

Wavelength Division Multiplexed Optical Interconnect Using Short Pulses

Bianca E. Nelson, *Student Member, IEEE*, Gordon A. Keeler, *Student Member, IEEE*, Diwakar Agarwal, *Student Member, IEEE*, Noah C. Helman, *Student Member, IEEE*, and David A. B. Miller, *Fellow, IEEE*

Abstract—We demonstrate operation of a wavelength division multiplexed chip-to-chip optical interconnect using surface-normal electroabsorption modulators, and a modelocked laser as a single broadband source. The link was successfully operated at 80 Mb/s. While this rate was limited by the repetition rate of the modelocked source, individual CMOS circuits and optoelectronic devices have been shown to work at data rates approaching 1 Gb/s.

Index Terms—CMOS integrated circuits, optical interconnection, short optical pulses, wavelength division multiplexing.

I. INTRODUCTION

OPTICAL interconnects for silicon electronics have been shown to have many advantages over traditional electrical interconnects, particularly in dense, high-capacity systems [1]. One approach to optical interconnects commonly used in telecommunications today is wavelength-division multiplexing (WDM). As bandwidth needs increase, WDM becomes an attractive solution for shorter distance interconnects as well, such as chip-to-chip or board-to-board interconnects. While telecommunications WDM systems traditionally use a separate laser for each wavelength channel, which necessitates wavelength monitoring and control, a simpler, less expensive alternative is needed for interconnects on scales of a few meters or less. One such solution is to use a single, broad-band laser source. A wavelength disperser can be used to separate the wavelengths in space, and a linear array of on-chip reflection modulators can then provide the wavelength channel definitions by spectral slicing [2]. In this way, the channel spacings are automatically maintained. This concept is illustrated in Fig. 1.

In this paper, we examine a system designed for 10-channel chip-to-chip WDM optical interconnection using a modelocked Ti:Sapphire laser and surface-normal electroabsorption modulators. Wavelength division multiplexing has the potential to simplify interconnect packaging in such a system in that all the channels can be combined and transported over a single fiber, while the use of a single optical source potentially decreases cost in a short-distance interconnect. As a WDM system, our




Fig. 1. A conceptual schematic of wavelength division multiplexing by spectrally slicing the broadband output of a modelocked source.

experiment takes advantage of the broad-band nature of a modelocked source; however, the corresponding short-pulse nature can be simultaneously exploited for other advantages including improved receiver sensitivity [3], system re-synchronization [4], and latency reduction [5]. Previous work [2], [6] has demonstrated the principle of using short pulses for wavelength division multiplexing. Here we demonstrate the first known operation of a full chip-to-chip optical interconnect of this type, including GaAs devices hybridly integrated to functional silicon logic, and analyze the losses in such a system.

II. SYSTEM COMPONENTS

A. CMOS Circuits

The 4-mm² chips used in this system were fabricated through Mosis in 0.5- μm silicon CMOS and contain both the transmitter and receiver circuits. The circuits are placed to line up with an overlaid 10 \times 20 array of optoelectronic devices, which is further divided into ten linear sub-arrays (or “rows”). Each row contains ten differential channels, or 20 devices. The rows are separated by 125 μm , and the devices within them spaced at 62.5 μm . These spacings were chosen to allow coupling to fiber ribbons in future experiments, and do not represent the maximum device or circuit density. For example, the transimpedance receivers on this chip have a footprint of approximately 15 \times 17 μm , much smaller than the overlying 40 \times 80- μm diode. By reducing the size of the diodes and the diode spacing, the device density could likely be increased to thousands of I/Os per square millimeter.

One row contains the main subarray of ten differential transmitter channels. The transmitters have various possible drivers,

Manuscript received October 31, 2002; revised January 28, 2003. This work was supported by the Joint Services Electronics Program under ONR Grant N00014-91-J-1050, by the Defense Advanced Research Projects Agency (DARPA) under Ultra-Photonics Program Grant F49620-97-1-0517, by DARPA Grant MDA972-98-1-0002, by MARCO/DARPA Grant GIT-B-12-D00-S57, and a subaward from the University of New Mexico. The work of N. C. Helman was supported by a Gerhard Casper Stanford Graduate Fellowship.

