Polylithic Integration of Electrical and Optical Interconnect Technologies for Gigascale Fiber-to-the-Chip Communication

Anthony V. Mule', *Member, IEEE*, Ricardo A. Villalaz, *Member, IEEE*, Paul Jayachandran Joseph, Azad Naeemi, *Senior Member, IEEE*, Paul A. Kohl, *Member, IEEE*, Thomas K. Gaylord, *Fellow, IEEE*, and James D. Meindl, *Life Fellow, IEEE*

***Abstract*—Polylithic integration of electrical and optical interconnect technologies is presented as a solution for merging silicon CMOS and compound semiconductor optoelectronics. In contrast to monolithic and hybrid integration technologies, polylithic integration allows for the elimination of optoelectronic and integrated optic device-related processing from silicon CMOS manufacturing. Printed wiring board-level and compound semiconductor chip-level waveguides terminated with volume grating couplers facilitate bidirectional optical communication, where fiber-to-board and board-to-chip optical coupling occurs through a two-grating (or grating-to-grating) coupling path. A 27% increase in the electrical signal I/O projected by and 33% increase in the number of substrate-level electrical signal interconnect layers implied by the International Technology Roadmap for Semiconductors (ITRS) projections for the 32-nm technology generation are required to facilitate 10 Tb/s aggregate bidirectional fiber-to-the-chip communication. Buried air-gap channels provide for the routing of chip or board-level encapsulated air-clad waveguides for minimum crosstalk and maximum interconnect density. Optical signals routed on-board communicate with on-chip volume grating couplers embedded as part of a wafer-level batch package technology exhibiting compatible electrical and optical input/output interconnects. Measurements of grating-to-grating coupling reveal 31% coupling efficiency between two slab, nonoptimized, nonfocusing volume grating couplers.**

***Index Terms*—High-speed interconnects, integrated optics, optical, optical interconnects, optoelectronic packaging, system level.**

I. Introduction

FOLLOWING the seminal paper by Goodman [1] suggesting the use of photons for clock distribution within silicon microelectronics, research over the past two decades investigating methods of integrating optical interconnect technologies can be categorized as either monolithic or hybrid in nature. Monolithic integration involves the hetero-epitaxial deposition of III–V compound semiconductors on silicon CMOS whereby optoelectronic sources and detectors are formed in lockstep with field effect transistors and electrical interconnection. In this manner, a silicon chip incorporating both light-emitting and light-detecting devices is constructed from a single fabrication process. Monolithic integration of GaAs on silicon can involve the use of a series of graded Ge_xSi_{1-x} layers to act as an aggregate interposer layer with a lattice constant intermediate to both III–V and silicon materials to lessen the impact of lattice mismatch by taking advantage of the 0.07% difference in lattice constants between Ge and GaAs [2]. Monolithic integration using silicon-based materials represents the longest term solution, as it would avoid the compatibility pitfalls of compound semiconductors. Silicon by itself is a poor light emitter due to its indirect bandgap, which results in the domination of light-emitting carrier recombination by nonradiative recombination routes. Erbium-doped silicon LEDs coupled with poly-Si/SiO_2 waveguides and SiGe detectors have been demonstrated, where photonic emission from Si:Er LEDs at $\lambda_o = 1540$ nm is achieved due to electronic transitions within the Er^{3+} ion [3]. Modulation of optical signals involving a silicon optical modulator based on metal–oxide–semiconductor technology has also been demonstrated with a modulation bandwidth greater than 1 GHz [4].

Hybrid integration involves the attachment of III–V materials or devices produced using a separate fabrication process, thereby allowing for the manufacture of silicon CMOS die in a conventional manner up to and including back-end-of-line metallization. Three approaches for achieving hybrid integration include wafer bonding, flip-chip attachment, and epitaxial liftoff. Fully-processed GaAs substrates have been bonded to silicon-on-insulator wafers prior to CMOS transistor formation, for example [5], while polyimide bonding of GaAs wafers with unprocessed epitaxial layers to fully-processed silicon CMOS substrates is also possible [6]. Examples of high-yield flip-chip attachment of quantum well modulator [7] and vertical cavity surface-emitting laser (VCSEL) [8] arrays to fully-processed silicon CMOS wafers have been demonstrated, where devices are isolated after flip-chip attachment by removing the substrate through a wet chemical etch. Protection of the front-side of devices during substrate removal is achieved by wicking an epoxy

Manuscript received March 16, 2004; revised August 26, 2004. This work was supported by the Semiconductor Research Corporation (SRC) and Microelectronics Advanced Research Corporation (MARCO).

