

Terabus: Terabit/Second-Class Card-Level Optical Interconnect Technologies

Laurent Schares, *Member, IEEE*, Jeffrey A. Kash, *Fellow, IEEE*, Fuad E. Doany, Clint L. Schow, Christian Schuster, *Senior Member, IEEE*, Daniel M. Kuchta, *Senior Member, IEEE*, Petar K. Pepeljugoski, *Senior Member, IEEE*, Jean M. Trehwella, Christian W. Baks, Richard A. John, Lei Shan, *Member, IEEE*, Young H. Kwon, *Member, IEEE*, Russell A. Budd, Punit Chiniwalla, Frank R. Libsch, *Member, IEEE*, Joanna Rosner, Cornelia K. Tsang, Chirag S. Patel, Jeremy D. Schaub, *Member, IEEE*, Roger Dangel, Folkert Horst, Bert J. Offrein, Daniel Kucharski, Drew Guckenberger, Shashikant Hegde, Harold Nyikal, Chao-Kun Lin, Ashish Tandon, Gary R. Trott, *Member, IEEE*, Michael Nystrom, David P. Bour, *Fellow, IEEE*, Michael R. T. Tan, *Member, IEEE*, and David W. Dolfi

Abstract—In the “Terabus” optical interconnect program, optical data bus technologies are developed that will support terabit/second chip-to-chip data transfers over organic cards within high-performance servers, switch routers, and other intensive computing systems. A complete technology set is developed for this purpose, based on a chip-like optoelectronic packaging structure (Optochip), assembled directly onto an organic card (Optocard). Vertical-cavity surface emitting laser (VCSEL) and photodiode arrays (4×12) are flip-chip bonded to the driver and receiver IC arrays implemented in $0.13\text{-}\mu\text{m}$ CMOS. The IC arrays are in turn flip-chip assembled onto a 1.2-cm^2 silicon carrier interposer to complete the transmitter and receiver Optochips. The organic Optocard incorporates 48 parallel multimode optical waveguides on a $62.5\text{-}\mu\text{m}$ pitch. A simple scheme for optical coupling between the Optochip and the Optocard is developed, based on a single-lens array etched onto the backside of the optoelectronic arrays and on 45° mirrors in the waveguides. Transmitter and receiver operation is demonstrated up to 20 and 14 Gb/s per channel, respectively. The power dissipation of 10-Gb/s single-channel links over multimode fiber is as low as 50 mW.

Index Terms—CMOS integrated circuits (CMOS ICs), integrated optoelectronics, optical interconnections, optical planar waveguides, optical receivers, optical transmitters.

I. INTRODUCTION

THE bandwidth and density requirements for interconnects within high-performance computing systems are growing fast, owing to increasing chip speeds, wider buses, and larger numbers of processors per system. Some of these trends are highlighted in the most recent update of the International Technology Roadmap for Semiconductors [1] and are summarized in [2]. These trends project that many high-performance chips or modules will be increasingly limited by off-chip or off-module bandwidth. While some relief is expected to come from larger caches and software, there will be an increasing need for technologies that provide improved chip-to-chip or module-to-module interconnections in order to continue the price-performance trends that servers and other high-end systems have shown over the years.

Parallel optical interconnects (POIs) promise to enable links with terabit/second-class data transfer capability in a small form factor, at higher density and with less constraint on link length than electrical interconnects. Several POIs based on multimode fiber (MMF) ribbons with aggregate data rates in the $>100\text{-Gb/s}$ -range have been demonstrated [3]–[5]. Such POIs are designed for links between racks of servers or between boards, over lengths ranging from about one meter up to hundreds of meters. There is little doubt that, for links on this length scale, fiber-based POIs will be increasingly used in high-performance computing systems over the next few years. The question when optical interconnects will penetrate further inside the box for on-board chip-to-chip links is currently debated [2], [6]–[12], and it is being discussed which interconnect architectures may benefit from the high bandwidth and density that optics has to offer.

If optics are to compete with copper-based electrical backplanes for on-board interconnects, significant advances in terms of speed, power consumption, density, and cost have to be made [13], particularly, in the light of the recent progress in high-speed electrical interconnects [14], [15]. However, there are also a number of challenges that make the design of high-density

Manuscript received October 18, 2005. This work was supported in part by the U.S. Defense Advanced Research Projects Agency under the C2OI Program Contract MDA972-03-3-0004.

L. Schares, J. A. Kash, F. E. Doany, C. L. Schow, D. M. Kuchta, P. K. Pepeljugoski, J. M. Trehwella, C. W. Baks, R. A. John, L. Shan, Y. H. Kwon, R. A. Budd, P. Chiniwalla, F. R. Libsch, J. Rosner, C. K. Tsang, and C. S. Patel are with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10591 USA (e-mail: schares@us.ibm.com; jeffkash@us.ibm.com).

C. Schuster was with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10591 USA. He is now with the Technical University of Hamburg-Harburg, D-21071 Hamburg, Germany.

J. D. Schaub was with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10591 USA. He is now with the IBM Austin Research Labs, Austin, TX 78758 USA.

R. Dangel, F. Horst, and B. J. Offrein are with the Zurich Research Laboratory, IBM Research GmbH, Rüschlikon 8803, Switzerland.

D. Kucharski and D. Guckenberger were with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10591 USA, and also with Cornell University, Ithaca, NY 14853 USA. They are now with Luxtera, Carlsbad, CA 92008 USA.

S. Hegde was with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10591 USA. He is now with Georgia Institute of Technology, Atlanta, GA 30332 USA.

H. Nyikal was with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10591 USA. He is now with Stanford University, Stanford, CA 94305 USA.

C.-K. Lin, A. Tandon, M. Nystrom, and D. P. Bour are with Agilent Technologies Laboratories, Palo Alto, CA 94304 USA (e-mail: dave_bour@agilent.com).

G. R. Trott, M. R. T. Tan, and D. W. Dolfi were with Agilent Technologies Laboratories, Palo Alto, CA 94304 USA. They are now with Avago Technologies, San Jose, CA 95131 USA.

Digital Object Identifier 10.1109/JSTQE.2006.881906

broadband on-board electrical interconnects difficult. The high-frequency losses of current backplane materials [9], [10], signal distortions due to vias/stubs [16], and the nonuniformity of the electrical packages and interconnects result in increased power consumption for equalization of I/O links. Based on the current technologies used for electrical and optical interchip links, there may be a critical bandwidth-length product above which the optical interconnects are favorable in terms of power consumption [17] and signal integrity.

A number of research programs have started to develop components and work on the integration for high-density on-board optical interconnects [18]–[29]. Two-dimensional (2-D) arrays with up to 540 optical transmitter and receiver elements have been demonstrated [30], high-speed driver and receiver circuits with low-power consumption have been designed [31], [32], low-loss polymer materials with optical waveguides have been developed [33]–[35], and schemes that allow optical coupling between optoelectronic modules and waveguides on backplanes compatible with manufacturing processes are being pursued [36], [37]. However, it is challenging to fulfill all these requirements together, and to develop simple packaging processes that permit the dense integration of high-speed components.

The Terabus project addresses these packaging, density, and speed issues, and a complete technology set is developed in order to realize a terabit/second-class optical bus for chip-to-chip interconnects on printed circuit boards. The strategy for reducing the size and increasing the speed is to develop optoelectronic modules that simultaneously push the data rate per line up to 20 Gb/s and the number of channels in the bus up to 48 in order to achieve data transfers approaching 0.5–1 Tb/s. Transmitter and receiver modules with a form factor of 1.2 cm² are designed for low-power operation and will transmit the data over an array of optical waveguides with a 62.5- μ m pitch on an organic card.

This paper is organized as follows. Section II presents an overview of the program and the motivation for certain design choices. Section III describes the Terabus components, namely an optical board with integrated waveguides (Section III-A), 2-D arrays of vertical-cavity surface emitting lasers (VCSELs) and photodiodes (Sections III-B and C), CMOS driver and receiver circuits (Sections D–F), a silicon carrier interposer (Section III-G), and the turning mirrors that couple light between the optoelectronic elements and the waveguides (Section III-H). Section IV describes the assembly, packaging, and thermal management aspects. Section V is devoted to the evaluation of the Terabus package, including optical coupling efficiency (Section V-A), waveguide loss and dispersion (Sections V-B and C), as well as high-speed characterization of the electrical signal path (Section V-D) and the optical components (Sections V-E–H). Section VI summarizes the results and comments on the future work.

