Gigabyte-per-Second Optical Interconnection Modules for Data Communications

Kenneth H. Hahn

Kirk S. Giboney

Robert E. Wilson

Joseph Straznicky

A ten-channel parallel optical link module operating at 1 Gbit/s per channel has been developed in the POLO (Parallel Optical Link Organization) program. Key components include vertical-cavity surface emitting laser and detector arrays, bipolar transceiver ICs, a high-speed ball-grid array package, polymer waveguides, and multichannel ribbon fiber connectors. Applications of the POLO module include computer clusters, switching systems, and multimedia.

The performance of advanced computer and communications systems is increasingly limited by the constraints of electrical interconnects. Future demand for interconnect bandwidth in computing and switching systems will be met by optical interconnects. Evolving communications standards such as ATM, gigabit Ethernet, and Fibre Channel require serial data rates approaching and often exceeding 1 Gbit/s. The next generation of high-performance processors will have clock speeds in excess of 300 to 400 MHz and aggregate processor bus bandwidths of more than 2 to 3 Gbytes/s. The increasing performance of such systems has driven the development of Gbyte/s interconnection standards such as SCI and Super HIPPI.

Such demands, when combined with stringent low-cost and high-performance specifications, cannot be met by any existing commercially available interconnect technology. Given the constraints of standard optical and electrical interconnections, parallel optical interconnection solutions operating at Gbyte/s data rates offer a number of advantages. The input and output data is inherently in parallel format, which reduces latency of multiplexing and demultiplexing functions and simplifies system integration. Parallel optical links amortize packaging costs over multiple channels while minimizing link latency and module footprint. By comparison, serial links will be expensive and bulky in multiple-channel applications. Copper wire has a very limited bandwidth-length product and is unsuitable for Gbyte/s data communications in the local area (10 to 300 m). In the telephone central office environment, for example, electrical interconnects between high-capacity switching systems are creating serious bottlenecks in terms of the sheer bulk of the cable required, the limited backplane real estate available for connections, and the resultant EMI created by large electrical cable bundles.[1]

The key performance parameters in such interconnections are bandwidth-length and interconnection density. Flexible film cables with impedance-controlled transmission lines offer high bandwidth and density. However, effective operating lengths are limited by attenuation and noise to less than 1 to 2 m. Twisted-pair and microcoaxial cables can accommodate transmission lengths of approximately 10 to 20 meters. However, they are bulky and relatively expensive. Connectors also limit performance. Today, an 18-twisted-pair cable for Gbyte/s Scalable Coherent Interface (SCI) measures 4 inches by 1.25 inches. In comparison, a 10-to-12-fiber optical connector will be less than 0.4 inch by 0.3 inch, representing an order of magnitude reduction in cross-sectional form factor. Optical fibers in ribbon form have much higher density, lower attenuation and skew, competitive cable cost, and future scalability to multi-Gbit/s line rates.

Parallel optical data links are expected to be widely used as interconnections for computer clusters, switching systems, and multimedia (**Figure 1**). Cost reduction and demonstration of reliable and robust operation will be critical to the success of parallel optical links in markets presently dominated by copper-wire and serial optical data links.

Figure 1

Applications of parallel optical data links.