A. V. Mule' is with the Intel Corporation, Portland Technology Development, Hillsboro, OR 97124 USA and also with the Microelectronics Research Center, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: tony.v.mule@intel.com).

R. A. Villalaz and A. Naeemi are with the Microelectronics Research Center, Georgia Institute of Technology, Atlanta, GA 30332 USA.

P. A. Kohl and P. J. Joseph are with the School of Chemical and Biomolecular Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0100 USA and also with the Microelectronics Research Center, Georgia Institute of Technology, Atlanta, GA 30332 USA.

T. K. Gaylord and J. D. Meindl are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA and also with the Microelectronics Research Center, Georgia Institute of Technology, Atlanta, GA 30332 USA.

Digital Object Identifier 10.1109/TADVP.2005.847838

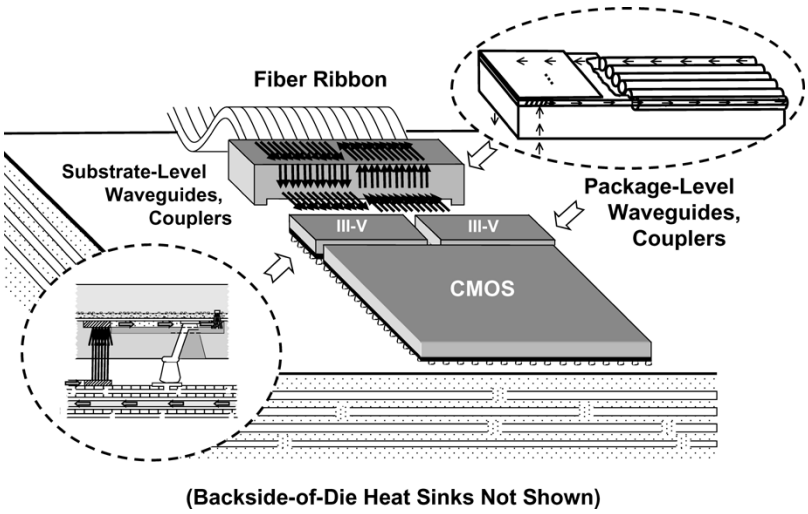

Fig. 1. Gigascale fiber-to-the-chip data communication using polylithic integration.

underfill between GaAs and silicon wafers following attachment, which may be removed after substrate removal with an oxygen plasma. Epitaxial lift-off proceeds in a manner similar to flip-chip bonding with the exception that the substrate upon which III–V device growth occurs is removed prior to host substrate attachment [9]. A thin transparent polyimide diaphragm enables alignment and selective transfer of single or multiple devices to the host substrate. The transfer of GaAs LEDs and inverted metal–semiconductor–metal (I-MSM) photodetectors onto silicon has been achieved using epitaxial lift-off [10]–[12].

Monolithic integration of compound semiconductor materials with silicon CMOS or the construction of silicon-based light emitters represent long-term solutions due to the need for *revolutionary* modification of the CMOS manufacturing process. Hybrid integration reduces additional processing associated with optical interconnect components to that associated with chip packaging, thereby resulting in *evolutionary* changes to CMOS manufacturing. In both cases, the inclusion of silicon chip-level light sources is implied, requiring a significant investment in the process technologies required to augment existing, optimized silicon CMOS manufacturing methodologies.