The authors are with Ginzton Laboratory, Stanford University, Stanford, CA 94305 USA (e-mail: bkeeler@stanfordalumni.org).

Digital Object Identifier 10.1109/JSTQE.2003.813322




Fig. 2. Optical eye diagrams of (a) the output of a modulator driven with pseudorandom electrical data at 950 MHz and (b) the output of a transimpedance receiver's read-out modulator when the receiver is driven with optical data at 700 MHz.

including pseudorandom bit sequence (PRBS) generators, though for the experiments here, we use primarily directly driven modulator diodes for detailed testing of individual channels. These diode pairs can also be used to measure the modulator reflectivity and contrast ratio as a function of voltage and wavelength.

The differential receivers occupy the remainder of the chip. Two different designs were used: clocked sense-amplifier (or “integrating”) receivers [7], and transimpedance-amplifier receivers [8]. Only the integrating receivers are used here because of their greater sensitivity. Adjacent to each row of receiver channels on the chip is a corresponding row of output (or “read-out”) modulators that allows optical monitoring of receivers without the need for further high-speed electrical connections. That is, each output modulator is driven by the received data from the corresponding receiver. Additionally, one specific receiver channel is directly connected to a high-speed SMA connector on the circuit board, allowing electrical monitoring of the receiver state. This is particularly useful when optical read-out is difficult.

The finished chips with integrated devices were wire-bonded into a high-speed electrical package, which was subsequently soldered to a custom-made printed circuit board (PCB). The circuit functionality was tested before the chips were implemented in the interconnect system. The individual receiver and transmitter circuits, together with their integrated optoelectronic diodes, were shown to have open eye diagrams at high speeds (Fig. 2). These results are limited by the bandwidth of the high-speed detector used—approximately 700 MHz.

B. Optoelectronic Devices

The present device and hybridization technologies are similar to those described in [9]. The optoelectronic devices are multiple-quantum-well (MQW) GaAs *p-i-n* diodes, similar to those first described in [10]. The MBE-grown wafer used to fabricate the devices contains 50 GaAs quantum wells in the intrinsic region with $\text{Al}_{0.25}\text{Ga}_{0.75}\text{As}$ barriers, between $\text{Al}_{0.25}\text{Ga}_{0.75}\text{As}$ *p* and *n* layers. A series of photolithographic steps was performed on the GaAs wafer to define the mesa structure of the diodes, which is shown in Fig. 3. Coplanar gold contacts are deposited to provide electrical connections to the *p* and *n* layers. When a static reverse bias is applied across the diodes, they act simply as photodetectors for wavelengths of about 860 nm and shorter. With changing applied bias, however, the diodes can also serve as electroabsorption modulators whose operation is based on the quantum-confined Stark effect [11]. That is, by changing the applied bias—and, therefore, shifting the absorption edge—a light




Fig. 3. Diagram of GaAs diodes before flip chip bonding. Gold *p*- and *n*-contacts are shown, as well as the indium solder bumps.




Fig. 4. Plot of measured contrast ratio vs. wavelength for modulator operation for three different voltage swings. For comparison, the gray solid curve represents the spectrum of the modelocked laser (in arbitrary units of intensity).

beam incident on the structure can be alternately absorbed or reflected. In use, the GaAs substrate is removed, as described below. Light is then incident on the *n*-layer, while the metallic *p*-contact serves as a mirror. The device, therefore, behaves as a reflection modulator. The devices are operated with a 3.3-V swing, corresponding to current CMOS logic levels. With such low voltages, however, the modulators have a relatively low contrast ratio (maximum 2:1) over a range of approximately 5 nm, centered near 850 nm, as seen in Fig. 4. Contrast ratio (CR) is defined as the ratio of reflected optical power in the “on” (reflecting) and “off” (absorbing) states of the modulator. The contrast ratio peak can be shifted slightly in wavelength by changing the offset voltage. To improve performance, the devices are operated in differential pairs.