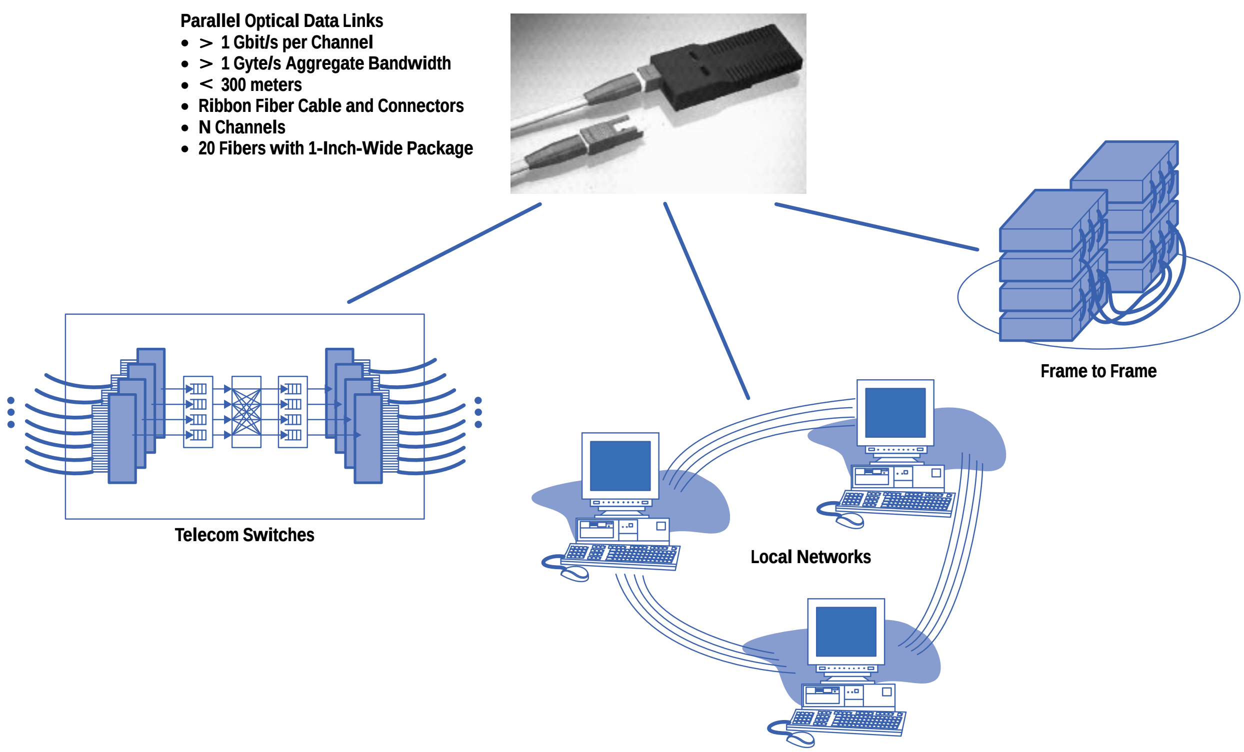

POLO Program

To demonstrate the technical feasibility of low-cost, high-performance parallel optical data links, a three-year collaborative program was launched by HP, AMP, Du Pont, SDL, and the University of Southern California (USC) in August 1994.[2] Led by HP, the POLO (Parallel Optical Link Organization) program has demonstrated 20-Gbit/s throughput in a one-inch-wide footprint, developed a high-density optical connector for ribbon fiber, and demonstrated the operation of POLO modules in a workstation communication testbed at USC. Two generations of modules have been developed, as shown in **Figure 2**. The first generation (POLO-1) module operated at 622 Mbits/s per channel with 980-nm vertical-cavity surface emitting lasers (VCSELs) and featured a prototype connector assembly and a 1.75-inch-wide leadframe package. The second genreation (POLO-2) module operates at 1 Gbit/s per channel and incorporates a higher-density (1-inch-wide) ball-grid array (BGA) package, a multichannel ribbon fiber connector from AMP, and 850-nm VCSELs and MSM (metal-semiconductor-metal) detector arrays.

Figure 2

(a) POLO-1 module on evaluation board. (b) POLO-2 module with push/pull connectors for ribbon fiber interface.

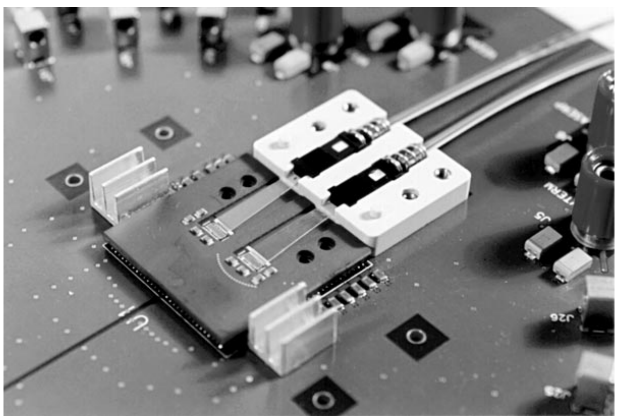

(a)

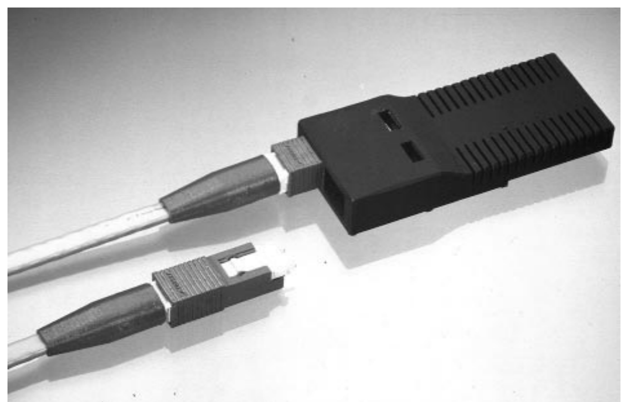

(b)