II. TERABUS—OVERVIEW OF THE PROJECT

Terabus is based on a chip-like optoelectronic packaging structure (Optochip) that is assembled directly onto an organic card with integrated parallel waveguides (Optocard), forming both electrical and optical connections (Fig. 1). The Optochip




Fig. 1. Schematic view of the Terabus package composed of an Optocard with optical waveguides and transmitter and receiver Optochips.

is a small module consisting of a 2-D array of 48 optoelectronic (OE) devices operating at 985 nm (VCSELs and PIN photodiodes) that is flip-chip bonded to a driver IC, which in turn is flip-chip attached onto a silicon carrier. The silicon carrier provides a unique platform that combines multilayer fine-pitch wiring and through-vias for high-performance electrical interconnects, with the ability to integrate heterogeneous components including integrated circuits and optoelectronic devices using the flip-chip bonding technology. The OE device arrays are backside emitting or illuminated, and include antireflection-coated microlens arrays etched into the substrate. In order to couple the light between the Optochips and Optocard, 45° mirrors are fabricated in the waveguides under the OE devices.

The Terabus project explores the hybrid integration of optics into a server environment. To this end, high bit rates and channel density, along with low power consumption are simultaneously required. The complete solution must also be possible to implement with high reliability and at reasonable cost in comparison to purely electrical alternatives. These requirements drive the overall design, and also raise challenges for the electrical packaging, IC and OE device designs, waveguide design, and optical coupling.

Some examples of the design choices to meet these requirements include the following:

- 1) extensive use of the flip-chip technologies in order to avoid the parasitics associated with wirebonds;
- 2) the choice of surface-laminar-circuitry (SLC) as an organic card because of the higher wiring density allowed by such build-up technologies [38];
- 3) the use of a silicon carrier for the Optochip package because the through-vias allow direct solder attachment of the Optochip to the Optocard along with high-density wiring to the IC [39];
- 4) the use of CMOS integrated circuits (CMOS ICs) to minimize IC power and cost;
- 5) an operating wavelength of 985 nm, which permits a simple optical design with emission through the GaAs and InP substrates without the need to thin the OE substrates. This wavelength also permits the direct integration of lenses into the substrates [3].

The Terabus components and packaging concepts are described in detail in the following sections.

III. COMPONENTS AND CIRCUITS

A. Optocard With Integrated Waveguides

The Optocard is a 15 cm × 15 cm printed circuit board made of the SLC technology. A top view of the Optocard is shown




Fig. 2. Overview of the Optocard (left). Cross section of the Optocard with integrated waveguides on a 62.5- μm pitch (right).

in Fig. 2. The build-up layers of SLC have a dielectric constant of 3.4 and a loss tangent of 0.027. Short differential striplines (typically <10 mm) of 20- μm width with a spacing of 50 μm are used to connect signal probe pads to the sites onto which the Optochips are mounted. The measured attenuation of these lines is 1.5 dB/cm at 20 GHz.

An acrylate layer is deposited on top of the SLC card by doctor blading, and waveguides are photolithographically patterned into this layer by UV exposure through a contact mask [40]. The unexposed regions are removed by a solvent. Upon completion, the cladding-core-cladding stack is thermally baked to complete the cure. The Optocard has 48 integrated multimode waveguides with a cross section of 35 $\mu\text{m} \times 35 \mu\text{m}$ on a 62.5- μm pitch. The waveguide link is 30-cm long and contains one 180° and two 90° bends with a minimum bend radius of 28.5 mm. Measurements of different waveguide samples show that no additional bending loss is observed for bend radii larger than 25 mm.

B. VCSELs

The VCSELs [41] are grown in a metal-organic chemical vapor deposition (MOCVD) reactor on semi-insulating GaAs substrates with multiple strained InGaAs quantum wells. The devices have an oxide-confined structure optimized for low series resistance, low parasitics, and high-speed operation at low current densities. The VCSELs with apertures of 4, 6, and 8 μm are fabricated. The VCSELs are optimized for operation at 70 °C with an emission wavelength around 985 nm. A 4 \times 12 array of 10-Gb/s eye diagrams at 70 °C is shown in Fig. 3. The VCSELs have diameters of 4 μm (rows A–I) and 6 μm (rows J and K), and their bandwidths are above 15 GHz. The bias is 2 mA for the 4- μm devices and 3 mA for the 6- μm VCSELs. The modulation is identical for all 48 devices, and the extinction ratios are above 6 dB on each channel. Fig. 3 also shows a zoom on a 10- and a 20-Gb/s eye of a 6- μm VCSEL at 70 °C.

C. Photodiodes

Photodiodes with mesa device structure are grown on an Fe-doped InP substrate. They are backside illuminated, which means that the light enters through the substrate lens and passes through the p-InGaAs contact before reaching the intrinsic layer. Optical absorption in the p-InGaAs layer is detrimental to the photodiode responsivity, which means that the p-InGaAs layer

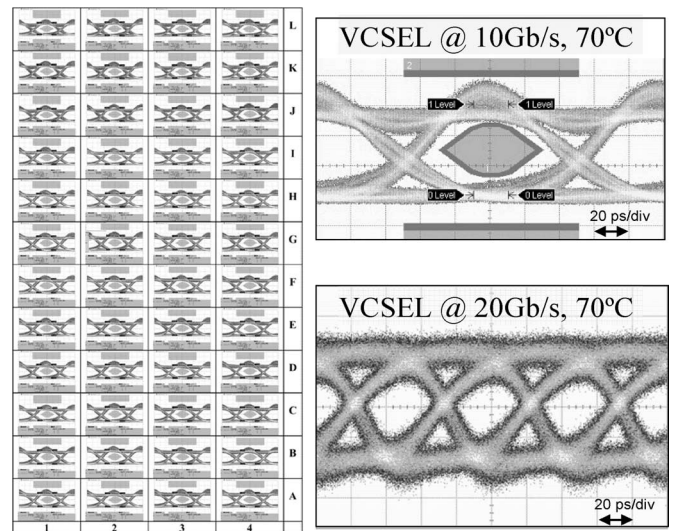


Fig. 3. Array (4 \times 12) of 10-Gb/s VCSEL eyes at 70 °C (left). Enlarged 10- and 20-Gb/s eyes with extinction ratios over 6 dB at 70 °C (right).

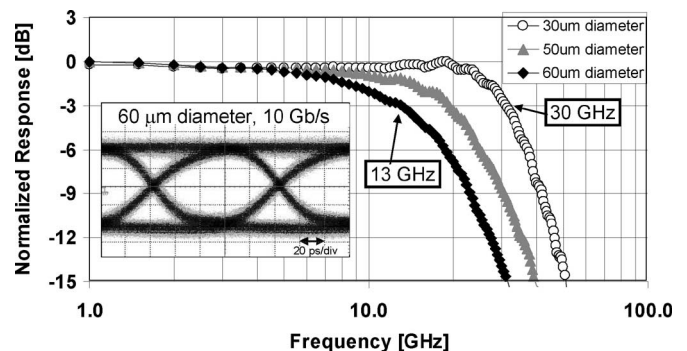


Fig. 4. Frequency response of photodiodes with diameters of 30, 50, and 60 μm at 1.5-V reverse bias. The inset shows a 10-Gb/s eye of a 60- μm photodiode.

needs to be as thin as possible. The responsivity is measured as 0.65 A/W at 985 nm. Photodiodes with diameters of 30, 40, 50, and 60 μm are fabricated. At a reverse bias of 1.5 V, the capacitances range from 90 fF for the smallest to 230 fF for the largest devices. The frequency response is calculated from a Fourier transform of impulse response measurements, using 2-ps pulses at 985 nm. Fig. 4 shows that the 3-dB bandwidths range from 13 GHz (for 60- μm -diameter photodiodes) up to




Fig. 5. Block diagram of an individual channel of the VCSEL driver IC.

30 GHz (for 30- μm diameter photodiodes) at a reverse bias of 1.5 V. Photodiodes of optimal size can be used to trade off lower bandwidth for increased alignment tolerance, depending on the channel bit rate for which a particular Terabus link is designed.

D. CMOS IC Arrays

The laser diode driver (LDD) [42] and receiver (RX) [32] IC arrays were fabricated by IBM in a standard 0.13- μm CMOS process. The LDD and RX arrays share a common electrical pad layout and a 3.9 mm \times 2.3 mm footprint, so that either chip can be attached to a common silicon carrier. The performance of both LDD and RX array ICs benefit from the Terabus packaging configuration: the flip-chip bonding of the OE element to the IC provides a very short electrical path that minimizes parasitic effects at this critical interconnection point. Both arrays consist of 48 individual amplifier elements and utilize two voltage supplies to minimize power dissipation. The power supply to each array is further divided into eight different domains such that blocks of six channels share the same power connections. This configuration enables the characterization of channel-to-channel crosstalk both within and between power blocks. The on-chip power supply decoupling is extensively employed and the layout of the amplifier array elements is carefully considered to minimize intra and interchannel crosstalk.