C. Hybrid Integration

In order to integrate the devices to silicon electronic circuits via flip-chip bonding, indium solder bumps are first evaporated onto the diode contacts. The GaAs wafer is then diced into individual arrays of 200 devices in preparation for integration. Using a commercial visible-alignment flip-chip bonder, the indium bumps are aligned to corresponding gold contacts that have been deposited on the CMOS chip. Pressure is applied to bond the indium to the gold. Because the indium is quite soft at room temperature, epoxy is flowed between the chips to provide




Fig. 5. SEM images of GaAs quantum-well diodes flip-chip bonded to silicon CMOS circuits.

mechanical strength. The GaAs substrate is then removed by selective wet etching, leaving isolated GaAs devices bonded to the CMOS chip. Scanning electron microscope (SEM) pictures of successfully bonded devices are shown in Fig. 5. To determine the device yield after bonding, a forward bias is applied across the entire device array; working devices emit light in the infrared. Yields have been measured as high as 98%. While currently the devices are bonded at the chip scale, wafer-scale bonding could also be achieved in a similar manner.

D. Optical Design

1) *Laser*: The laser used for the system is a Spectra-Physics modelocked Ti:Sapphire laser, tunable over a range of approximately 710–980 nm. The laser was tuned to a central wavelength of 851 nm, corresponding to the contrast ratio peak of the absorption modulators (as shown in Fig. 4). Low-jitter pulses are output from the laser at a repetition rate of approximately 80 MHz and typical pulsewidth of approximately 150 fs, corresponding to a transform-limited spectral bandwidth of roughly 5 nm full-width at half-maximum (FWHM).

2) *Baseplates*: The optical components were assembled on a system of stainless steel slotted baseplates, first described in [12]. Optics such as lenses and waveplates are mounted and centered in small steel cylindrical cells, having a standard outer diameter of 25 mm. The cells rest on precision-milled slots in the plate and are held in place by magnets, which sit in the bottom of the slots and are separated from the cells by a small air gap. In this way, the mechanical degrees of freedom are reduced to rotation about and translation along the optical axis, allowing for fast and precise system alignment. The mechanical stability offered by the baseplates is also important as, for example, a rotation of the grating by only 0.01° leads to a $10\text{-}\mu\text{m}$ shift in the position of the light on the chip. To steer beams and deviate them from the optical axis, matched pairs of Risley prisms—wedged pieces of glass—are used. Beamsplitters are placed at the junction of two perpendicular slots and are glued in place after proper alignment and positioning. The baseplate used here was custom-designed to accommodate mounting of the grating positioner. A small portion of the optical setup is shown in Fig. 6.

3) *Gratings*: The dispersive elements used in this system are 158 lines/mm gold-coated reflective echelle gratings with




Fig. 6. Photograph of a baseplate with various optical elements including lenses, a beamsplitters, a waveplate, and the reflection grating. The CMOS chip on a printed circuit board can also be seen.

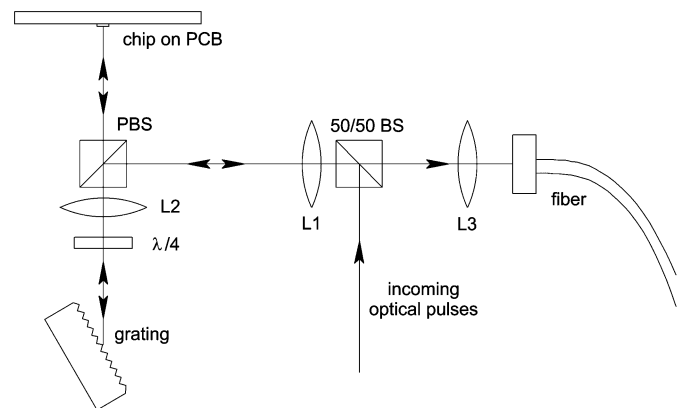


Fig. 7. Schematic showing transmitter portion of interconnect system. The receiver portion is similar.

a 63-degree blaze angle. When used in the Littrow configuration, the gratings have a reflection efficiency of approximately 60%, which varies slightly with polarization. The gratings were used in this configuration—with the input and reflected light at substantially the same angle—to facilitate incorporation into the baseplate system and to reduce alignment sensitivity. The Newport goniometer stages on which gratings were mounted allow rotation of ± 10 degrees about the central position for fine-tuning. By placing the grating at the focal distance from a lens, the reflected light's angular wavelength separation is converted to a spatial wavelength separation. A lens with a focal length of 60 mm was chosen to achieve the desired wavelength separation of 0.5 nm, or roughly 200 GHz, per channel. A change in grating position of 1 degree (easily read out from the goniometer scale) results in only a 3% change in wavelength separation at the chip, allowing for precise matching of channel spacing between the transmitter and receiver chips.

