PARALLEL FIBER-OPTIC SCI LINKS

David R. Engebretsen

University of Minnesota

Daniel M. Kuchta

Richard C. Booth

John D. Crow

Wayne G. Nation

IBM

Parallel fiber-optic links promise to be a key enabling technology for distributed computing, substantially increasing transmission distance and bandwidth.

The drive to include more and more processors and I/O devices in parallel-processor complexes has created a need for box-to-box interconnection links capable of delivering bandwidth comparable to system backplanes. Point-to-point links overcome the speed and cost constraints of shared buses. However, copper-based point-to-point links can reliably carry high-speed data only a limited distance. Parallel fiber-optic links substantially increase transmission distance and surpass copper's bandwidth capability.

Our research addresses the integration of parallel fiber-optic data links into system designs that require a high-bandwidth link over moderate distances. To demonstrate the feasibility of using parallel fiber-optic technology as a fundamental building block in large-scale commercial and parallel machines, we constructed a link testbed. This testbed integrates Optoelectronics Technology Consortium's (OETC) parallel fiber-optic technology and IBM's high-bandwidth Scalable Coherent Interface-Link (SCI-Link) technology. (OETC is an ARPA-sponsored industrial consortium of AT&T, IBM, Honeywell, and Lockheed Martin, chartered to develop optical bus technology.) OETC's prototype transmitter and receiver modules comprise an optical link with a 2-Gbytes-per-second transmission rate and transmission distances of up to 100 meters.

Based on the IEEE 1596-1992 SCI standard,[1] IBM SCI-Link is a general-purpose communications protocol for multiple system functions, including message passing, I/O bus, and various forms of shared memory. Like the OETC link, SCI-Link has a transmission rate of 500 Mbits per second per line. To improve data integrity, IBM modified the SCI standard protocol by including in the SCI-Link protocol a 32-bit CRC (cyclic redundancy code), hardware packet retry, duplicate packet suppression, and end-to-end acknowledgments. IBM designed the SCI-Link protocol for either copper or parallel fiber-optic medium; hence, the same SCI-Link module design has been successful in both copper and parallel optical testbeds.

The OETC/SCI-Link testbed, which began operating in September 1994, is, to our knowledge, the first functional demonstration of the reliable framing and transmission of SCI packets over a parallel optical bus. It has demonstrated low bit-error rates (less than 10^{-15}) over distances 10 times that of a copper bus with microcoaxial cables, using a fiber-ribbon cable 10 times smaller than the copper cable. The testbed not only demonstrates this emerging technology's feasibility but promises computer system designers new layout flexibility.

High-bandwidth distributed computing

Some forms of distributed computing require high-bandwidth links, supporting distances of tens of meters. In the context of our work, high bandwidth means 1.0 Gbyte per second and higher. Obviously, this rate will increase as the technology develops. High-bandwidth links allow system processors and I/O devices packaged in multiple physical boxes to meet performance requirements. Shared buses can meet the same high-bandwidth requirements, but they cannot achieve the distances distributed computing requires.

High-bandwidth, moderate-distance link applications exist in two general classes of system topologies: processor-to-processor and processor-to-I/O interconnections. High-bandwidth, moderate-distance, processor-to-processor links enable coupling, which allows processor scalability beyond a traditional symmetric multiprocessor topol-

0772-1732/96/$5.00 © 1996 IEEE

ogy. Systems using highly scalable, point-to-point topology fall into two classes: shared-memory, nonuniform-memory-access (NUMA) systems and clustered, nonorthogonal, random-memory-access (NORMA) systems. Current commercial systems using point-to-point links with one of these topologies include the Cray T3D, the Convex Exemplar, and the IBM SP series. Numerous additional NUMA and NORMA machines will enter the market in the near future.

High-bandwidth, moderate-distance links can also increase I/O connectivity by enabling the I/O bus to escape the central processor complex and extend to attached or remote I/O subsystems. An example of a commercial product using serial fiber to increase I/O connectivity is the IBM AS/400 fiber-optic I/O bus.[2]

Why parallel fiber optics?

Serial fiber-optic links are already widely used in data communications and I/O channel connections, where data transfer rates reach 100 Mbytes per second, and distances can reach 10 km. In this environment, fiber's low loss and high bandwidth are its main advantages. In moderate-distance applications, especially those which transmit a significant fraction of the processor bus bandwidth and whose latency requirements limit distances to under 100 meters, parallel fiber-optic synchronous buses are practical. Data transfer rates in excess of 1 Gbyte per second are possible for distances up to 100 meters. Keeping the data in parallel format even in optical cables eliminates the complexities of multiplexing.