Module Design and Performance Summary

The performance of the POLO-2 module is summarized in **Table I**. The maximum length is limited by the modal bandwidth of standard multimode fiber. While interchannel skew in ribbon fiber can limit length for synchronous operation of parallel channels, recent results indicate that synchronous parallel transmission for more than 1 km is possible at 622 Mbits/s per channel by cutting each fiber in sequence from the same fiber pull, limiting interchannel skew to less than 1 ps/m.[3]

Figure 3 shows a rendering of the POLO-2 module. The key components integrated into the package have been extensively described in the literature, including vertical-cavity surface emitting lasers (VCSELs)[4] and Polyguide™ polymer optical waveguides.[5]

Table I

POLO-2 Module Performance Summary

Number of channels	10 transmit and 10 receive (9 data plus 1 clock or 10 data)
Data rate per channel	0 to 1 Gbit/s
Length	< 300 m
Electrical interface	Differential ECL, latched or unlatched by clock channel
Multichannel module package	Ceramic ball-grid array
Module width	2.5 cm
Wavelength	850 nm
Connector	Lightray MPX (based on MT ferrule)
Optical interface	62.5/125-μm graded-index ribbon fiber to polymer waveguide
Power dissipation	< 2W or < 100 mW/channel
Receiver sensitivity	−17 dBm (−20 dBm at detector), single channel only

Figure 3

POLO-2 module design.

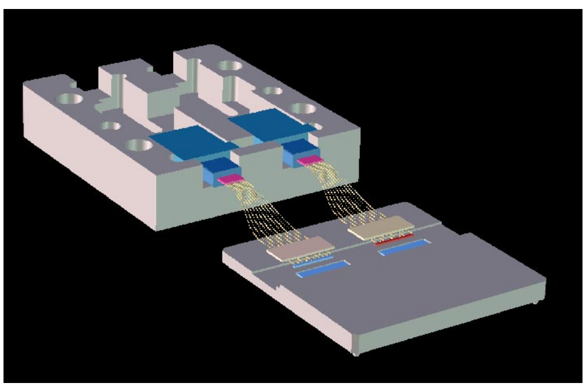

The design of the optical-electrical interface is shown in **Figure 4**. The VCSEL and detector arrays are packaged with the transceiver ICs in a 324-pin ceramic BGA package. Polyguide waveguides couple light between the VCSEL and detector arrays and the ribbon fiber using 45° out-of-plane mirrors and fiber-to-waveguide connectors. The ceramic package features impedance-controlled traces and integrated resistors for termination of input ECL signals. The use of a 45° optical

Figure 4

(a) Optical-electrical interface design. (b) Coupling of VCSEL and detector arrays with optical waveguides.

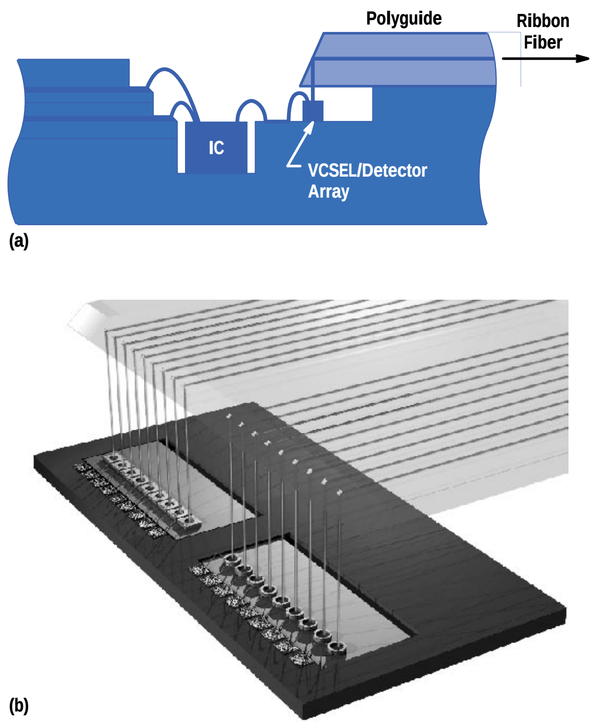

interface allows the VCSELs and detectors to be packaged in close proximity to the transceiver ICs, allowing control of electrical parasitics and GHz-bandwidth operation. Because the waveguides are multimode, simultaneous alignment of ten channels to the VCSEL and detector arrays is possible with loose alignment tolerances.