In this paper, the concept of polylithic integration of electrical and optical interconnect technologies is presented as an alternative solution to monolithic and hybrid integration. An overview of polylithic integration for fiber-to-the-chip communication is discussed, where differences between low-cost fiber-optic fiber-to-the-chip transceiver technologies and gigascale fiber-to-the-chip communication are highlighted. The cost in implementing a 1–10 Tb/s communication system in terms of additional electrical signal I/O and substrate-level electrical signal interconnect layers is calculated based on projections provided by the International Technology Roadmap for Semiconductors (ITRS). Following this overview, a wafer-level batch packaging technology exhibiting encapsulated, chip-length, air-clad and index-defined waveguides in conjunction with volume grating coupler-based optical I/O is described. Board-to-chip coupling of optical signals is demonstrated through the measurement of grating-to-grating coupling efficiency between two nonoptimized, nonfocusing volume diffraction grating couplers.

II. Polylithic Integration: Overview

In contrast to monolithic and hybrid integration technologies, *polylithic integration* of electrical and optical interconnect technologies for fiber-to-the-chip communication eliminates completely the need for modifying the silicon CMOS manufacturing process. Polylithic fiber-to-the-chip communication involves ultrahigh bandwidth *electrical* communication between separate silicon and compound semiconductor die mounted atop a common interconnection substrate, as depicted in Fig. 1. Each substrate is packaged using a wafer-level batch packaging technology (such as that in [13], for example), where the package associated with compound semiconductor die has been augmented to incorporate waveguide and grating-based optical I/O interconnects. The packaging of each die with a common wafer-level batch packaging technology and subsequent near-neighbor flip-chip bonding allows for ultrahigh bandwidth electrical chip-to-chip communication by taking advantage of ultra-low lead parasitics [13] and short chip-to-chip electrical interconnect distances. Fiber-to-the-chip communication begins with optical data originating from a fiber ribbon source being fed into a fiber-to-substrate connector, where each fiber is aligned to optical waveguides terminated by volume grating couplers. Connector-level volume grating couplers communicate in a surface-normal fashion with substrate-level volume grating couplers to pass optical data streams to and from the fiber ribbon. The fiber-to-substrate connector is completely passive in nature, thereby reducing connector fabrication costs. Substrate-level volume grating couplers couple light to and from substrate-level waveguides that direct optical data streams between the fiber-to-substrate connector and III–V compound semiconductor die. Optical data passes to and from the III–V die in a manner similar to that between connector and substrate, where substrate-level volume grating couplers communicate with volume grating couplers embedded within the III–V wafer-level package. Optical data streams received by package-level volume grating couplers are routed through package-level waveguides to optoelectronic detectors for optical-to-electrical conversion. Following conversion, the data is multiplexed between the required number of package pins and

TABLE I
SUMMARY OF FIBER-TO-THE-CHIP TRANSCEIVER TECHNOLOGIES

NAME	REF.	PERFORMANCE (PER TX/RX)	COUPLING METHOD	PACKAGING TECHNOLOGIES	POWER DISSIPATION /TEMP RANGE	THERMAL MANAGEMENT
Optobus	[14-16]	400 Mb/s/ch x 10 (Optobus I); 800 Mb/s/ch x 10 (Optobus II)	Butt-coupling	Surface mount, Wire bonding, flex circuitry	< 2 W, 25-70°C	TAB lead frame, MCM
OETC	[17]	500 Mb/s/ch x 32	45° Fiber Mirror	Wire bonding	8.2 W, 0-70°C	Heat spreader + heat sink
POLO	[18-20]	622 Mb/s/ch x 10 (POLO-1); 1 Gb/s/ch x 10 (POLO-2)	45° Waveguide Mirror	POLO-1: QFP, wire bonding; POLO-2: BGA, wire bonding	< 2.5 W, NA	POLO-1: Heat sink POLO-2: NA
PONI	[20,27]	1.25-2.5 Gb/s/ch x 12	Butt-coupling	BGA, flex circuitry	1.5 W, 0-70°C	Thermocouple + Heat sink
PAROLI	[22]	1.25-2.5 Gb/s/ch x 12	45° Fiber Mirror	Surface mount, wire bonding	1.2 – 3.9 W, 0-85°C	Heat sink
PARABIT	[23,24]	ParaBIT-0: 700 Mb/s/ch x 20 ParaBIT-1: 1.25 Gb/s/ch x 24	45° Waveguide Mirror	ParaBIT-0: QFP, wire bonding; ParaBIT-1: PGA, wire bonding	ParaBIT-0: 8 W, NA; ParaBIT-1: NA	MCM