Current serial optical technology cannot provide the high bandwidth needed for certain applications. Processor-to-remote-I/O links require bandwidths exceeding 250 Mbytes per second, and this requirement will increase. For processor-to-processor applications, the key requirement is low communication latency for shared-memory NUMA and message-passing cluster applications. A fundamental approach to providing low latency is to transmit with high bandwidth. Hence, processor-to-processor links require bandwidths of up to 1 Gbyte per second. This requirement will also increase in the future.

Fiber-optic links offer several advantages over copper-based links, including greater transmission distances, reduced cable and connector bulk, and improved electrical isolation. At a transmission rate of 500 Mbits per second per line, parallel copper cables have a practical limit of about 10 meters, whereas parallel fiber-optic cables can reach up to 100 meters. This practical limit is set by the increasing price and bulk of the higher quality cables needed to achieve additional incremental distance.

System designers increase the processing power and I/O capability of distributed systems by adding interconnected physical boxes, such as racks, frames, and towers. Each additional box contains more processors and memory, more I/O capability, or both. We measure transmission distance from electronic module to electronic module. Internal cable routing to escape a box may consume as much as 2 meters. Hence, with copper's 10-meter maximum distance, all boxes must be within about 6 meters of each other in a fully connected configuration. Therefore, adding more than about eight boxes, while allowing for maintenance access, becomes impractical. Interconnecting more boxes requires greater cable lengths. An alternative is to organize boxes in a daisy chain. But this topology produces additional latency by using repeaters or switches unnecessarily, and it reduces reliability. Thus, fiber's greater transmission distance, along with its other advantages, positions parallel fiber-optic links as a key enabling technology for distributed computing.

Fiber-optic links offer several advantages over copper-based links, including greater transmission distances, reduced cable and connector bulk, and improved electrical isolation.

OETC link

The OETC project has proceeded in two phases. Phase 1's goal was to build prototype parallel optical transmitter and receiver modules and construct a user testbed. By user testbed, we mean a testbed interfaced to a bus simulator (or an operational interconnection bus) and operated by system designers (not optoelectronic technicians). Phase 2's goal is to produce modules in quantity for premanufacture testing. We briefly describe the OETC technology here; Wong et al.[3] give details.

The OETC link is a 32-channel fiber-optic bus consisting of a transmitter module, a fiber ribbon cable with connectors, and a receiver module. The modules accept logic level synchronous data lines and a clock line, and recreate the logical data lines at the remote output.

Transmitter module. The transmitter module houses the driver chip and laser emitter chip. The driver chip receives incoming data with negative, differential ECL (emitter-coupled logic) inputs, retimes the data, Manchester encodes it, and converts voltage levels into current levels to drive the lasers. OETC fabricated the driver chip using an enhancement/depletion mode GaAs MESFET (gallium arsenide metalized semiconductor field-effect transistor) process with gate lengths less than 1 μm. Each channel can accept non-return-to-zero (NRZ) data at up to 500 Mbits per second; the lasers are driven with Manchester-encoded data at up to 1 Gbit per second.

A single input latch, with setup and hold times that permit as much as ± 500 ps of skew on the incoming data, performs data retiming. Of the 32 transmitter channels, one is a dedicated clock channel. The driver chip delays the clock channel to position the clock edge at the center of the data signal. For diagnosis, the clock delay is adjustable up to 400 ps in 100-ps increments. The transmitter components' combined power dissipation is 7W, and they require four voltage levels.

The driver chip's laser driver is a shunt-current steering design. It accepts two feedback voltage signals from the lasers