VCSELs and MSM detectors

VCSELs are ideal sources for optical data links. The devices are processed and characterized at the wafer level, and one-dimensional or two-dimensional arrays are easily fabricated. Light is emitted perpendicular to the substrate with a circular beam that enables efficient, direct, fiber or waveguide coupling. For parallel links, VCSEL arrays can be fabricated to match the pitch of the optical waveguide array. Large-area top emitting 850-nm VCSELs are used in the POLO-2 module. The threshold currents of these 18-μm-diameter VCSELs are about 3 to 4 mA. The lasers are typically prebiased near threshold to guarantee a high extinction ratio for all channels, and modulated to peak output power of $\sim$2 mW. The low relative intensity noise (RIN) and reflection sensitivity of the VCSELs allows Gbit/s data rates in multimode fiber links with low BER. RIN is typically less than -130 dB/Hz under typical operating conditions. We have previously shown that large-area VCSELs emit in multiple transverse modes, leading to reduced coherence.[6] This reduces the susceptibility of the multimode fiber link to modal noise, making these sources ideal for such applications. **Figure 5** shows an eye diagram of an 850-nm VCSEL biased below threshold and driven with a pseudorandom binary sequence (PRBS) at 1 Gbit/s. The eye is wide open and the measured BER is $<10^{-13}$.

An attractive feature of VCSELs is their ability to scale to higher data rates. Modulation of greater than 3 Gbits/s per channel has been successfully demonstrated. **Figure 6** shows the frequency response of a 980-nm VCSEL at two bias currents, showing a small-signal -3-dB electrical frequency response of 6.6 GHz at the larger bias.

Figure 5

Eye diagram of an 850-nm VCSEL biased below threshold and driven with a pseudorandom binary sequence at 1 Gbit/s.

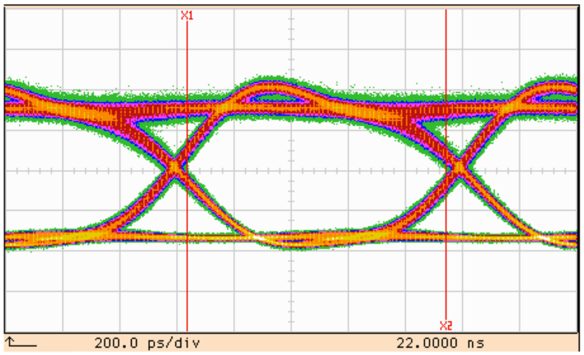

Figure 6

Frequency response of 980-nm VCSEL at two bias currents.

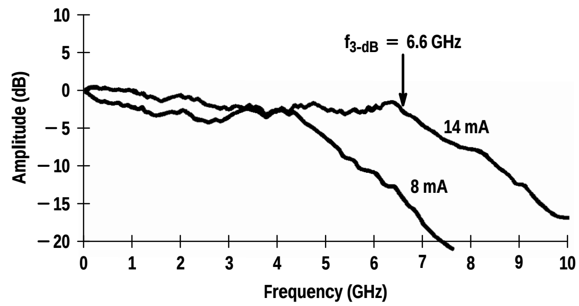

A 1-by-10 MSM detector array is shown in **Figure 7**. These devices are fabricated with a straightforward two-mask process. The interleaving metal fingers on a bulk GaAs layer have very low capacitance, allowing large area detectors with a measured -3-dB frequency of 1.5 GHz. The detector area of 200×200 μm^2 provides greater than ± 50-μm alignment tolerance to the optical waveguides. The measured responsivity at 850 nm is >0.4A/W and fall times are <200 ps.

Figure 7

1-by-10 MSM detector array.

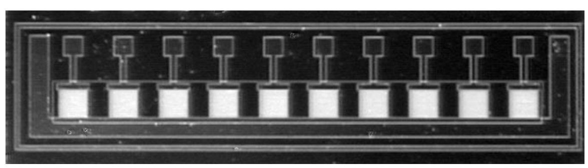

Transmitter and Receiver ICs

Transmitter and receiver ICs fabricated with Hewlett-Packard's HP-25 silicon bipolar process are used in the POLO module. The transmitter IC contains ten laser drivers that use common reference voltages to set the VCSEL prebias and modulation currents. The transmitter input and receiver output interfaces are differential ECL.

Since the receiver determines the link architecture, several versions of the receiver IC have been designed to provide maximum user flexibility, including arrays of latched digital receivers, unlatched digital receivers, and analog transimpedance amplifiers for linear testing. The latched receiver has nine data channels and one clock channel. The output data is synchronized by the clock channel at the receiver output, removing any accumulated skew and jitter. The unlatched receiver allows the module to operate as ten independent serial links.