OETC: Optoelectronic Technology Consortium; POLO: Parallel Optical Link Organization

PONI: Parallel Optics for Network Interconnects; PAROLI: Parallel Optical Links

PARABIT: Parallel Inter-Board Optical Interconnect Technology MCM: Multi-Chip Module

QFP: Quad flat-pack BGA: Ball grid array

PGA: Pin Grid Array NA: Not available

are sent to a silicon CMOS processor through substrate-level electrical interconnections.

The first steps toward creating ultrahigh bandwidth communication between high-performance CMOS microelectronics and optical fibers have already been taken in several optical transceiver technologies developed for fiber-optic communications. Several transceiver technologies achieving fiber-to-the-chip communication between optical fiber ribbons and III–V transmitter and receiver arrays are summarized in Table I. The majority of those technologies listed in Table I incorporate an intermediate optical waveguide path between external fiber and internal active devices. Optical signal coupling between active devices and optical fibers is implemented through the use of either butt-coupling or reflection mirrors. Power dissipation is limited to $\leq$8 W, and is managed through thermal via connections and heat sinks and spreaders. Finally, electrical I/O interconnections are facilitated through wire bonding, quad flat pack, and/or ball grid array packages. For example, a ten-channel fiber optic transceiver module known as the Parallel Optical Link Organization (POLO) module is depicted in Fig. 2 [18]–[20]. In the POLO-1 and POLO-2 modules, waveguides formed from DuPont's polyguide photopolymer connect 62.5/125 μm optical fibers to III–V VCSEL/PIN arrays, where 45° out-of-plane mirrors direct light to and from III–V components at terminal ends of the waveguides. Alignment of waveguide regions to optoelectronic die is aided by photolithographically defined alignment marks. Waveguide widths are tapered from $w_{wg} = 80$ μm (at the VCSEL end) to