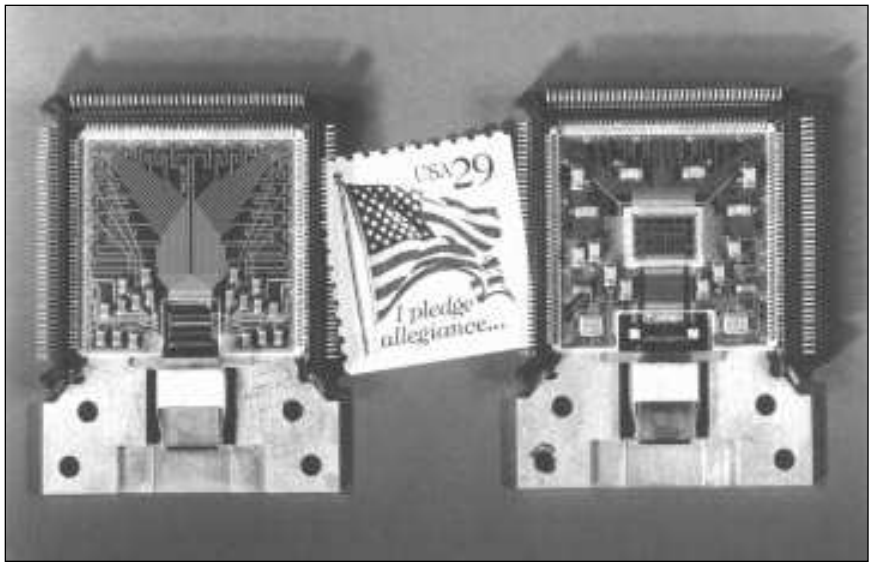

Figure 1. OETC receiver (left) and transmitter (right) modules.

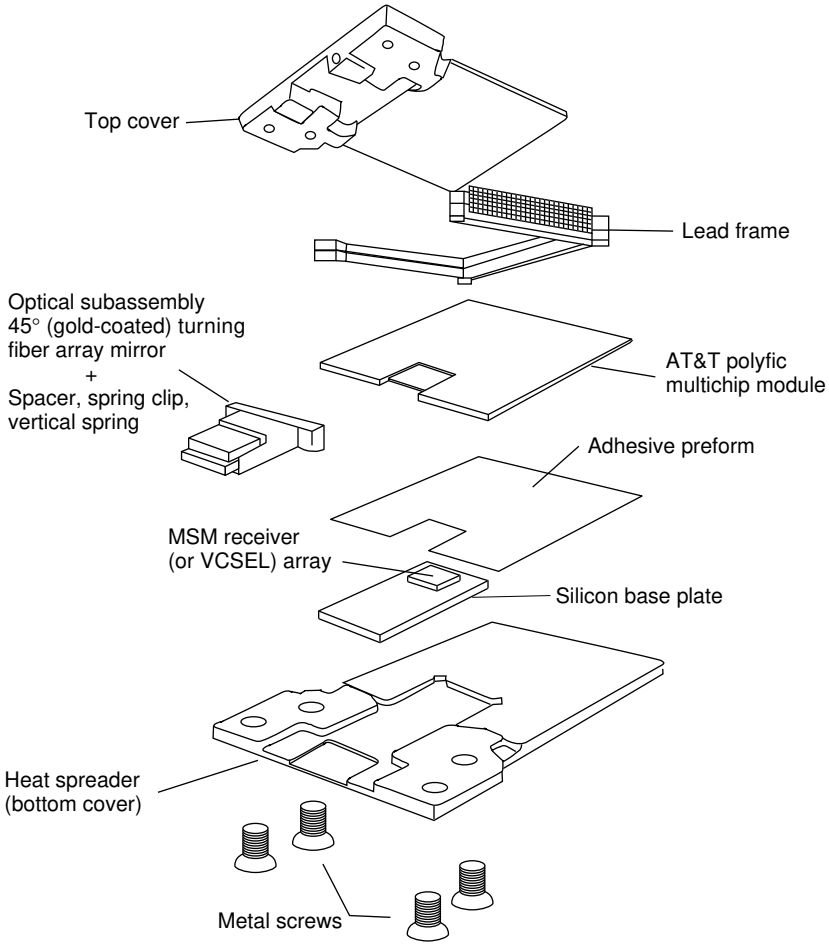

Figure 2. OETC module package.

to set the high- and low-level laser modulation currents and track them over temperature variations. The lasers are 850-nm, vertical-cavity, surface-emitting lasers (VCSELs). Morgan et al.[4] describe their fabrication and structure. OETC chose multitransverse-mode devices because they are capable of speeds greater than 3 G bits per second,[5] have high output power, and have coherence low enough to minimize modal noise.[6,7] An advantage of VCSELs is their low-cost wafer level testability. To demonstrate this, OETC constructed a computer-controlled, wafer level test station with 4-second test times per laser. The test station has shown an array yield of greater than 90 percent.