E. VCSEL Driver Circuits

The 48-channel LDD array is powered by a 1.8-V supply for the input amplifier circuitry and a 3.3-V supply for the output stage and bias. As shown in Fig. 5, each driver circuit contains a differential preamplifier followed by a dc-coupled transconductance output amplifier that supplies the modulation current to the VCSEL. Each driver has differential inputs with a 100- Ω floating termination and fully differential signal paths except for the output stage. Transformer peaking is utilized in all of the predrivers to achieve a large voltage swing and to provide fast transition times to the output stage.

The transconductance output stage has a single-ended current output with an adjacent ground pad to provide a low-inductance




Fig. 6. Block diagram of an individual channel of the receiver IC.

return path for the modulation current. The transconductance amplifier with its high output impedance is well suited for current modulation, and requires less voltage headroom while providing more tolerance to variations in laser series resistance compared to an impedance-controlled voltage driver.

Additionally, the output stage incorporates a fall time compensation (FTC) circuit that improves the optical eye symmetry at high data rates. This circuit decreases the fall time of the driver output by momentarily increasing the tail current in the output stage during high-to-low transitions, providing a preemphasis to the falling edge of the modulation current to compensate for the characteristically slow fall times exhibited by the VCSELs. An eye diagram of the electrical output of the driver with the FTC circuit enabled is shown in Fig. 5 with the falling edge preemphasis clearly visible. Although enabling the FTC circuit results in an asymmetrical electrical eye diagram, when paired with a VCSEL, the symmetry of the optical eye diagram is improved, as Fig. 15 in Section V-E illustrates.

The LDD circuits are designed to be driven with 250-mV peak-to-peak (mVpp) differential input signals. Two variations of the basic driver circuit were designed: a low-power design optimized for 10 Gb/s that is capable of output modulation current swings of 5–6 mA, and a high-speed version that can supply 11 mA of modulation current at data rates up to 20 Gb/s. Further details of the LDD circuit design can be found in [42].

F. Receiver Circuits

The 48-channel receiver array is powered by dual 1.8-V supplies for the amplifier circuits and a separate 1.5–3-V supply for the photodiode bias. Each receiver element is comprised of a low-noise differential transimpedance amplifier (TIA) followed by a limiting amplifier (LA) and an output buffer (Fig. 6). The array is configured so that the TIA and LA circuits occupy the central region of the chip and share one 1.8-V supply, whereas the output buffers are located at the chip edges and are powered with a separate 1.8-V supply. This physical and power supply isolation was implemented to prevent switching noise arising from the large signals at the chip outputs from interfering with the small signals present at the inputs of the sensitive front-end circuits.




Fig. 7. Schematic cross section of a silicon carrier showing the etched cavity, electrical through-vias, and the microstrip surface wiring (left). Layer definition and dimensions of a pair of differential microstrip lines (right).

The input of the TIA is ac-coupled using on-chip three-dimensional (3-D) interdigitated vertical parallel plate capacitors that provide a high capacitance per unit area and a low parasitic capacitance to the substrate. The TIA is a modified common-gate circuit similar to the one described in [32], and utilizes inductive peaking in series with both the TIA inputs and loads to enhance the circuit bandwidth.

The limiting amplifier consists of five cascaded differential Cherry–Hooper gain stages with an offset cancellation feedback loop around the final four stages. The output buffer circuit also employs inductive peaking at its input, and is designed to drive an ac-coupled, off-chip 50-Ω load. The gain of the receiver is 86 dB · Ω, providing up to 600-mV_{pp} differential output signal at the minimum input current of 30 μA.

G. Silicon Carrier

The silicon carrier serves as a packaging platform that is bonded directly onto the Optocard. It contains densely spaced microstrip lines and deep silicon-etched vias through the carrier for electrical signal routing between the Optocard and the Tx/Rx arrays. A cavity etched in the middle of the carrier allows it to hold the OE-on-IC arrays. The carrier measures 1.0 cm × 1.2 cm and has a thickness of 300 μm. The fabrication process is described in [43].

The ability to transfer power and signals from the top surface of the silicon carrier to the Optocard is one critical enabling technology element and relies on a robust process for fabricating electrical through-via connections. Fabrication of through-vias is a multistep process integrating the following: via definition, sidewall insulation, via metallization, connection to terminals or surface wiring on the silicon carrier, and wafer thinning. The through-vias are formed prior to adding the fine-pitch wiring and the cavity. The vias are on a 225-μm pitch with a diameter of 70 μm and a depth of 300 μm. The sidewalls of the vias are electroplated with copper. The metallization is found to be continuous, but the vias are not fully filled with Cu, owing the high aspect ratio (>4:1) and mismatch of the thermal expansion coefficients of Cu and Si. In order to enhance the electrical contact and the stability during temperature cycling, a composite paste is added into the vias.

The silicon carrier has three levels of back-end-of-line (BEOL) CMOS wiring to distribute power, ground, and signals. The differential microstrip transmission lines are routed on the signal level (which is the topmost wiring level), and they interconnect the IC bond pad on the silicon carrier to the through-vias




Fig. 8. Optical coupling scheme between OEs and waveguides (left). Side view of the coupling mirror (right).

that subsequently connect to the Optocard. A schematic cross section of the carrier is shown in Fig. 7, together with the layer dimensions of a pair of differential microstrip signal lines.

A 1.5 mm × 4.2 mm rectangular cavity is etched into the middle of the carrier in order to house the OE-on-IC arrays. The purpose of this cavity is twofold. First, the optical path length between the OEs and the waveguides is minimized, which makes it possible to design a simple optical coupling scheme and hence minimize the coupling losses. Second, the proximity of the OEs to the waveguides allows for a reduced height of solder between the Optochip and the Optocard, which results in increased module reliability and manufacturability.

A top view of the silicon carrier design with a clear central region for OE cavity, differential microstrip lines, and through-vias is shown in Fig. 10. The through-vias for signal, power, and ground are distributed on three sides of the silicon carrier. The fourth side is left free to accommodate space for the waveguides underneath the carrier on the Optocard.

H. Optical Coupling Scheme and Mirror Fabrication

The optical system for coupling light from the OE devices to the waveguides is based on an array of 4 × 12 relay lenses integrated into the OE device. The lenses are etched on the back surface of the GaAs/InP substrate and are aligned to each individual VCSEL/photodiode device on the opposite surface. As shown in Fig. 8, each lens images an OE active area (VCSEL or photodiode) onto a waveguide core. Laser-ablated mirrors are fabricated at either end of the waveguides to allow the 90° coupling into and out of the plane of the Optocard. As shown in Fig. 8, the 45° surface of the mirror is coated with a gold layer in order to achieve high reflectivity.

Based on optical modeling, a lens with a radius of curvature of 110–120 μm and a conic constant of −2 in the GaAs/InP substrate is determined to provide efficient coupling to the 35 μm × 35 μm waveguide core. An optical underfill material of index 1.5, comparable to the index of the polymer waveguides, is used to couple the light between the antireflection-coated OE lens surface and the waveguide core.

Fig. 9 shows a schematic diagram of the 48 waveguides on a 62.5-μm pitch on the Optocard. For illustration, the waveguides are overlaid by a staggered 4 × 12 array of OEs. Mirrors have been fabricated on two of the four rows. The mirrors in the outermost (left) row of Fig. 9 are ablated as three long mirrors for ease of fabrication. Each long mirror couples light between four waveguides and four OE devices in the first row of the Optochip. The smaller mirrors in the third row of Fig. 9 are about 125 μm in width on a 250-μm pitch. This mirror arrangement allows us to




Fig. 9. Schematic diagram of mirror fabrication at the end of the 48 waveguides, showing one row of large mirrors covering four channels, and one row of small mirrors for 12 individual channels (*left*). A 4×12 OE array is superimposed. Photograph of the fabricated sample illuminated by a white light source (*right*).




Fig. 10. Optochip components. A CMOS IC is first bonded to a silicon carrier. Next, an OE is flip-chip bonded onto the IC/silicon-carrier assembly.

couple light into and out of the waveguides at a $125\text{-}\mu\text{m}$ spacing. A fabricated sample of this 2-D mirror array is shown in Fig. 9. The mirrors are illuminated by a white light source incident onto the opposite side of the waveguides (right-hand side). It can be seen that some light is leaking through the neighboring channels of the illuminated channel in the large mirrors, owing to the fact that the individual mirrors of row three only partially extend over the neighboring waveguides and allow some light through to the row with the large mirrors. We are currently refining our laser-ablation process to allow the fabrication of smaller mirrors that can be staggered at $62.5\text{-}\mu\text{m}$ spacing.