4) *Optical System Design*: The path of the signal beam is detailed in Fig. 7. (Some of the same optical components are also shown in Fig. 6.) The incoming short pulse light beam passes through a nonpolarizing beamsplitter (50/50 BS) before

passing through a polarizing beamsplitter (PBS) and a telescope system consisting of a 40-mm and 60-mm focal length lens (L1 and L2, respectively). This telescope resizes the beam to a diameter compatible with the active area of the modulators. In order to be transmitted through the polarizing beamsplitter on the return pass, the beam makes two passes through a quarter-wave plate. After reflection from the grating, the light again passes through the 60-mm focal length lens, which serves to focus the spectral strip of light onto the chip, while simultaneously parallelizing the wavelengths. The modulators modulate and reflect portions of the spectrum, which pass back to the grating. At this second grating reflection, the wavelength channels are recombined into one beam, which now reflects at the polarizing beamsplitter and is focussed into the interconnect fiber by a 20-mm focal length lens (L3). Alternatively, the beam can simply be passed to the receiver side by a series of mirrors to create a free-space interconnect.

The signal light that reaches the receiver baseplate is similarly demultiplexed by the second grating onto the receiver chip. Instead of a spectral stripe, the light now appears as discrete spots approximately the size of the modulator active area. These spots can be positioned to land on any of several receiver channels or arrays of channels.

To monitor the state of a receiver channel optically, a separate read-out beam is brought onto the receiver baseplate along a different path and focussed onto the proper read-out modulator. The reflected portion of this beam is sent to a high-speed detector via a second fiber. Note that, given the limited number of possible beam paths on the baseplate, only one receiver channel can be monitored at a time. Polarization optics are used throughout this system to minimize losses and to make the system as compact as possible.

III. EXPERIMENTAL RESULTS

A typical spectral slicing plot, showing a portion of the spectrum reflected from the transmitter chip, is shown in Fig. 8. Modulator data from two different points in the PRBS are shown. The static peaks between diodes are reflections from the silicon CMOS chip between the modulator diodes, as the surface of this chip is not antireflection coated. This additional nonsignal reflection is a possible source of noise error in the overall system performance.

To demonstrate link performance, the electrical bit patterns input to the transmitter and output from the electrical output of a receiver were compared. Because only one receiver channel had an electrical output available, only single-channel link operation could be demonstrated in this way. The available optical power prevented simultaneous multichannel tests using other methods. In this case, a single modulator pair on the transmitter chip modulated a portion of the spectral stripe. This channel was driven with an external 32-bit data sequence, frequency-locked to the 80-MHz laser repetition rate. The modulated signals were sent through free space to the second chip, and the data received at the corresponding receiver channel was read out through a high-speed electrical connection. Fig. 9 shows both the electrical input to the transmitter channel as well as the received data at the other end, indicating successful operation of the link at 80 Mb/s and hence connection between the silicon




Fig. 8. Spectral slicing plot of reflection from three channels for two different points in the pseudorandom sequence (dashed and solid lines). Each channel is optically differential, and therefore uses two modulators. For the two curves shown, all three channels have changed state.




Fig. 9. Oscilloscope plot showing electrical input to modulators (top) and electrical output of receivers (bottom).

logic levels of the two chips. The differential received optical power required was approximately $6 \mu\text{W}$ per diode. Glitches in the bottom rail of the receiver output trace are due only to electrical feed-through in the output and do not indicate errors.