The VCSEL chip has 34 lasers—32 centered on the chip with a 140-µm pitch for output and one on each end of the chip for laser bias control. Two silicon photodiodes, mounted over the extra VCSELs, control laser bias. One control laser is biased with a DC current corresponding to a logic-one level; the other, with DC current corresponding to a logic-zero level. The laser driver chip amplifies DC photocurrents from the photodiodes and feeds them back to the driver chip, permitting operation of the lasers from 10° C to 70° C.

The VCSELs launch light into the fiber through a coupler that uses two silicon chips with V grooves to hold 32 graded-index multimode fibers. One end of this assembly is ground and polished to a 45-degree angle and coated with gold to increase reflectivity. This end sits over the VCSELs. The coupler's other end mates directly to a MACII-32 cable connector without pigtail. Coupling efficiency is better than 80 percent with a 12.5-µm alignment tolerance.

Receiver module. The receiver chip is a 32-channel optoelectronic integrated circuit, fabricated with standard 1-µm GaAs MESFET technology.[8,9] Each data channel incorporates a metal-semiconductor-metal (MSM) detector, preamplifier, postamplifier, level-restore circuit, decision/decoder circuit, and an off-chip driver. The fiber-to-MSM coupling element is the same type as the laser-to-fiber coupling element. Each data

channel receives Manchester-encoded data at rates of up to 1 Gbit per second. The receiver chip performs data retiming, Manchester decoding, and converting to differential ECL logic levels.

The entire receiver chip is fully testable at the wafer level. On a computer-controlled, wafer level test station, testing an optical channel for rise and fall times, optical sensitivity, delay, jitter, and output amplitude levels took less than 20 seconds, with array yields of better than 90 percent. Each channel's average sensitivity is –19 dB relative to 1 mW with a standard deviation of 1.0 dB (sample size is 10,000). The chip consumes only 2W and can operate on a single power supply.

Data retiming, maximized to allow the link to work to at least 100 meters, removes laser turn-on and fiber skew. The chip retimes data by receiving and distributing the clock channel to latches on the 31 data channels. The total skew of data coming off the package is less than 200 ps. The delay through the receiver module is 2 ns.

Figure 1 shows the OETC transmitter and receiver modules with the covers removed. The white-edged rectangle in the center of the transmitter module is the laser driver integrated circuit. The optical components are just below center on both modules.

Packaging and fiber. The OETC driver, laser emitter, and receiver chips are mounted in modified 164-pin JEDEC (Joint Electron Device Engineering Council), premolded-plastic, quad flat packs and interconnected with a polymer film integrated circuit (polyfic) chip carrier. On both the transmitter and receiver packages, we remove the leads from one side to allow connection to the fiber-optic ribbon. The lead pitch is 25 mils along the other three sides, resulting in 123 leads on a module package of 1×1.5 inches. The polyfic for the transmitter module contains integrated thin-film resistors for termination of the differential inputs. The polyfic for both the receiver and transmitter modules contains solderable pads for discrete decoupling capacitors. Figure 2 illustrates the package components.

The fiber in the OETC link is a ribbon of thirty-two 62.5/125 graded-index multimode fibers on a 140-µm pitch. With a custom, 5-µm polyimide protective coating, the fibers' diameter is 135 µm. The fiber ribbon uses Berg Electronics' enhanced MACII multifiber array connector.[10]

Module testing. Before inserting them into the testbed, we put the OETC modules through several screening tests involving bit-error rate measurements on 20 data channels simultaneously. We did this by feeding a pseudorandom pattern into one transmitter channel and feeding the corresponding output from the receiver back into the adjacent transmitter channel. This enabled us to monitor a single bit error on any of the 20 channels. The links went through temperature cycling and voltage variations for as long as 88 hours without error. During some of these tests, we fed the unused channels a pseudorandom data pattern from a separate source to simulate cross-talk effects.

SCI-Link

IEEE standard SCI, the basis of IBM's SCI-Link, specifies a comprehensive scheme for building scalable parallel-processing systems. It includes a high-level protocol for coupling processors and memories in shared-memory environments. The protocol allows maintenance of the distributed, shared memory as a cache-coherent single address space. The standard also supports a message-passing paradigm.