IV. PACKAGING

A. Optochip and Optochip-to-Optocard Assembly

The assembly of the Optochip consists of four steps: bonding of the IC to the silicon carrier, bonding of the OE to the IC/silicon carrier assembly, attachment of the Optochip to the Optocard, and addition of an optical underfill material. The first two steps are shown in Fig. 10.

In order to achieve a high-accuracy placement between parts, the following solder hierarchy is used: the ICs and OEs are sequentially flip-chip attached to the silicon carrier using eutectic AuSn (80% Au, 20% Sn) solder, and the Optochip is flip-chip attached to the Optocard using eutectic SnPb (63% Sn, 37% Pb) solder. The AuSn solder is applied to the IC at both the silicon carrier and the OE bond sites. The melting temperature of




Fig. 11. Optochip-on-Optocard assembly (*left*). Bottom view of the silicon carrier with 4×12 VCSEL-array (*right*).

eutectic AuSn is $278\text{ }^\circ\text{C}$ during reflow and more than $400\text{ }^\circ\text{C}$ after it. This permits multiple bonding steps (i.e., the IC remains attached to the silicon carrier while the OE is being attached). The attachment is performed using a flip-chip bonding tool (Suss FC150) with an alignment accuracy of $<2\text{ }\mu\text{m}$. The IC and the silicon carrier are aligned and then bonded at $305\text{ }^\circ\text{C}$. The attachment of the OE is done in the same manner but through the cavity in the silicon carrier.

Shear tests are performed on IC/silicon carrier and OEIC/silicon carrier bonds and result in an average bond strength force of more than 400 g for the IC-to-silicon carrier bond and more than 1 kg for the OE-to-IC bond.

Before attaching the Optochip to the Optocard, the eutectic solder is transferred onto the Optocard using injection-molded solder (IMS) technique [44]. Using IMS, the solder is deposited in a “C”-shape on the Optocard around the waveguide mirrors. The IMS process provides solder columns approximately $200\text{ }\mu\text{m}$ in height, providing adequate height clearance for the Optochip over the $150\text{-}\mu\text{m}$ waveguide height. The final attachment uses a differential temperature between the Optochip and the Optocard and a process that provides solder height correction. A shear strength of more than 10 kg is achieved for the bonds between the Optochip and the Optocard.

Flip-chip bonding does not readily permit active alignment of the OE devices, since this would require electrical powering or sensing of the photocurrent. Therefore, only passive alignment structures are used. The alignment scheme uses features on the OE device to directly contact features on the waveguide. This reduces the number of tolerances that build up if each part is required to be aligned to a central alignment feature. Then, a key constraint for the system is that the OE chip must be visible or in contact with the waveguide during the assembly.

The full Optochip-on-Optocard assembly is shown on the left-hand side of Fig. 11. The right-hand side shows a bottom view of the Optochip with silicon carrier having under-bump metallurgy (UBM) pads. The lenses of a 4×12 VCSEL array are visible through the cavity in the silicon carrier.

B. Thermal Management

Due to the high degree of integration of the Terabus package, thermal challenges arise and need to be addressed. Temperature control in the OE devices is critical; in particular, high-speed performance and lifetime of VCSELs are strongly temperature

sensitive, as is the photodiode leakage. Our strategy to deal with these issues is threefold. First, based on thermal modeling of the full Optochip, the OE devices are optimized for operation at 70 °C at the contact pads. Second, the ICs are designed for low-power operation, with a total link power consumption of less than 100 mW per channel being targeted in an initial phase and 50 mW per channel being a more aggressive objective. Third, the thermal simulations of the full package suggest that an additional cooling system able to handle a heat flux of up to 60 W/cm² is necessary [45]. During the evaluation phase, cooling is performed by putting a heat pipe that is in contact with the backside of the IC.

V. EVALUATION AND RESULTS

A. Optical Coupling Efficiency and Tolerances

A high coupling efficiency between the Optochip and the Optocard is required in order to comply with the optical power budget of the Terabus link, which specifies a maximum coupling loss of 1.5 dB and a mirror loss of 1 dB at both the transmitter and receiver ends. In this section, the losses and tolerances for optical coupling into and out of the waveguides are measured for both the transmitter and the receiver assemblies.

Coupling at the receiver side is measured for a full Optochip–Optocard assembly with waveguide mirrors on a 125- μm pitch. Light from a 980-nm cw-laser is coupled into the ends of the waveguides at the card edge using a single-mode fiber with index-matching fluid. Using the responsivity of the photodiodes and the loss in a short piece of the waveguide, the measurement of the photocurrent allows us to calculate the coupling efficiency. The average coupling and mirror loss for four 60- μm diameter photodiodes on a 125- μm pitch is 2.4 dB, with a best-case loss of 1.6 dB. This fulfills the optical power budget requirement. The coupling efficiency between the 40- μm photodiodes and the large mirrors ranges between 2.3 and 3.3 dB. These values on an average exceed the power budget, as do the losses between the individual mirrors and the 40- μm photodiodes. We expect that a refinement of the mirror and the waveguide fabrication process will result in an improved coupling. We also note that the 40- μm photodiodes, having a higher bandwidth than the larger devices, may only be necessary at data rates of 15 Gb/s and beyond.

A transmitter Optochip–Optocard assembly has been built, and coupling into the 12 operational channels on either the 125- μm or the 250- μm waveguide pitch is achieved. However, a large coupling loss (–7 dB best-case) is observed, which is mostly attributed to the lack of an index-matching material in this assembly (>3.5 dB estimated loss) and the reduced collection efficiency into the smaller waveguides (22 μm \times 35 μm) from an earlier fabrication run.

The coupling efficiency is also measured between an actively aligned transmitter and an Optocard with waveguide mirrors on a 125- μm pitch. The light at the output of the 2-cm long waveguides is collected by a fiber with a 100- μm core. A subset of five 125- μm spaced channels is measured and shows combined coupling and mirror losses of 2 dB on average, which is within the power budget specifications. The optical power coupled into the fiber is above 1 mW for each measured channel.




Fig. 12. Measurements and simulations of coupling tolerances between VCSEL and waveguide (*left*), and between waveguide and photodiode (*right*).

An alignment tolerance analysis has been performed at both the VCSEL–mirror–waveguide and the photodiode–mirror–waveguide interfaces [43]. The dependence of the coupling efficiency on the alignment offset has been measured in the plane parallel to the waveguides (xy -direction) and in the focal (z) direction. Fig. 12 shows the relative coupling loss as a function of the offset in the x -direction. The tolerance required for less than 0.7 dB (85%) of change in the coupling efficiency is better than ± 13 μm on the VCSEL-side and better than ± 14 μm for coupling to the photodiodes of 60- μm diameter. Similar values have been observed in the y -direction [43]. These tolerances can be readily achieved with the flip-chip bonding tools. In the focal direction, a large tolerance of ± 50 μm is measured.

B. Waveguide Loss Measurements

The propagation loss is measured on a 30-cm-long sample, which contains acrylate waveguides fabricated on an SLC substrate with multiple bends. The laser-ablated mirrors are fabricated at either end, providing 90° coupling out of the plane of the SLC. The mirrors on a 2-D array allow us access to 24 waveguides on a 125- μm pitch. Light from a 980-nm cw-laser is coupled into a single-mode fiber and imaged onto the waveguide core using relay optics. The light output at the opposite end is measured with a large-area photodetector. A reference measurement on a similar 2-cm waveguide sample allows us to calibrate the coupling and mirror losses. The average loss of the 30-cm waveguides is found to be –4.8 dB, with a best channel loss of –3 dB, corresponding to an average of 0.16 dB/cm and a best-case loss of 0.10 dB/cm. These loss measurements are consistent with the loss measurements taken on the edge-coupled linear waveguides.

C. Waveguide Dispersion

The modal dispersion of the waveguides is investigated by propagating 2-ps short pulses from a Ti:Sapphire laser at 990 nm through waveguides of different lengths [46]. The input and output pulses are measured with a 14-GHz photodiode (Picometrix D-25) on a high-speed sampling scope. The impulse responses before and after propagation through a 1-m-long waveguide sample are shown on the left of Fig. 13. The pulse broadening plotted versus the waveguide length on the right graph of Fig. 13 is calculated by deconvoluting the response of the photodiode and the sampling head.




Fig. 13. Impulse response of a 2-ps pulse at 990 nm before and after propagation through a 1-m-long waveguide (*left*). Pulse broadening after 8.3-, 17-, 100-, and 250-cm waveguides (*right*).