Both ac-coupled and dc-coupled versions of the receivers have been fabricated and tested. Because each channel determines its own threshold, the ac-coupled system has much higher channel-to-channel dynamic range. However, data needs to be encoded because of a low-frequency cutoff. The dc-coupled version can handle any data pattern, but channel-to-channel uniformity in received power (within several dB) is required because a single threshold is used across all channels.

Polyguide™-Ribbon Fiber Optical Interface

The use of polymer waveguides allows the waveguide design to be easily tailored to system requirements, including waveguide dimensions, pitch, and numerical aperture. For example, the waveguide pitch is 360 μm at the p-i-n detector interface and 500 μm at the VCSEL interface, but a smooth taper allows a waveguide pitch of 250 μm at the ribbon fiber interface. The width and numerical aperture of the polymer waveguide are optimized to increase coupling efficiencies and optical alignment tolerances at each interface.

The Polyguide™ waveguides are assembled with an MT-style ferrule and aligned to the VCSEL and MSM detector arrays on the ceramic package. **Figure 8** shows a 10-channel polymer waveguide integrated with an MT-style ferrule. Guide pins in the MT ferrule allow for accurate optical alignment of this assembly with ribbon fiber.

Figure 8

Polymer waveguide assembled with an MT-style ferrule for the multichannel optical interface.

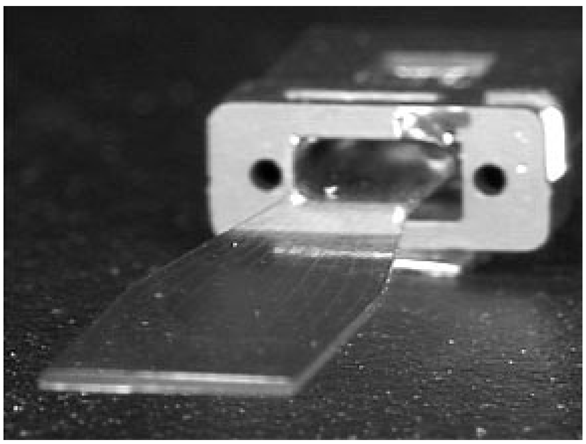

Assembly with BGA package

The POLO-2 module is the first fiber-optic module based on a ball-grid array (BGA) electrical interface. A significant advantage is the high pin density of the BGA. For example, the use of a BGA enabled a 3× reduction in package size compared to the leadframe package of POLO-1. Other advantages of BGA technology include compatibility with standard surface mount processes, high thermal conductivity, and low electrical parasitics.

The VCSELs and p-i-n detectors, laser driver and receiver ICs, and bypass capacitor arrays are mounted on the BGA substrate and wire-bonded (**Figure 9**). Polyguide™ waveguides are aligned to the optoelectronics and attached to the package, forming the optical interface to the ribbon fiber. The 18-by-18 BGA is on standard 0.050-inch pitch, resulting in a total module width of 1 inch. Integrated 50Ω resistors in the ceramic package allow termination of the input ECL signals.

Push-Pull Connector for Ribbon Fiber Interface

POLO-2 features a multichannel ribbon fiber connector developed by AMP. Only 9 mm wide and capable of handling 12 fibers, this connector is based on the precision molded MT array ferrule housed inside a push-pull SC-style housing. The ribbon fiber cable uses 62.5/125-μm fiber and meets the requirements of GR-001435 *Generic Requirements for Multifiber Optical Connectors for Type IR Media (Ribbonized Fiber Enclosed in Reinforced Jacket)*. The design and construction of the push-pull connector is also in accordance with the optical, environmental, and mechanical testing requirements of the same Bellcore generic requirement specifications.Thus, the uniformity of the insertion loss across 10 channels will be kept below 0.6 dB throughout the service life, which includes 200 durability mating cycles, and the optical insertion loss

Figure 9

(a) 324-pin BGA package for optoelectronic integration. (b) Wire-bonded IC and detector arrays.

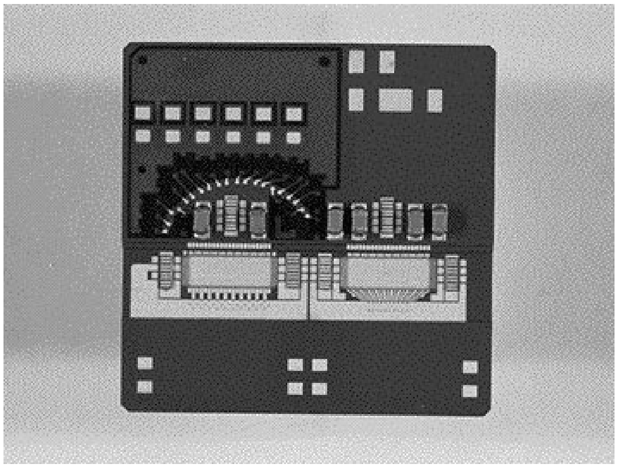

(a)