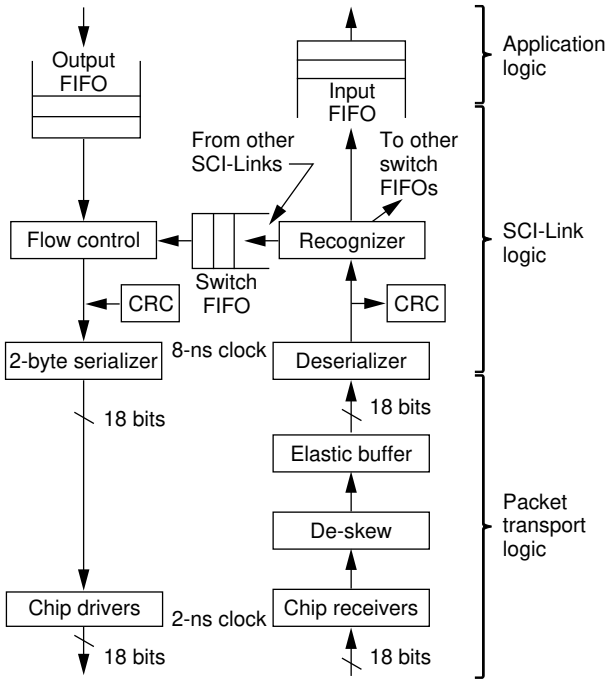

Figure 3. SCI-Link architecture.

Both IEEE SCI and IBM's SCI-Link use signaling technologies based on IEEE P1596.3, Low-Voltage Differential Signaling (LVDS). Thus, results from our testbed directly apply to LVDS designs compatible with P1596.3, independent of compliance with the SCI logical protocols.

Figure 3 diagrams the components of the SCI-Link architecture. The application logic generates requests and responses, which queue for transmission in the output FIFO. The flow-control logic removes commands from the output FIFO, formats the packet with correct CRC, and sends the packet at a rate of 8 bytes per 8-ns cycle to the 2-byte serializer. The flow-control logic also accepts and transmits packets received at other interfaces (if this SCI-Link is one of several on a switch chip). Acknowledgment packets and other flow-control packets also originate from the flow-control logic. The 2-byte serializer forwards the packet at a rate of 2 bytes per 2-ns cycle to the chip drivers, which drive the packet off chip.

SCI-Link transmits differential NRZ (that is, unencoded) data at a rate of 500 Mbits per second per logical bit. The data channel's width is 16 data bits plus a clock and a flag line, yielding a link bandwidth of 1 Gbyte per second. When packets arrive at the receiver end of the link, the chip receivers send unlatched packet contents to the de-skew logic, which dynamically adjusts each bit line to remove bit-

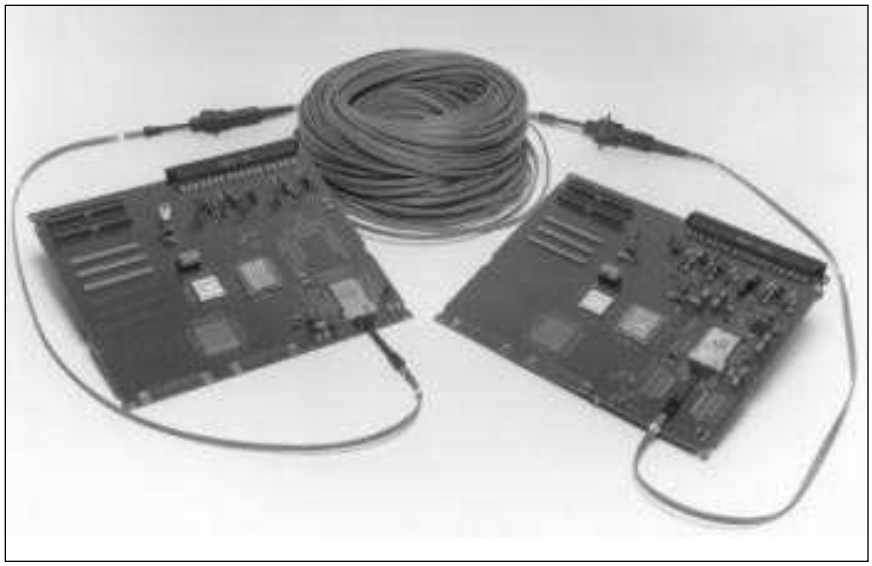

Figure 4. Testbed receiver (left) and transmitter (right) cards with 100-meter cable.

to-bit skew induced by the interconnection.[11] Special synchronization packets, transmitted intermittently, dynamically calibrate this logic. The skew observed on fiber cables is low compared to copper cables (1-10 ps per meter versus 40 ps per meter) and quite stable environmentally. Hence, this chip set uses logic primarily to compensate for skew that might result from remote, arbitrary placement of the OETC modules with respect to the SCI-Link modules on the card.