The transfer function of different waveguide lengths can be calculated by a fast Fourier transform of the impulse responses. The associated 3-dB bandwidths of 30-cm- and 1-m-long links are found to be above 50 and 39 GHz, respectively. For links shorter than 1 m, waveguide dispersion will not be significant for data rates up to 40 Gb/s. The waveguide dispersion has also been measured at 850 nm in the same samples [46], and a similar pulse broadening has been observed as at 990 nm. We also note that time-domain measurements at 850 nm have been performed on similar waveguides in a different experiment, and little signal degradation at 12.5 Gb/s has been observed [20].

For 2.5-m-long samples, the bandwidth limitation caused by modal dispersion decreases to about 23 GHz. However, note that in organic waveguides of this length, the loss rather than its dispersion becomes the limiting factor, owing to the relatively high intrinsic absorption in the currently known organic materials around 985 nm. Further progress in material technology must be made before waveguides much longer than about 1 m become realistic for multi-gigabit/second backplane communication at this wavelength. Ideally, the waveguide absorption at 985 nm should be decreased to <0.05 dB/cm in materials that are compatible with low-cost printed-circuit board manufacturing processes.

D. Electrical Signal Path

Fig. 14 shows the electrical signal path, which is fully differential and consists of striplines on the Optocard, silicon carrier through-vias, and microstrip lines on the silicon carrier. A link consisting of 7-mm surface wires on the Optocard, 300- μ m deep through-vias, and 5-mm transmission lines on the silicon carrier has been characterized.

The S-parameter measurements in Fig. 14 show that the transmission loss is about 4.5 dB at 20 GHz, of which 2.5–3 dB is due to the silicon carrier microstrip lines. The measured reflections (S11 and S22) are lower than -8 dB. More measurements on silicon carrier transmission lines are presented in [43]. Time-domain measurements are performed using a 40-Gb/s pattern generator and a 50-GHz scope. Fig. 14 shows 20-Gb/s eye diagrams with a PRBS $2^{31} - 1$ pattern before and after the electrical signal path. The vertical eye opening after the full link is about 62% (-4.2 dB) of the input opening at 20 Gb/s. The large timing jitter that can be observed in both the input and the

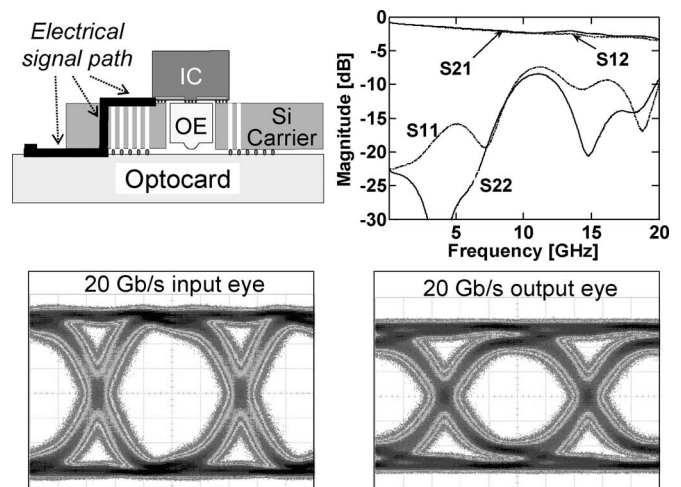


Fig. 14. Frequency- and time-domain measurements of electrical signal path, consisting of a 7-mm stripline on an SLC-Optocard, a 300- μ m through-via, and a 5-mm microstrip line on the silicon carrier. Differential S-parameters (*top right*). Single-ended 20-Gb/s input and output eye diagrams (time scale: 10 ps/div) (*bottom*).




Fig. 15. Single-channel 20-Gb/s eye diagrams of a transmitter OEIC, without preemphasis (*left*), and with falling-edge preemphasis (*right*). The time axis has a 10-ps/div scale, and the OMA is 0 dBm.

output eyes is mostly due to the pattern generator used in these measurements and not due to the interconnect.

E. Optochip Characterization

High-speed testing at the Optochip level is performed by wire-bonding the Optochips onto a printed circuit card with a cavity in the middle. It is possible either to power up all 48 channels at a time or to turn on banks of six channels. Differential wedge probes (GSG–GSG) are brought into contact with the probe pads on top of the silicon carrier and allow simultaneous access to four channels per wedge. A cavity in the test board under the silicon carrier allows us to optically probe the VCSELs or photodiodes with single MMF or fiber ribbons on the backside of the Optochip. The receiver Optochips are characterized by the connecting transmitter and the receiver channels over a 5-m-long 50- μ m MMF link. In all the measurements given later, photodiodes with a 30- μ m diameter are used since they have the largest bandwidth, although their use results in a small reduction in the received optical power.

F. Transmitter Optochip

Fig. 15 shows 20-Gb/s eye diagrams of the high-speed driver paired with a 6.5- μ m diameter VCSEL. It uses 2.5- and 3.3-V supplies and consumes 120 mW including the VCSEL power.




Fig. 16. Single-channel 15- and 20-Gb/s eye diagrams of an assembled transmitter Optochip. The time axis has a 20-ps/div scale, and the OMA is -1 dBm.




Fig. 17. Single-channel 10-Gb/s receiver eyes for OMA values of 0 dBm (left) and -10 dBm (right). The time axis has a 20-ps/div scale.

The optical modulation amplitude (OMA) is 0 dBm and the average VCSEL current is 9 mA. The preemphasis visibly reduces the falling edge tail and increases vertical eye opening by more than 40%. To the best of our knowledge, this is the fastest directly modulated VCSEL transmitter with a CMOS driver demonstrated to date [42].

Fig. 16 shows the 15- and 20-Gb/s eye diagrams of the low-power Tx-Optochip with $8.5\text{-}\mu\text{m}$ VCSELs. The OMA is -1 dBm, and the extinction ratio is 3 dB. The rise and fall times (between 20% and 80%) are 15 ps, and the rms timing jitter is 1.3 ps. The combined power consumption of the VCSEL and the driver IC from a 2.0- and a 3.1-V supply is 76 mW, which corresponds to 3.8 mW per gigabit/second at 20 Gb/s. The 20-Gb/s eye shows some ISI closure owing to the reduced bandwidth of the low-power driver, but error-free 20-Gb/s operation [bit error rate (BER) $<10^{-12}$] has been measured with a reference receiver.

A preliminary frequency-domain crosstalk analysis of the transmitter is performed by modulating one aggressor channel and observing the optical output of adjacent channels with a fiber probe. The aggressor channel is differentially driven at data rates of 10, 15, and 20 Gb/s. A PRBS $2^{31} - 1$ sequence is chosen because it exhibits a more continuous spectrum than the shorter sequences. The optical outputs of the aggressor and the neighboring victim channels are observed with a fast photodiode connected to an 18-GHz spectrum analyzer. The single-channel crosstalk is below 40 dB up to 18 GHz, with the measurement being limited by the dynamic range of the spectrum analyzer.

G. Receiver Optochip

Fig. 17 shows single-ended 10-Gb/s eye diagrams of the receiver Optochip for OMA values of 0 and -10 dBm. While the single-ended output amplitude remains constant at 170 mVpp in both cases, the rise/fall times (20%–80%) increase from 37 ps (OMA = 0 dBm) to 42 ps (OMA = -10 dBm), and a slightly more rms timing jitter is observed at the lower OMA.




Fig. 18. Single-channel sensitivity of a receiver Optochip at 7.5, 10, and 12.5 Gb/s (left). Single-channel receiver eye opening at 10 Gb/s for OMA values of -7 and -10 dBm (right). The pattern is PRBS7.




Fig. 19. Receiver eye diagram at 14 Gb/s for a single-channel link over MMF.

BER measurements of the receiver Optochip are shown in Fig. 18. The OMA sensitivity at BER = 10^{-12} of the receiver is -12 dBm at 7.5 Gb/s and -10.8 dBm at 10 Gb/s (PRBS $2^7 - 1$). At 12.5 Gb/s, the sensitivity decreases to -7.7 dBm, owing to the limited TIA bandwidth in the current design. The photodiode responsivity in these measurements is 0.55 A/W at a reverse bias of 1.5 V. The receiver eye opening is measured at 10 Gb/s. The bathtub curves in Fig. 18 show that the eye opening extrapolated to BER = 10^{-12} is more than 40 ps at an OMA of -7 dBm, and decreases to 27 ps at an OMA of -10 dBm.

Fig. 19 shows a 14-Gb/s eye diagram of a receiver Optochip after transmission over a short 5-m MMF link. The single-ended amplitude of the receiver eye is 170 mVpp. The rms timing jitter values of the transmitter and the receiver are 3.1 and 3.8 ps, respectively. Error-free link operation (BER $<10^{-12}$) is observed with a total link power consumption of 130 mW.