To show the feasibility of a flexible link, the same system demonstration was performed using the fiber link. While successful operation was also achieved for this interconnect, data rates were limited to less than 10 Mb/s because of additional signal power losses incurred by use of the fiber.

IV. DISCUSSION

The main objective of this system was to investigate the feasibility of a WDM chip-to-chip optical interconnect, as well as to compare the issues involved in implementing a WDM link versus a standard free-space link. The primary difficulties in implementing this system were found to be optical power losses and alignment, both of which are due primarily to the use of gratings and spectral slicing.

This WDM system requires more overall optical power than an equivalent non-WDM short-pulse chip-to-chip interconnect system [3], [4], because less of the input optical power is modulated in the WDM case. There are two primary sources of such loss: reflection from gratings, and spectral slicing. The gratings used here have only about 60% efficiency, and the input light makes one reflection from a grating before reaching the transmitter chip. This reflection from the grating falls across the transmitter chip in a continuous band of light, extending past the ends of the linear array of modulators. The spectral bandwidth of the laser pulses is chosen such that only the central portion of the wavelengths—the 5-nm FWHM portion—is used. This means that 29% of the intensity, contained in the “tails” of the spectrum, falls outside the modulator array area and is wasted. However, the benefit gained is that the light intensity across the array is somewhat more uniform than if a narrower-bandwidth pulse were used. Additionally, the active modulator areas fill only 40% of the linear array, leading to additional losses from light falling between modulators. Altogether, then, only about 17% of the light incident on the first grating is modulated. These losses could be reduced by factor of three by using a multiwavelength “comb” source similar to that demonstrated in [13], tailored to the wavelength needs of this system.

The next important consideration is how much of the differential signal power from the transmitter chip is incident on the receiver chip. Signal power is here defined as the difference between the amounts of optical power reflected from a single modulator in the “on” and “off” states. The WDM-specific losses at work here are two grating reflections of about 60% each, and the coupling into the optical fiber (when used). A single-mode fiber was used for this experiment, and the maximum fiber transmission of signal power achieved was about 10%. This loss was due in part to the distortion of the beam shape upon reflection from the uneven surfaces of the modulator diodes, as well as diffraction losses for those wavelengths reflected from the edges of the integrated devices. The total chip-to-chip transmission of differential signal power was then about 36% in the free-space link case, or about 4% when using the fiber link. Using lower-loss (de)multiplexers and more efficient fiber coupling would greatly decrease this loss.

An additional issue to explore is whether a wavelength division multiplexed optical interconnect requires more received differential signal power than the equivalent non-WDM system. When the same type of transmitters and receivers used in this experiment were implemented in an equivalent non-WDM link, the required received differential signal power was 35 μW per diode at 400 Mb/s, indicating a requirement of 7 μW per diode at 80 Mb/s. This is in close agreement with the 6 μW per diode at 80 Mb/s required in the free-space WDM system, indicating essentially no power penalty in using wavelength division multiplexing. The more than ten times reduction in operating speed of the fiber-based WDM interconnect is due in large part to the 10% efficiency of the fiber link itself.

The mechanical alignment between the transmitters and receivers proved to be one of the greatest challenges in this system, particularly when the additional alignment sensitivity of the interconnect fiber was included. Slight changes to the balance between and position of the signal spots on the two detectors

within a differential pair caused great variation in system performance. This is largely due to the spectral slicing nature of the system. The existence of nonsignal light between the diodes and the fact that the signal beam spot size was the same as that of the detectors both contributed to the problem. This sensitivity could again be relieved by using a multiwavelength “comb” source rather than a continuous-spectrum source.

V. SUMMARY AND CONCLUSION

We have demonstrated operation of a wavelength division multiplexed optical interconnect using a broadband source and surface-normal electroabsorption modulators. Such use of a single optical source and single interconnect fiber can significantly reduce the system cost and complexity; this is a necessity if this type of interconnect is to be packaged for short distances. The system demonstrated here used only a continuous spectrum, which led to various problems in efficiency of the link. These results suggest that use of a modelocked laser generating a spectral comb matched to the modulator array would eliminate many of the loss and coupling efficiency problems. Furthermore, while only single-channel operation was demonstrated, the spectral width of the source is sufficient for multiple channels, and the modulators appear to have a large enough spectral bandwidth for such a system. The other practical difficulties in the present demonstration were mostly related to optical alignment, and we expect these could be eliminated in a system employing planar packaging of the optical components [14]. With the use of comb laser sources and such integrated optics, this WDM approach could be practically interesting for dense interconnects.