Another important feature of the SCI-Link modules is the elastic buffer, which allows multiple systems built around an SCI-Link to tolerate differences in nominal clock frequency. The elastic buffer resynchronizes the packet into the local chip's clock domain, using intermittent null symbols on the link as elasticity symbols. Then, the deserializer latches the packet contents into the 8-ns clock domain.

Next, the packet recognition logic (the recognizer) determines if a packet is destined for the application logic of this port or another port on the chip. If destined for this port, the packet queues into the input FIFO. If destined for another port, it queues into that port's switch FIFO.

We built the SCI-Link chip used in the testbed with IBM's 0.8-µm bi-CMOS fabrication process. We packaged the sparsely populated 12.7×12.7-mm die in a 32-mm solder ball carrier. The chip contains a pseudorandom number generator, which creates packet data for a synthetic packet workload and forms the pseudorandom data into correctly framed packets with valid CRC. The packet recognizer contains separate counters for packets received with good CRC and bad CRC. Once the recognizer calculates a packet's CRC, it discards the packet contents. Thus, the testbed's SCI-Link chip is a link exerciser rather than a full-protocol chip. It does not contain logic for high-level protocols or for switching packets from one port to another.

OETC/SCI-Link testbed

We constructed the OETC/SCI-Link testbed to test the OETC modules' effectiveness as the physical medium for a SCI-Link system and to compare parallel fiber optics with parallel copper links. The testbed consists of an optical transmitter card, a parallel fiber-optic cable, an optical receiver card, and a controlling workstation. We use the workstation to initialize data packet generation in the optical transmitter card and to monitor error counts in the optical receiver card. Figure 4 shows the optical card set and a 100-meter parallel fiber-optic cable. Figure 5 shows a close-up of the OETC receiver module mounted on the testbed card.

We designed a card layout that could operate as either an OETC transmitter or receiver. Constructed with standard, production level, epoxy-resin technology, the card has eight signal planes, six power planes, and 50-ohm, single-ended impedance. Because the testbed operates at a 500-MHz clock rate, we designed the card carefully to minimize the effects of skew, cross talk, and signal reflections. We controlled the length of all 500-MHz signal traces to reduce the maximum bit-to-bit skew between the SCI-Link module and the OETC modules to less than 100 ps.

Different teams, with different initial objectives, designed the OETC modules and the SCI-Link module. We defined the OETC/SCI-Link testbed only after each team met its initial objectives. Hence, constructing a testbed using the modules required some additional effort.

A challenge common to both the optical transmitter card and the optical receiver card was interfacing the OETC and SCI-Link electrical levels. The OETC modules use ECL logic, while SCI-Link uses IEEE P1596.3 LVDS. Hence, a direct interface between modules was not possible. Rather than using complex, 500-MHz, level-shifting logic, we level-shifted all power supplies for the OETC modules by a positive 2.5V with respect to ground. This approach put the voltage transfer characteristics at compatible levels and allowed direct module-to-module connection of the 500-MHz signals.

The optical transmitter card consists of an SCI-Link module operating in transmit mode, an OETC transmitter module, and support circuitry. The SCI-Link module generates correctly framed packets of pseudorandom data with a valid CRC. The SCI-Link module feeds the entire 18-bit (16 data bits, 1 flag, and 1 clock bit) output directly to the OETC transmitter module, which transmits the bits on 18 of the 31 OETC data channels. SCI-Link uses a dual-edge clock, consistent with the IEEE SCI standard, while OETC uses a single-edge clock. To resolve this inconsistency, we transmit the SCI-Link clock over the optical link as an ordinary data bit. We take the OETC transmitter clock from the same 500-MHz clock that drives the SCI-Link chip and transmit it as the optical clock that latches data at the OETC receiver. To account for

on-card wiring delays, we phase-shifted the 500-MHz clock by trimming the coaxial cables that carry the clock signal.

The optical receiver card consists of an SCI-Link module operating in receive mode, an OETC receiver module, and support circuitry. The OETC receiver module receives the 19 incoming optical signals and clocks them with the single-edge optical clock. The 18 original SCI-Link signals then go to the SCI-Link module.