In densely integrated parallel receivers, channel-to-channel crosstalk may induce a power penalty. This penalty can be a result of either imperfect optical coupling due to misalignment between the waveguides or fibers and the photodiodes, or due to on-chip crosstalk between the electrical circuit signal lines or through the substrate. It is therefore important that, first, the optical coupling scheme is carefully designed, and second, the high-speed on-chip transmission lines are shielded [47] and the individual receiver channels decoupled by adding capacitors. A set of initial receiver crosstalk measurements is carried out in the frequency domain. A transmitter is small-signal modulated and coupled to an aggressor channel on the receiver using an MMF probe. The electrical response of a neighboring victim channel is measured and compared to the response of the

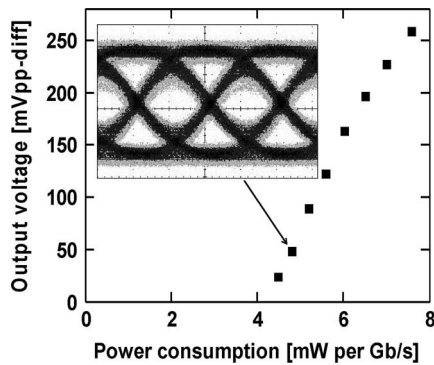


Fig. 20. Power consumption per bit rate for a single-channel 10-Gb/s error-free MMF link. Inset: received 10-Gb/s eye of a 48-mW link (time scale: 20 ps/div), with a power consumption of 22 mW for the transmitter and 26 mW for the receiver.

aggressor. The optical power injected onto the photodiodes is 1 mW. While the worst-case electrical single-channel crosstalk is -23 dB between 10 MHz and 10 GHz, several victim–aggressor pairs exhibited less than -30 dB of crosstalk. If all the channels are turned on simultaneously, we expect these crosstalk values to result in a small power penalty that is accounted for in the optical power budget.

H. Power Consumption and Optical Link Budget

Microprocessors and other ICs mounted on the organic card can generate large amounts of heat. Therefore, optical links operating at very low power are necessary to avoid further increase in the total heat generated on the card beyond the capability of the server’s cooling system. Several factors affect the power consumption of our modules. The transmitter power depends to a large extent on the optical output power (or OMA) of the VCSELs required to overcome a certain link loss. In the case of transmission over optical waveguides, this loss will inevitably be higher than for transmission over fiber. The receiver power consumption strongly depends on the output voltage swing that is required to drive the electrical interface following the receiver Optochip. A single-ended eye diagram of a 10-Gb/s low-power link over 5-m of MMF is shown in Fig. 20. It shows that, by reducing the Rx-supply voltages to <1 V, a 10-Gb/s link can be achieved with a total link power below 50 mW but at the expense of a reduced differential output signal of less than 50 mVpp.

Links with more transmission loss or with higher required output voltages ask for a larger transmitter and/or receiver power. For instance, a differential output swing of above 400 mVpp is observed with a total power consumption of 100 mW for a single-channel fiber-based Optochip-to-Optochip link, including 6 dB of attenuation to simulate the effects of waveguide and coupling losses. Considering achieved values for the 10-Gb/s receiver sensitivity of -10.8 dBm and for the transmitter OMA of above $+1$ dBm, the optical power budget is currently about 12 dB at 10 Gb/s. This budget has to account for coupling and mirror losses, for transmission loss due to the material attenuation in the waveguide link, and for power penalties due to relative intensity noise at the transmitter and the receiver crosstalk.

VI. CONCLUSION

The constituent technologies for a terabit/second-class waveguide-based optical interconnect between chip-like packages have been developed. They include 4×12 -channel CMOS transmitters and receivers with VCSEL and photodiode arrays of <9 mm² footprint each, flip-chip bonded to a silicon carrier interposer of 1.2-cm² size, which is in turn flip-chip bonded to an organic card with 48 integrated waveguides at a 16-channel/mm density that is of 30-cm length. The operating wavelength of 985 nm permits a simple optical design with emission/illumination through lenses directly etched into the substrate of the VCSEL and photodiode arrays. Out-of-plane mirrors have been fabricated in the waveguides on a 125- μ m pitch. A transmitter performance of up to 20 Gb/s per channel and a receiver operation of up to 14-Gb/s has been demonstrated. A 10-Gb/s low-power link over MMF is shown to operate in an error-free manner with less than 5 mW per gigabit/second total power consumption.

The next phase of Terabus will focus on parallel system-level demonstration of the components developed to date. Terabus is an initial step toward a complete technology for chip-to-chip or board-to-board optical buses. Such systems would permit greater bandwidths between processors or modules in high-performance computer systems. While much additional work needs to be carried out before a complete commercial technology becomes realistic, the results summarized earlier are promising and demonstrate that such interconnects are possible.

REFERENCES

- [1] International Technology Roadmap for Semiconductors (ITRS), Update (2004). [Online]. Available: <http://www.itrs.net/Common/2004Update/2004Update.htm>
- [2] A. F. Benner, M. Ignatowski, J. A. Kash, D. M. Kuchta, and M. B. Ritter, “Exploitation of optical interconnects in future server architectures,” *IBM J. Res. Develop.*, vol. 49, no. 4–5, pp. 755–776, 2005.
- [3] L. A. Buckman-Windover, J. N. Simon, S. A. Rosenau, K. S. Giboney, G. M. Flower, L. W. Mirkarimi, A. Grot, B. Law, C.-K. Lin, A. Tandon, R. W. Gruhlke, H. Xia, G. Rankin, M. R. T. Tan, and D. W. Dolfi, “Parallel optical interconnects >100 Gb/s,” *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2055–2063, Sep. 2004.
- [4] D. M. Kuchta, Y. H. Kwark, C. Schuster, C. Baks, C. Haymes, J. Schaub, P. Pepeljugoski, L. Shan, R. John, D. Kucharski, D. Rogers, M. Ritter, J. Jewell, L. A. Graham, K. Schrödinger, A. Schild, and H.-M. Rein, “120-Gb/s VCSEL-based parallel-optical interconnect and custom 120-Gb/s testing station,” *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2200–2212, Sep. 2004.
- [5] 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,” presented at the 55th Electron. Compon. Technol. Conf., Orlando, FL, 2005.
- [6] D. A. B. Miller, “Rationale and challenges for optical interconnects to electronic chips,” *Proc. IEEE*, vol. 88, no. 6, pp. 728–749, Jun. 2000.
- [7] A. F. J. Levi, “Optical interconnects in systems,” *Proc. IEEE*, vol. 88, no. 6, pp. 750–757, Jun. 2000.
- [8] J. Trezza, H. Hamster, J. Iamartino, H. Bagheri, and C. DeCustatis, “Parallel optical interconnects for enterprise class server clusters: Needs and technology solutions,” *IEEE Commun. Mag.*, vol. 41, no. 2, pp. S36–S42, Feb. 2003.
- [9] D. Huang, T. Sze, A. Landin, R. Lytel, and H. L. Davidson, “Optical interconnects: Out of the box forever?,” *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 614–623, Mar.–Apr. 2003.
- [10] E. Mohammed, A. Alduino, T. Thomas, H. Braunisch, D. Lu, J. Heck, A. Liu, I. Young, B. Barnett, G. Vandentop, and R. Mooney, “Optical