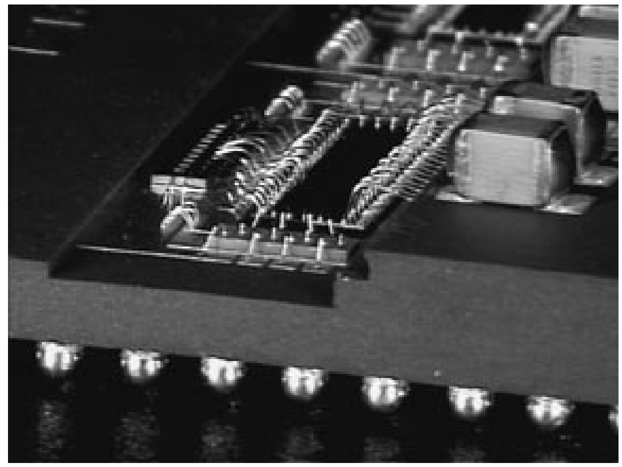

(b)

for the interface will be less than 2 dB at the end of the service life. **Figure 2b** shows the assembled POLO-2 module with ribbon fiber connectors and a plastic housing that provides a receptacle for the connectors.

System Results

The module is mounted on an evaluation board for characterization. To prevent the transfer of any mechanical loads from the ribbon fiber cable to the internal module components, the module housing mounts rigidly to the printed circuit board. Supply voltages of -5V and -3V are required for transmitter and receiver operation. An additional -2V supply is also required for ECL termination.

To test BER with worst-case cross talk conditions, all of the transmit and receive channels of one module are operated in loopback mode, with the transmitter and receiver of one module connected by a single ribbon fiber. A multichannel data generator (HP 80000) is used to modulate the ten transmitter channels with independent PRBS streams. **Figure 10** shows the eye patterns of all 10 channels in simultaneous operation at 1 Gbit/s at the receiver output.

The BER for each channel is $<10^{-11}$, and an extended measurement of one channel results in BER $<10^{-14}$ with 300 m of low-skew ribbon fiber. While some pattern dependent jitter is observed, the eyes are clearly open at 1 Gbit/s. The rise and fall times are <350 ps, and channel-to-channel skew (excluding ribbon fiber skew) is <100 ps. The eye opening (timing margin for BER $<10^{-9}$) is typically $>70\%$ for most channels. **Figure 10b** shows ten simultaneous output eye patterns of the module on a single oscilloscope trace. The aggregate timing margin for all channels is better than 50%.

We have also obtained operation of the POLO-2 module at 1 Gbit/s per channel with an ac-coupled receiver. Initial measurements show significantly improved channel-to-channel dynamic range. Latched receivers have been operated previously with the POLO-1 module at 622 Mbits/s. Performance comparable with unlatched systems has been demonstrated. A 622-MHz clock channel synchronized nine data channels, eliminating accumulated skew and jitter at the receiver output.

Figure 10

Output eye patterns of unlatched, dc-coupled module at 1 Gbit/s per channel. (a) 10 channels shown separately. (b) 10 channels aggregated.

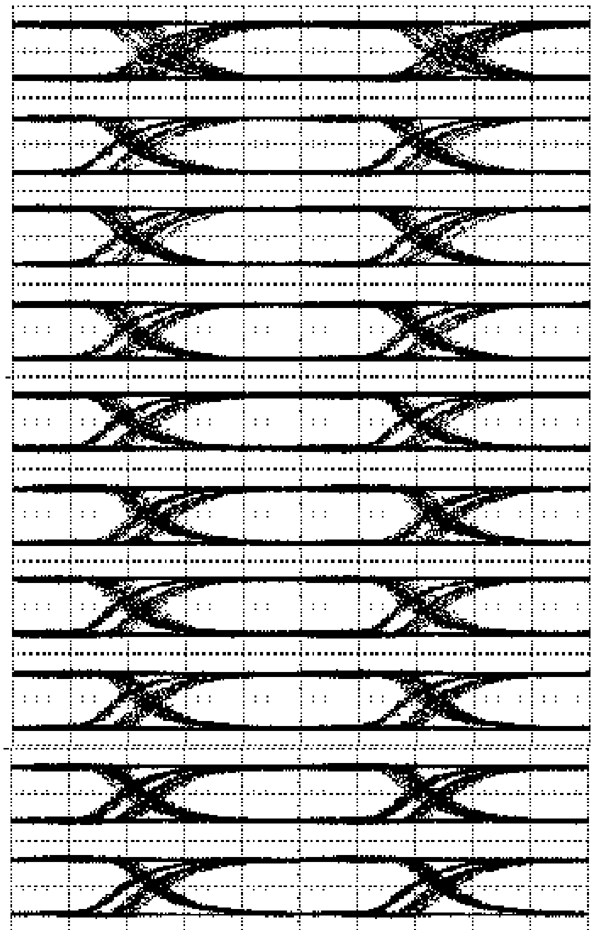

(a)