Testbed results

Laboratory testing of the full testbed has demonstrated reliable operation at 512 MHz over 100 meters of parallel fiber. The longest error-free interval was 148 hours, translating to a bit-error rate of less than 10^{-15}. Sustained testing resulted in an average time between CRC errors of 33 hours. These results are conservative because longer parallel fiber cables were not available for testing, and the SCI-Link modules do not operate reliably above 512 MHz. Therefore, we could not determine the OETC modules' breaking point.

The testbed allows us to compare copper- and fiber-based links on the same card design and in the same environment. The SCI-Link modules driving a parallel copper link succeeded in operating without error up to 10 meters at 512 MHz. To achieve greater distance with copper, we had to decrease the clock rate. Thus, achieving higher speeds requires sacrificing distance.

PHASE 2 PRODUCTION of the OETC modules shows high yield (over 90 percent) for the laser and receiver chip set. Thus, for the first time, we see a potential for making optical-array line drivers and receivers as inexpensively as their electrical counterparts. We expect organizations to install more than 50 OETC buses in user testbeds in telecommunications switching and computer interconnection applications. The OETC/SCI-Link testbed gave our system designers the opportunity to generate a list of desirable features for the next generation of optical modules. Among these features are direct LVDS logic interfaces, which we have included in the phase 2 modules.

A key to the acceptance of optical buses is that they must be cost-competitive with copper technology. Packaging is crucial to low cost. A consortium of IBM, 3M, and Lexmark is developing Jitney, a follow-up to OETC, to advance inexpensive optical-bus packaging. Vitesse, Honeywell, and IBM are supplying chips for Jitney, and the NIST-ATP (National Institute of Standards and Technology's Advanced Technology Program) is partially funding the project. Jitney combines the OETC chip set with plastic, lead-frame IC technology. It uses molded-plastic optical elements to couple optoelectronic chips to fibers, a snap-together optical-alignment scheme, and an optical cable with connectors attached on the manufacturing line. Also under development are CMOS SCI-Link chips that will serve as testbeds for Jitney.

Projects like OETC/SCI-Link and Jitney are leading the progress of computer architecture toward widespread use of parallel fiber optics for processor-to-processor and processor-to-I/O interconnections.

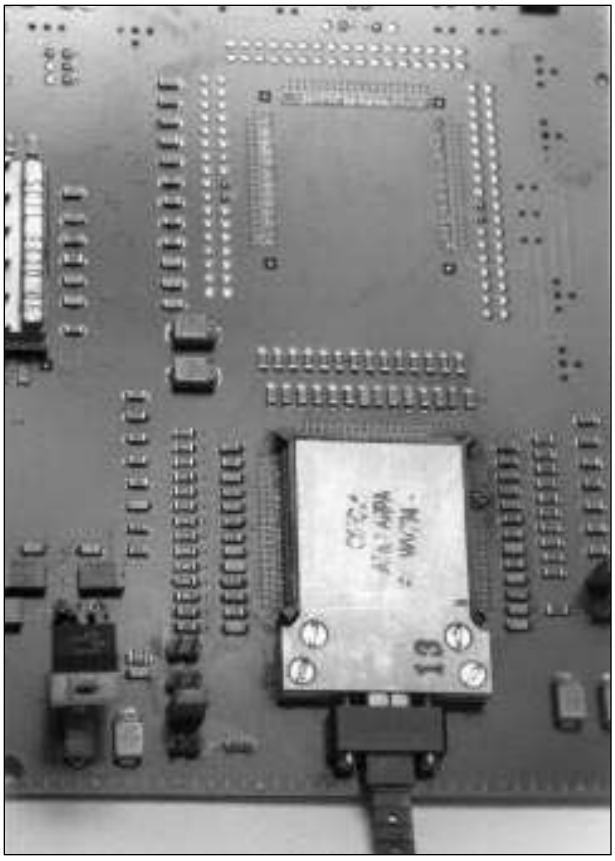

Figure 5. OETC receiver module mounted on testbed card.

Acknowledgments

ARPA (under contract MDA972-92-C-0075) and IBM Corporation supported this work. IBM employed David Engebretsen at the time he performed this work.

Ken Valk designed the lab software used to run the testbed. Del Cecchi and Curt Preuss offered valuable advice on the testbed card design.