- interconnect for ultra-short-reach applications," *Intel Technol. J.*, vol. 8, no. 2, pp. 115–128, May 2004.
- [11] M. J. Koblinsky, B. A. Block, J.-F. Zheng, B. C. Barnett, E. Mohammed, M. Reshotko, F. Robertson, S. List, I. Young, and K. Cadien, "On-chip optical interconnects," *Intel Technol. J.*, vol. 8, no. 2, pp. 129–141, May 2004.
- [12] R. Hemenway, R. R. Grzybowski, C. Minkenberg, and R. Luijten, "Optical packet-switched interconnect for supercomputer applications," *J. Opt. Netw.*, vol. 3, no. 12, pp. 900–913, 2004.
- [13] C. Berger, B. J. Offrein, and M. Schmatz, "Challenges for the introduction of board-level optical interconnect technology into product development roadmaps," presented at the Photonics West, San Jose, CA, Jan. 21–26, 2006, Paper 6124-18.
- [14] M. Meghelli, S. Ryllov, J. Bulzacchelli, W. Rhee, A. Ryljakov, H. Ainspan, B. Parker, M. Beakes, A. Chung, T. Beukema, P. Pepeljugoski, L. Shan, Y. Kwark, S. Gowda, and D. Friedman, "A 10 Gb/s 5-tap-DFE/4-tap-FFE transceiver in 90 nm CMOS," presented at the IEEE Int. Solid-State Circuits Conf., San Francisco, CA, Feb. 2006.
- [15] A. Adamiecki, M. Duell, and J. H. Sinsky, "25-Gb/s electrical duobinary transmission over FR-4 backplanes," *Electron. Lett.*, vol. 41, no. 14, pp. 826–827, Jul. 2005.
- [16] C. Schuster, Y. H. Kwark, G. Selli, and P. Muthana, "Developing a physical model for vias," presented at the DesignCon, Santa Clara, CA, Feb. 6–9, 2006, Paper 8-WA2.
- [17] H. Cho, P. Kapur, and K. C. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2021–2033, Sep. 2004.
- [18] Y. Ishii, S. Koike, Y. Arai, and Y. Ando, "SMT-compatible large-tolerance 'OptoBump' interface for interchip optical interconnections," *IEEE Trans. Adv. Packag.*, vol. 26, no. 2, pp. 122–127, May 2003.
- [19] S.-Y. Cho, S.-W. Seo, N. M. Jokerst, and M. A. Brooke, "Board-level optical interconnection and signal distribution using embedded thin-film optoelectronic devices," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2111–2118, Sep. 2004.
- [20] G.-L. Bona, B. J. Offrein, U. Bapst, C. Berger, R. Beyeler, R. Budd, R. Dangel, L. Dellmann, and F. Horst, "Characterization of parallel optical-interconnect waveguides integrated on a printed circuit board," in *Proc. SPIE Micro-optics, VCSELs, and Photonic Integration*, 2004, vol. 5453, pp. 134–141.
- [21] P. Demmer, R. Mödinger, J. Bauer, F. Ebling, H. Schröder, P. Beil, H. Albrecht, A. Beier, K. Pfeiffer, M. Franke, E. Griese, R. Reuber, and J. Kostelnik, "New generation interconnection technology: Printer circuit boards with integrated optical layers," presented at the 9th IEEE Workshop Signal Propagation on Interconnects, Garmisch-Partenkirchen, Germany, May 10–13, 2005.
- [22] P. Lukowicz, J. Jahns, R. Barbieri, P. Benabes, T. Bierhoff, A. Gauthier, M. Jarczynski, G. A. Russell, J. Schrage, W. Süllau, J. F. Snowdon, M. Wirz, and G. Tröster, "Optoelectronic interconnection technology in the HOLMS system," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 624–635, Mar.–Apr. 2003.
- [23] J. J. Liu, Z. Kalavjian, B. Riehl, W. Chang, G. J. Simonis, A. Apsel, and A. Andreou, "Multichannel ultrathin silicon-on-sapphire optical interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 380–386, Mar.–Apr. 2003.
- [24] R. Bockstaele, M. De Wilde, W. Meeus, O. Rits, H. Lambrecht, J. Van Campenhout, J. De Baets, P. Van Daele, E. van den Berg, M. Klemenc, S. Eitel, R. Annen, J. Van Koetsem, G. Widawski, B. Bareel, J. Goudeau, P. Le Moine, R. Fries, P. Straub, F. Marion, J. Routin, and R. Baets, "A parallel optical interconnect link with on-chip optical access," in *Proc. SPIE Photonics Europe*, Strasbourg, France, Apr. 2004, pp. 124–133.
- [25] J.-S. Kim and J.-J. Kim, "Stacked polymeric multimode waveguide arrays for two-dimensional optical interconnects," *J. Lightw. Technol.*, vol. 22, no. 3, pp. 840–844, Mar. 2004.
- [26] K. B. Yoon, I.-K. Cho, S. H. Ahn, M. Y. Jeong, D. J. Lee, Y. U. Heo, B. S. Rho, H.-H. Park, and B.-H. Rhee, "Optical backplane system using waveguide-embedded PCBs and optical slots," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2119–2127, Sep. 2004.
- [27] B. S. Rho, S. Kang, H. S. Cho, H.-H. Park, S.-W. Ha, and B.-H. Rhee, "PCB-compatible optical interconnection using 45°-ended connection rods and via-holed waveguides," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2128–2134, Sep. 2004.
- [28] S. H. Hwang, M. H. Cho, S.-K. Kang, H.-H. Park, H. S. Cho, S.-H. Kim, K.-U. Shin, and S.-W. Ha, "Passively assembled optical interconnection system based on an optical printed-circuit board," *IEEE Photon. Technol. Lett.*, vol. 18, no. 5, pp. 652–654, Mar. 2006.
- [29] S. Hiramoto and T. Mikawa, "Optical design of active interposer for high-speed chip level optical interconnects," *J. Lightw. Technol.*, vol. 24, no. 2, pp. 927–934, Feb. 2006.
- [30] M. B. Venditti, E. Laprise, J. Faucher, P.-O. Laprise, J. E. A. Lugo, and D. V. Plant, "Design and test of an optoelectronic-VLSI chip with 540-element receiver-transmitter arrays using differential optical signaling," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 361–379, Mar.–Apr. 2003.
- [31] C. Kromer, G. Sialm, C. Berger, T. Morf, M. Schmatz, F. Ellinger, D. Erni, G.-L. Bona, and H. Jäckel, "A 100 mW 4 × 10 Gb/s transceiver in 80-nm CMOS for high-density optical interconnects," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, Feb. 2005, pp. 334–336.
- [32] D. Guckenberger, J. D. Schaub, D. Kucharski, and K. T. Kornegay, "1 V, 10 mW, 10 Gb/s CMOS optical receiver front-end," presented at the IEEE RFIC Symp., Long Beach, CA, Jun. 2005.
- [33] L. Eldada and L. W. Shacklette, "Advances in polymer integrated optics," *IEEE J. Sel. Topics Quantum Electron.*, vol. 6, no. 1, pp. 54–68, Jan.–Feb. 2000.
- [34] A. Yeniay, R. Gao, K. Takayama, R. Gao, and A. F. Garito, "Ultra-low-loss polymer waveguides," *J. Lightw. Technol.*, vol. 22, no. 1, pp. 154–158, Jan. 2004.
- [35] Y. Kuwana, S. Takenobu, K. Takayama, S. Yokotsuka, and S. Kodama, "Low loss and highly reliable polymer optical waveguides with perfluorinated dopant-free core," presented at the Opt. Fiber Commun. Conf., Anaheim, CA, Mar. 5–10, 2006, Paper OWF2.
- [36] L. Wang, X. Wang, W. Jiang, J. Choi, H. Bi, and R. Chen, "45° polymer-based total internal reflection coupling mirrors for fully embedded intraboard guided wave optical interconnects," *Appl. Phys. Lett.*, vol. 87, no. 14, Sep. 2005, 3 p.
- [37] T. Lamprecht, F. Horst, R. Dangel, R. Beyeler, N. Meier, L. Dellmann, M. Gmür, C. Berger, and B. J. Offrein, "Passive alignment of optical elements in a printed circuit board," presented at the 56th Electron. Compon. Technol. Conf., San Diego, CA, 2006, Paper 371.
- [38] E. D. Blackshear, M. Cases, E. Klink, S. R. Engle, R. S. Malfatt, D. N. de Araujo, S. Oggioni, L. D. LaCroix, J. A. Wakil, G. G. Houghham, N. H. Pham, and D. J. Russell, "The evolution of build-up package technology and its design challenges," *IBM J. Res. Develop.*, vol. 49, no. 4–5, pp. 641–661, 2005.
- [39] J. U. Knickerbocker, P. S. Andry, L. P. Buchwalter, A. Deutsch, R. R. Horton, K. A. Jenkins, Y. H. Kwark, G. McVicker, C. S. Patel, R. J. Polastre, C. Schuster, A. Sharma, S. M. Sri-Jayantha, C. W. Surovic, C. K. Tsang, B. C. Webb, S. L. Wright, S. R. McKnight, E. J. Sprogis, and B. Dang, "Development of next-generation system-on package (SOP) technology based on silicon carriers with fine-pitch chip interconnection," *IBM J. Res. Develop.*, vol. 49, no. 4–5, pp. 725–753, 2005.
- [40] R. Dangel, U. Bapst, C. Berger, R. Beyeler, L. Dellmann, F. Horst, B. Offrein, and G.-L. Bona, "Development of a low-cost low-loss polymer waveguide technology for parallel optical interconnect applications," in *IEEE/LEOS Summer Topical Meetings Tech. Dig.*, San Diego, CA, 2004, 2 p.
- [41] C. K. Lin, A. Tandon, K. D. Djordjev, S. Corzine, and M. R. T. Tan, "High-speed 985 nm bottom-emitting VCSEL arrays for chip-to-chip parallel optical interconnects," *J. Lightw. Technol.*, to be published.
- [42] D. Kucharski, Y. H. Kwark, D. M. Kuchta, K. T. Kornegay, M. R. T. Tan, C.-K. Lynn, and A. Tandon, "A 20Gb/s VCSEL driver with pre-emphasis and regulated output impedance in 0.13 μm CMOS," presented at the IEEE Int. Solid-State Circuits Conf., San Francisco, CA, Feb. 2005.
- [43] C. S. Patel, C. K. Tsang, C. Schuster, F. E. Doany, H. Nyikal, C. W. Baks, R. Budd, L. P. Buchwalter, P. S. Andry, D. F. Canaperi, D. C. Edelstein, R. Horton, J. U. Knickerbocker, T. Krywaczuk, Y. H. Kwark, K. T. Kornegay, J. H. Magerlein, J. Rosner, and E. Sprogis, "Silicon carrier with deep through-vias, fine pitch wiring, and through cavity for parallel optical transceiver," presented at the 55th Electron. Compon. Technol. Conf., Orlando, FL, 2005.
- [44] P. A. Gruber, D. Y. Shih, L. Belanger, G. Brouillette, D. Danovitch, V. Oberson, M. Turgeon, and H. Kimura, "Injection molded solder technology for Pb-free wafer bumping," in *Proc. 54th Electron. Compon. Technol. Conf.*, Las Vegas, NV, 2004, pp. 650–654.
- [45] S. Hegde, L. Shan, D. M. Kuchta, Y. H. Kwark, C. W. Baks, F. E. Doany, J. A. Kash, and J. M. Trehwella, "Thermal-mechanical analysis of terabus

