optical fiber waveguides and optical materials, and fundamental defect centers in glass. Since 2007, he has been a Member of the Technical Staff in OFS Laboratories, Somerset, NJ. He has authored or coauthored more than 60 journal and conference papers. His current research interests include design and processing of new transmission fibers and specialty fibers for high-power fiber lasers.

Dr. Taunay is a member of the Optical Society of America and served two terms as an Associate Editor of *Optics Express*.

Benyuan Zhu received the Ph.D. degree in applied physics from Bath University, Bath, U.K., in 1996.

In 1999, he joined Bell Laboratories, Holmdel, NJ. He is currently a Distinguished Member of Technical Staff in OFS Laboratories, Somerset NJ. He has been mainly involved in high-speed dense wavelength-division multiplexing transmission systems. He has authored or coauthored more than 100 journal and conference papers, and one book chapter in the field of optical fiber communications. His current research interests include 100G & above optical coherence transmission systems, novel fiber, and advanced optical amplifier technologies.

Man F. Yan received the S.B. and Sc.D. degrees from the Massachusetts Institute of Technology, Cambridge, and the M.Eng. degree from the University of California, Berkeley.

In 1976, he joined Bell Laboratories where he later became a Distinguished Member of Technical Staff and Technical Manager. In 2001, he joined OFS Laboratories when it acquired the Optical Fiber Division from Lucent Technologies. He is currently a Technical Manager at OFS Laboratories, Murray Hill, NJ. His current research interests include design and processing of transmission and specialty fibers.

George E. Oulundsen received the Ph.D. degree in chemical engineering from the University of Massachusetts, Amherst.

He is currently a Distinguished Member of Technical Staff in the R&D Group, OFS, Sturbridge, MA. He was involved in the development of various multimode processes and products at OFS over the past 13 years, and has spent the last several years improving multimode fiber measurements, and developing new products and processes.

Durgesh S. Vaidya received the B.S. degree in chemical engineering from the Indian Institute of Technology, Mumbai, India, and the Ph.D. degree in chemical engineering from the State University of New York, Buffalo.

He is currently a Senior Manager at OFS, Sturbridge, MA, where his research is focused on multimode optical fiber and optical connectivity solutions. At Bell Labs and now OFS, he has led many product and process development projects introducing new design, materials, and fabrication methods to improve optical system performance.

Wenlin Luo, biography not available at the time of publication.

Neinyi Li received the B.S. and M.S. degrees in electrical engineering from National Cheng-Kung University, Tainan, Taiwan, in 1988 and 1990, respectively, and the Ph.D. degree in electrical and computer engineering from the University of California, San Diego, in 1997.

In 1996, he had an Internship at Sandia National Laboratories, Albuquerque, NM. From 1997 to 1999, he was an Assistant Research Professor at the University of New Mexico. In 1999, he subsequently joined Emcore Photovoltaic Division, Albuquerque, NM, as a Staff Scientist, where he was involved in the development of quadruple-junction solar cells, 1.3 mm vertical-cavity surface-emitting lasers (VCSELs), and low-turn-on voltage heterojunction bipolar transistors using InGaAsN. He is currently a Senior Director of Emcore Digital Product Division, Albuquerque, where he is involved in high-speed VCSEL/photodiode chip development/production for high-speed fiber optic links and Mach–Zehnder modulator/RTD chips development/production for telecom tunable modules.