References

1. *IEEE Std. 1596-1992, Scalable Coherent Interface (SCI),* IEEE, Piscataway, N.J., 1992.
2. B.L. Beukema et al., "Input/Output Bus Using Fiber Optics," *IBM AS/400 Technology J., Version 2,* order no. S325-6020-00, IBM Corp., Armonk, N.Y., Jan. 1992, pp. 172-177.
3. Y. Wong et al., "Technology Development of a High Density 32-Channel 16 Gb/s Optical Data Link for Optical Interconnection Applications for the Optoelectronic Technology Consortium (OETC)," *IEEE J. Lightwave Technology,* Vol. 13, No. 6., June 1995, pp. 995-1016.
4. R.A. Morgan et al., "Transverse Mode Control of Vertical Cavity Top-Surface Emitting Lasers," *IEEE Photonics Tech. Letter,* Vol. 5, No. 4, Apr. 1993, pp. 374-377.
5. D.M. Kuchta et al., "Multiple Transverse Mode VCSELs for High-Speed Data Communications," *Proc. IEEE Lasers and Electro-*

Optics Society 1993 Annual Meeting, IEEE, Piscataway, N.J., 1993, pp. 372-373.
6. D.M. Kuchta and C.J. Mahon, "Mode Selective Loss Penalties in VCSEL Optical Fiber Transmission Links," *IEEE Photonics Tech. Letter*, Vol. 6, No. 2, 1994, pp. 288-290.
7. R.J.S. Bates, D.M. Kuchta, and K.P. Jackson, "Improved Multimode Fibre Link BER Calculations Due to Modal Noise and Non-Self-Pulsating Laser Diodes," *Optical and Quantum Electronics*, Vol. 27, No. 4, Apr. 1995, pp. 203-224.
8. D.L. Rogers, "Integrated Optical Receivers Using MSM Detectors," *IEEE J. Lightwave Technology*, Vol. 9, No. 12, Dec. 1991, pp. 1635-1638.
9. D.L. Rogers, "GaAs Opto-Electronic Integrated Circuits for High-Speed Optical Communications," *Microelectronic Eng.*, Vol. 15, Oct. 1991, pp. 679-684.
10. W.H. Knausenberger et al., "Optical Interconnection Technology for AT&T Shelf and Frame Level Equipment," *Proc. 26th Annual Connector and Interconnection Technology Symp.*, Int'l Institute of Connector and Interconnection Technology, Deerfield, Ill.,1993, pp. 81-100.
11. D. Cecchi, M. Dina, and C. Preuss, "A 1GB/S SCI Link in 0.8µm BiCMOS," *Digest of Technical Papers, 1995 IEEE Int'l Solid-State Circuits Conf.*, IEEE, Piscataway, N.J., 1995, pp. 326-327.

David R. Engebretsen is a PhD student at the University of Minnesota, where he researches specialized processor architectures and optimizing compilers designed for those processors. He is studying under an IBM graduate fellowship. Engebretsen received his BSEE and MSEE degrees from the University of Minnesota. He is a member of the IEEE Computer Society and the ACM.

Daniel M. Kuchta is a research staff member at the IBM T.J. Watson Research Center. His current research interests include high-speed modulation of VCSELs and low-cost components for fiber-optic data communications. Kuchta received the BS, MS, and PhD in electrical engineering and computer science from the University of California, Berkeley.

Richard C. Booth is a senior engineer at the IBM AS/400 Division. He has held numerous technical leadership and management positions on System/34 and System/36 processor designs and massively parallel processor designs. He has authored technical papers on processor design, multiprocessor architecture, scalable symmetrical multiprocessors, and high-speed link technology. Booth received his BSEE degree from the University of Minnesota.

John D. Crow is a member of IBM's research staff at the T.J. Watson Research Center, where he is currently involved in multiprocessor optical interconnection design and prototyping as manager of the OETC and Jitney projects. He coedited and contributed to *Integrated Optoelectronics* (Academic Press, 1995). Crow received his BSEE, MSEE, and PhD from the University of California, Berkeley. He is a fellow of the IEEE.

Wayne G. Nation is an advisory engineer at the IBM AS/400 Division. He has coauthored over 25 papers related to his research interests in interconnection networks, parallel algorithms, scalable parallel-processing systems, and high-speed interconnectivity. Nation received the BSEE from the University of Kentucky and the MSEE and PhD from Purdue University. He is a member of the IEEE Computer Society.

Send correspondence about this article to Richard C. Booth, IBM Corp., 3605 US 52 North, Rochester, MN 55901; rcbooth@rchland.vnet.ibm.com.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Service Card.

Low 156 Medium 157 High 158