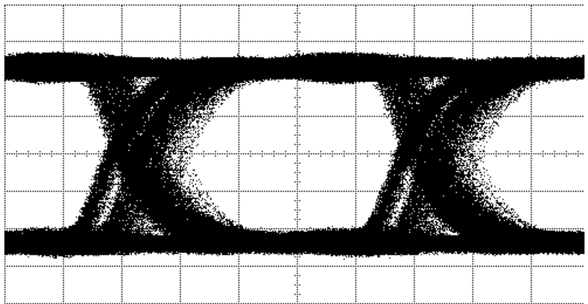

(b)

Figure 11

Link adapter board for interfacing a POLO-2 module with a computer bus.

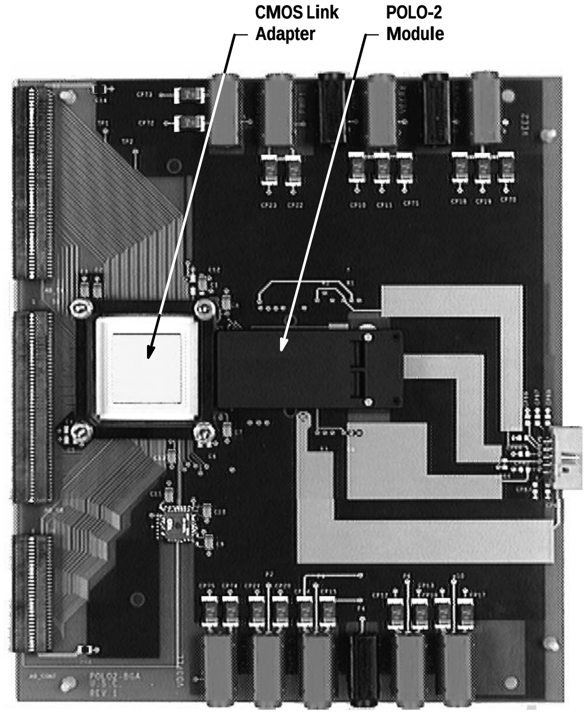

Network Interface Testbed

A prototype POLO module with evaluation board has been successfully integrated into a Gbit/s experimental workstation network at the University of Southern California (USC). The network uses experimental high-speed network interface boards called Jetstream, which were developed at Hewlett-Packard Laboratories Bristol.[7] One of these boards is inserted into each of two Hewlett-Packard 9000 Series 700 workstations, which form the two nodes of the network. Eight channels (4 transmit, 4 receive) of the POLO module, each running at 1 Gbit/s, were exercised between the two workstations, which were connected by 500 m of low-skew fiber ribbon. The POLO module successfully transmitted and received multi-Gbit/s data packets error-free in this network.

In addition, a link adapter board for a host interface has been fabricated at USC (**Figure 11**). This board contains a CMOS link adapter chip, which will directly interface to the POLO module and to external synchronous FIFO buffers. This will allow the use of the hardware interface with generic bus architectures such as PCI or other open bus standards. USC has recently demonstrated 1-GHz clocking and 1-Gbyte/s throughput in the link adapter chip.

Future Developments

The bandwidth demands of processing and communications systems will continue to multiply in the foreseeable future. While ribbon fiber has the highest bandwidth, density, and length capability, its implementation at very short distances (less than 1 to 2 meters) is limited by termination costs (i.e., parallel optical modules and connectors). To demonstrate the cost/performance superiority of parallel optics for short-distance applications such as shelf-to-shelf and board-to-board interconnections, greater system integration and functionality need to be demonstrated to the end user. The functionality includes interface compatibility with common communications standards and direct integration with network and processor buses. The integration and packaging include small footprint, low power consumption, and cheap optical subassemblies packaged as standard electronic components.

Conclusion

Parallel optical links that offer the highest bandwidth-length and bandwidth-density performance available have been designed and demonstrated. 1 Gbyte/s duplex operation over several hundred meters of ribbon fiber has been obtained with a 1-inch-wide optical interconnection module with ten transmit and ten receive channels.