- high-speed optoelectronic package,” in *Proc. 55th Electron. Compon. Technol. Conf.*, Orlando, FL, 2005, pp. 1854–1858.
- [46] F. E. Doany, P. K. Pepeljugoski, A. C. Lehman, J. A. Kash, and R. Dangel, “Measurement of optical dispersion in multimode polymer waveguides,” in *IEEE/LEOS Summer Topical Meetings Tech. Dig.*, San Diego, CA, 2004, 2 p.
- [47] D. Guckenberger, C. Schuster, Y. H. Kwark, and K. Kornegay, “On-chip crosstalk mitigation for densely packed differential striplines using via fence enclosures,” *Electron. Lett.*, vol. 41, no. 7, pp. 412–414, Mar. 2005.

Laurent Schares (S’99–M’03) received the Diploma and Ph.D. degrees in physics in 1981, from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 1998 and 2004, respectively.

He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 2004, where he has been engaged in several projects on optical interconnects, high-speed optical receivers, and equalized multimode fiber links. He is the author or coauthor of more than 40 papers published in technical journals and conference proceedings.

Dr. Schares was the recipient of the IEEE Japan Chapter Student Paper Award in 2002.

Jeffrey A. Kash (SM’94–F’05) received the Ph.D. degree in physics from the University of California, Berkeley, in 1981.

He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1981. Presently, he manages the Optical Link and Systems Design Group, which explores advanced optical links for future servers. He is also the IBM Principal Investigator for the Terabus Project. His previous research has been concerned with the optical properties of hot electrons in semiconductors, including the development of picosecond imaging circuit analysis, which is being used by major semiconductor manufacturers today to debug CMOS ICs.

Dr. Kash is a Fellow of the American Physical Society.

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

He was a Postdoctoral Fellow at the California Institute of Technology, Pasadena. He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1985, as a Research Staff Member, and worked on laser spectroscopy, applied optics, projection displays, and laser material processing for electronic packaging. He is the author or coauthor of numerous technical papers, and is the holder of over 40 U.S. patents. His current research interests include optoelectronic packaging and high-speed optical link and systems design.

Clint L. Schow, photograph and biography not available at the time of publication.

Christian Schuster (S’98–M’00–SM’05) received the Diploma degree in physics from the University of Konstanz, Konstanz, Germany, in 1996, and the Ph.D. degree in electrical engineering from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 2000.

From 2001 to 2006, he was a Research Staff Member with the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he was involved in high-speed optoelectronic package and backplane interconnect modeling, design, and optimization. Starting in October 2006, he will be a Professor of electrical engineering at the Technical University of Hamburg-Harburg, Hamburg, Germany.

Dr. Schuster is a member of the German Physical Society (DPG). He was the recipient of the IEEE TRANSACTIONS ON ELECTROMAGNETIC COMPATIBILITY Best Paper Award in 2001, IBM Research Division Awards in 2003 and 2005, and an IEC DesignCon Paper Award in 2005.

Daniel M. Kuchta (S’85–M’86–SM’97), photograph and biography not available at the time of publication.

Petar K. Pepeljugoski (S’91–M’92–SM’03) received the B.Sc. degree from the Ss. Cyril and Methodius University, Skopje, Macedonia, in 1982, the M.S. degree from the University of Belgrade, Belgrade, Yugoslavia, in 1986, and the Ph.D. degree from the University of California, Berkeley, in 1993.

Currently, he is a Research Staff Member with the IBM T. J. Watson Research Center, Yorktown Heights, NY. His previous research has been concerned with the designing, modeling, prototyping, and characterization of high-speed multimode fiber LAN links and parallel interconnects. He was involved in the development of the 10-Gb Ethernet Standard (IEEE 802.3ae) from its inception. He is the author or coauthor of more than 40 papers published in journals and conference proceedings.

Dr. Pepeljugoski was the recipient of the IEEE recognition by the 10-Gb Ethernet Alliance. He was also the recipient of the Certificate of Appreciation by the Telecommunications Industry Association for his contribution to the Working Group on modal bandwidth of multimode fibers and the development of the U.S. Standards for fiber optic technology. He is a member of the OFC/NFOEC Technical Program Committee.

Jean M. Trewhella, photograph and biography not available at the time of publication.

Christian W. Baks, photograph and biography not available at the time of publication.

Richard A. John, photograph and biography not available at the time of publication.

Lei Shan (M’02) photograph and biography not available at the time of publication.

Young H. Kwark (S’80–M’82) received the B.S.E.E. degree from the Massachusetts Institute of Technology, Cambridge, and the M.S.E.E. and Ph.D. degrees from Stanford University, Stanford, CA, all in electrical engineering.

Currently, he is a Research Staff Member with the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he is concerned with circuit design for optical links and wireless applications. His current research interests include package characterization for high-performance computing platforms.

Russell A. Budd, photograph and biography not available at the time of publication.

Punit Chiniwalla, photograph and biography not available at the time of publication.

Frank R. Libsch (S’80–M’82) received the Ph.D. degree in electrical and computer engineering from Lehigh University, Bethlehem, PA, in 1989.

He has been a Research Staff Member with the IBM T. J. Watson Research Center, Yorktown Heights, NY, since 1989. He has managed groups engaged in advanced flat panel display devices, processes, and test research. He is currently the Manager of the Advanced Optical Interconnect and Packaging Group, which is responsible for defining and executing IBM’s exploratory high-bandwidth waveguide optical interconnect road map for server technology. He is the author or coauthor of over 75 papers published in international journals and conference proceedings, and the author of several book chapters. He is the holder of over 30 U.S. and foreign patents.

Dr. Libsch is a member of the SID, Materials Research Society, New York Academy of Sciences, and Sigma Xi.

Joanna Rosner, photograph and biography not available at the time of publication.

Shashikant Hegde, photograph and biography not available at the time of publication.

Cornelia K. Tsang, photograph and biography not available at the time of publication.

Harold Nyikal, photograph and biography not available at the time of publication.

Chirag S. Patel received the B.S., M.S., and Ph.D. degrees in electrical engineering from the Georgia Institute of Technology, Atlanta.

He has been a Research Staff Member with the Department of Science and Technology, IBM T. J. Watson Research Center, Yorktown Heights, NY, since 2001. His research has been concerned with the design and characterization of advanced and exploratory systems based on silicon carrier technology. He is the author or coauthor of more than 40 papers published in international journals and conference proceedings.

Chao-Kun Lin, photograph and biography not available at the time of publication.

Jeremy D. Schaub (S'93–M'95), photograph and biography not available at the time of publication.

Ashish Tandon, photograph and biography not available at the time of publication.

Roger Dangel, photograph and biography not available at the time of publication.

Gary R. Trott (M'83), photograph and biography not available at the time of publication.

Folkert Horst, photograph and biography not available at the time of publication.

Michael Nystrom, photograph and biography not available at the time of publication.

Bert J. Offrein, photograph and biography not available at the time of publication.

David P. Bour (S'84–M'85–SM'97–F'00), photograph and biography not available at the time of publication.

Daniel Kucharski, photograph and biography not available at the time of publication.

Michael R. T. Tan (M'86), photograph and biography not available at the time of publication.

Drew Guckenberger, photograph and biography not available at the time of publication.

David W. Dolfi, photograph and biography not available at the time of publication.