With the development of appropriate optoelectronic modulators and detectors for 1.55 μm , this system concept could be further extended to use in networks. That is, WDM optical network data could be brought directly down onto a CMOS chip containing all of the network electronics, providing for an extremely compact system. While the use of WDM in long-distance applications has become standard, it is clear that the full potential of wavelength division multiplexing has yet to be reached. Significant opportunities for WDM in short distance interconnects and for integrated WDM network chips exist, particularly as silicon CMOS speeds extend beyond 10 GHz.

ACKNOWLEDGMENT

The authors would like to acknowledge Dr. P. Atanackovic, Dr. T. Pinguet, M. Wistey, and Prof. J. Harris, Jr. for MBE wafer growth. N. Helman gratefully acknowledges the support of a Gerhard Stanford Graduate Fellowship.

REFERENCES

- [1] D. A. B. Miller, “Rationale and challenges for optical interconnects to electronic chips,” *Proc. IEEE*, vol. 88, pp. 728–749, June 2000.
- [2] E. A. De Souza, M. C. Nuss, W. H. Knox, and D. A. B. Miller, “Wavelength-division multiplexing with femtosecond pulses,” *Opt. Lett.*, vol. 20, pp. 1166–1168, 1995.
- [3] G. A. Keeler, D. Agarwal, B. E. Nelson, N. C. Helman, and D. A. B. Miller, “Performance enhancement of an optical interconnect using short pulses from a modelocked diode laser,” in *CLEO 2002 Tech. Dig.*, 2002, p. 163.

- [4] G. A. Keeler, B. E. Nelson, D. Agarwal, and D. A. B. Miller, "Skew and jitter removal using short optical pulses for optical interconnection," *IEEE Photon. Technol. Lett.*, vol. 12, pp. 714–716, June 2000.
- [5] D. Agarwal and D. A. B. Miller, "Latency in short pulse based optical interconnects," in *2001 IEEE LEOS Annu. Meeting Conf. Proc.*, 2001, pp. 812–813.
- [6] A. V. Krishnamoorthy *et al.*, "The AMOEBA switch: An optoelectronic switch for multiprocessor networking using dense-WDM," *IEEE J. Select. Topics Quantum Electron.*, vol. 5, pp. 261–327, Mar.–Apr. 1999.
- [7] T. K. Woodward, A. V. Krishnamoorthy, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, L. M. F. Chirovsky, S. P. Hui, B. Tseng, D. Kossives, D. Dahringer, D. Bacon, and R. E. Leibenguth, "Clocked-sense-amplifier-based smart-pixel optical receivers," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 1067–1069, Aug. 1996.
- [8] T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, L. A. D'Asaro, L. M. F. Chirovsky, S. P. Hui, B. Tseng, D. Kossives, D. Dahringer, and R. E. Leibenguth, "1-Gb/s two-beam transimpedance smart-pixel optical receivers made from hybrid GaAs MQW modulators bonded to 0.8- μ m silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 422–424, Mar. 1996.
- [9] A. V. Krishnamoorthy and K. W. Goossen, "Optoelectronic-VLSI: Photonics integrated with VLSI circuits," *IEEE J. Select. Topics Quantum Electron.*, vol. 4, pp. 899–912, Nov.–Dec. 1998.
- [10] K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. D. Bacon, D. Dahringer, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 360–362, Apr. 1995.
- [11] D. A. B. Miller, D. S. Chemla, T. C. Damen, A. C. Gossard, W. Wiegmann, T. H. Wood, and C. A. Burrus, "Band-edge electroabsorption in quantum well structures: The quantum-confined stark-effect," *Phys. Rev. Lett.*, vol. 53, pp. 2173–2176, 1984.
- [12] J. L. Brubaker, F. B. McCormick, F. A. P. Tooley, J. M. Sasian, T. J. Cloonan, A. L. Lentine, S. J. Hinterlong, and M. J. Herron, "Optomechanics of a free space photonic switch: Components," in *Proc. SPIE*, vol. 1533, 1991, pp. 88–96.
- [13] H. Shi, J. Finlay, G. A. Alphonse, J. C. Connolly, and P. J. Delfyett, "Multiwavelength 10-GHz picosecond pulse generation from a single-stripe semiconductor diode laser," *IEEE Photon. Technol. Lett.*, vol. 9, pp. 1439–1441, Nov. 1997.
- [14] J. Jahns, "Planar packaging of free-space optical interconnections," *Proc. IEEE*, vol. 82, pp. 1623–1631, Nov. 1994.