Acknowledgments

The authors would like to thank Eric Wong and Tsu-Yau Chuang for assembly and testing, and Mike Tan and Chris Kocot for VCSEL and photodetector development. Thanks also to Jean Norman and Susan Davies for packaging, Tim Schoen for printed circuit board design and layout, Larry Hilliard and Lewis Dove at the Colorado Springs Technology Center for module assembly, and Alan Peters for board assembly. We would also like to thank Jim Dudley, Chun Lei, and Nobu Nakamoto at the Optical Communications Division for their work on VCSELs and photodetectors, Rick Schneider for p-substrate and oxide-confined VCSEL development, Brian Lemoff for Polyguide™ support, and Dennis Brzezinski for helpful discussions on system applications. We are grateful to David Dolfi, Steve Newton, Waguih Ishak, Ron Kaneshiro, and Steve Joiner for their support of this work.

References

1. G.J. Grimes, S.R. Peck, and B.H. Lee, “User perspectives on intrasystem optical interconnection in SONET/SDH transmission terminals,” *1992 IEEE Global Telecommunications Conference*, IEEE, New York, 1992, pp. 201-207.

2. K.H. Hahn, K.S. Giboney, R.E. Wilson, J. Straznicky, E.G. Wong, M.R. Tan, R.T. Kaneshiro, D.W. Dolfi, E.H. Mueller, A.E. Plotts, D.D. Murray, J.E. Marchegiano, B.L. Booth, B.J. Sano, B. Madhaven, B. Raghavan, and A.F.J. Levi, “Gigabyte/s Data Communications with the POLO Parallel Optical Link,” *Proceedings of the 46th Electronic Components and Technology Conference*, 1996, pp. 301-307.

3. S. Siala, A.P. Kanjamala, R.N. Nottenburg, and A.F.J. Levi, “Low-skew multimode ribbon fibers for parallel optical communication,” *Electronics Letters*, Vol. 30, October 1994, pp.1784-1786.

4. M.R. Tan, K.H. Hahn, Y.M. Houng, and S.Y. Wang, "SELs for short-distance optical links using multimode fibers," *Conference on Lasers and Electro-Optics 1995*, Optical Society of America, Washington D.C., 1995, pp. 54-55.

5. B.L. Booth, "Polymers for integrated optical waveguides," *Polymers for Lightwave and Integrated Optics*, C.P. Wong, editor, Academic Press, New York, 1993. Also B.L. Booth, "Optical interconnection polymers," *Polymers for Lightwave and Integrated Optics: Technology and Applications*, L.A. Hornak, editor, Marcel Dekker, New York, 1993.

6. K.H. Hahn, M.R. Tan, Y.M. Houng, and S.Y. Wang, "Large-area multitransverse-mode VCSELs for modal noise reduction in multimode fiber systems," *Electronics Letters*, Vol. 29, August 1993, pp. 1482-1483.

7. A. Edwards, et al, "User-space protocols deliver high performance to applications on a low-cost Gb/s LAN," *ACM SIGCOMM*, 1994.

Kenneth H. Hahn

A principal project scientist with HP Laboratories, Ken Hahn is currently on leave, attending Harvard Business School. He has been a project leader and a principal investigator of the POLO Consortium and has worked on high-speed modulators, polymer optical waveguides, and gigabit/s optical links. His PhD degree in physics is from Stanford University (1991). He is married and enjoys travel, golf, and tennis.

Kirk S. Giboney

Kirk Giboney is a hardware design engineer with HP Laboratories, responsible for parallel optical link module design and testing. He received his PhD degree in electrical engineering from the University of California at Santa Barbara in 1995, and has forty publications to his credit. Before coming to HP he was with McDonnell Douglas Astronautics Company and Hughes Research Laboratories. His interests include cycling, flying, hiking, archery, and the arts.

Joseph Straznicky

A project scientist at HP Laboratories since 1980, Joe Straznicky is responsible for the development of high-performance optical receivers, transmitters, and subsystems in various technologies. He received his Diplom Ingenieur in cybernetics in 1969 from the Technical University of Brno, Czech Republic. Before joining HP he was with the University of California at Los Angeles and Hughes Aircraft Company.

Robert E. Wilson

Rob Wilson received his BSME degree in 1978 from Stanford University. A development engineer with the HP Communications Semiconductor Solutions Division, he is a packaging specialist, responsible for optical subassemblies for small-form-factor modules. He is married, has three children, and is interested in photography and elementary school science teaching.

▶ Go to Next Article

▶ Go to Journal Home Page