Bianca E. Nelson (S'97) was born in Milwaukee, WI, in 1973. She received the B.S. degree in physics and astronomy from the University of Iowa, Iowa City, in 1996 and the M.S. and Ph.D. degrees in applied physics from Stanford University, Stanford, CA, in 1998 and 2003, respectively. Her Ph.D. research was done under Prof. David A. B. Miller.

Her research involved wavelength division multiplexed optical interconnect systems and passive optical devices.



Gordon A. Keeler (S'97) was born in Thunder Bay, Canada, in 1973. He received the H.B.Sc. degree in physics from Lakehead University, Thunder Bay, Canada, in 1996 and the M.S. and Ph.D. degrees in applied physics from Stanford University, Stanford, CA, in 1998 and 2003, respectively. His Ph.D. research was done under Prof. David A. B. Miller.

His research involved optical interconnects and optoelectronic devices.



Diwakar Agarwal (S'95) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Bombay, India, in 1994 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1996 and 2002, respectively. His Ph.D. research was done under Prof. David A. B. Miller.

His research interests involve optical–electrical electrical–optical interfaces, and optical interconnect systems.



Noah C. Helman (S'02) received the A.B. degree in physics from Harvard University, Cambridge, MA, in 1998 and the M.S. degree in applied physics from Stanford University, Stanford, CA, in 2002. He is currently working toward the Ph.D. degree in applied physics under Prof. David A. B. Miller at Stanford University.

His research involves semiconductor modulators in optical interconnect systems.



David A. B. Miller (M'84–SM'89–F'95) received the B.Sc. degree from the University of St. Andrews, St. Andrews, U.K., and the Ph.D. degree from Heriot-Watt University, Edinburgh, U.K., in 1979.

He was with Bell Laboratories, from 1981 to 1996, as a department head from 1987, latterly of the Advanced Photonics Research Department. He is currently the W. M. Keck Professor of Electrical Engineering at Stanford University, Stanford, CA, and the Director of the Ginzton and Solid State and Photonics Laboratories, Stanford, CA. His research interests include quantum-well optoelectronic physics and devices, and fundamental and applications of optics in information, sensing, switching, and processing. He has published more than 200 scientific papers, and holds over 40 patents.

Dr. Miller has served as a Board Member for both the Optical Society of America (OSA) and IEEE Lasers and Electro-Optics Society (LEOS), and in various other society and conference committees. He was President of the IEEE Lasers and Electro-Optics Society in 1995. He was awarded the Adolph Lomb Medal and the R. W. Wood Prize from the OSA, the International Prize in Optics from the International Commission for Optics, and the IEEE Third Millennium Medal. He is a Fellow of the Royal Societies of London and Edinburgh, OSA, and APS, and holds an honorary degree from the Vrije Universiteit Brussels, Belgium.

Dr. Miller has served as a Board Member for both the Optical Society of America (OSA) and IEEE Lasers and Electro-Optics Society (LEOS), and in various other society and conference committees. He was President of the IEEE Lasers and Electro-Optics Society in 1995. He was awarded the Adolph Lomb Medal and the R. W. Wood Prize from the OSA, the International Prize in Optics from the International Commission for Optics, and the IEEE Third Millennium Medal. He is a Fellow of the Royal Societies of London and Edinburgh, OSA, and APS, and holds an honorary degree from the Vrije Universiteit Brussels, Belgium.