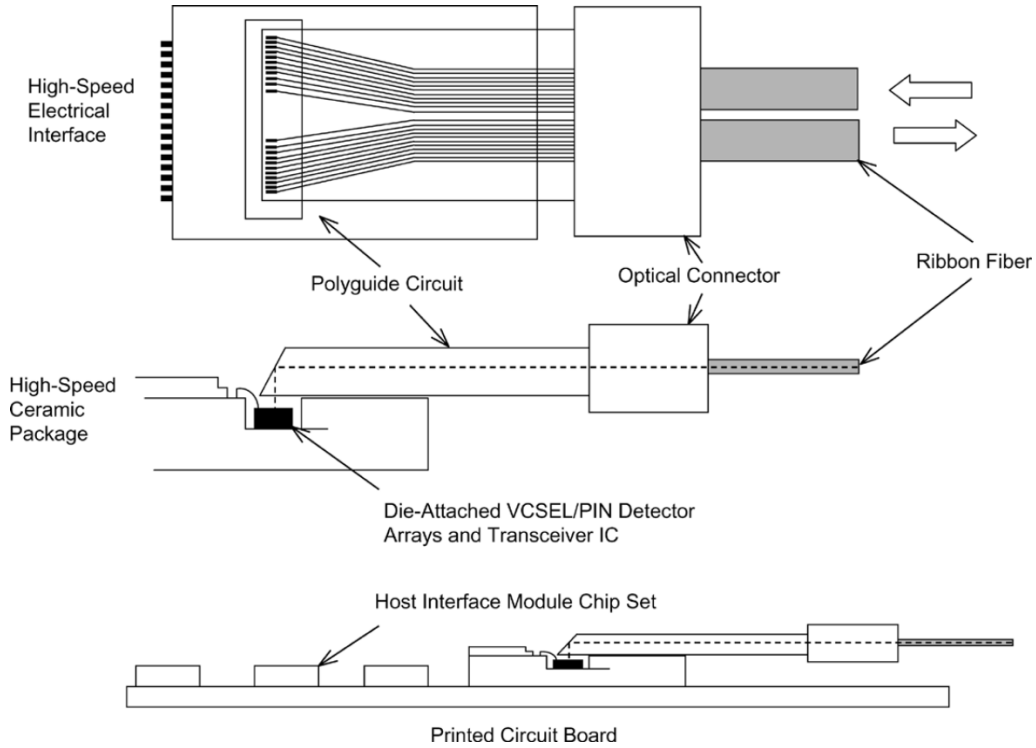

Fig. 2. Components of the POLO-2 module [19].

$w_{wg} = 47$ μm (at the fiber end). As demonstrated in Fig. 2, the waveguide pitch is also tapered to address the pitch mismatch between optical fibers and optoelectronic devices. For the POLO-2 module, alignment tolerances for source-to-waveguide and detector-to-waveguide alignment are 10–20 μm and > 50 μm, respectively. Optical fibers interface with waveguides through an MT ferrule in the POLO-1 module and an MPX push–pull connector in the POLO-2 module. The MT ferrule of POLO-1 snaps into a molded plastic housing attached to the printed wiring board upon which the interconnection substrate also resides. Source and detector arrays within the POLO-1 module are flip-chip bonded to a ceramic quad flat-pack (QFP) lead frame multichip module (MCM) interconnection substrate, where each die is wire-bonded for connection with driver and receiver integrated circuits. The interconnection substrate of the POLO-2 module incorporates a ball grid array (BGA) package for attachment to the substrate. Thermal management of the POLO-1 module is achieved using a heat sink attached to the interconnection substrate. The ten-channel POLO-1 module achieves 622 Mb/s/ch, while the second generation POLO-2 module achieves 1 Gb/s/ch.

An extension of the low-cost packaging technologies associated with fiber-optic transceivers would be required to enable Tb/s fiber-to-the-chip communication between a gigascale microprocessor and compound semiconductor die packaged with an optical I/O-enabled wafer-level packaging technology. Such extensions would pertain to electrical I/O density and performance at high frequencies, package heat removal capacity to allow operation at higher temperatures, and the reduction of optical I/O fabrication and packaging complexity, for example. Key attributes of the technology depicted in Fig. 1 which require enhancement of transceiver-based fiber-to-the-chip technologies and address the constraints imposed by high-performance microprocessor operation are as follows:

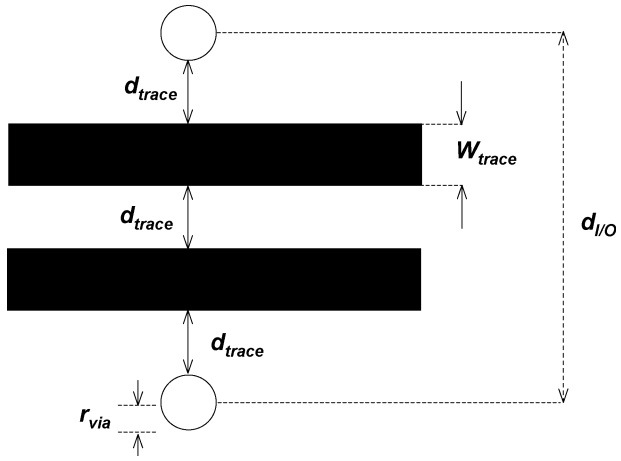

Fig. 3. Top-down view of two escape vias on interconnection substrate (circles) of radius $r_{\rm via}$ enclosing four escape traces of width $W_{\rm trace}$ and edge-to-edge spacing $d_{\rm trace}$.

1) the integration of gigascale CMOS and III–V die on the *same* interconnection substrate;
2) the use of an ultrahigh density of electrical I/O wafer-level packaging technology for flip-chip attachment of CMOS and III–V die to the interconnection substrate;
3) the coupling of optical input/output signals to and from the interconnection substrate before coupling to III–V devices;
4) the use of volume grating couplers for fiber-to-substrate and substrate-to-package coupling.

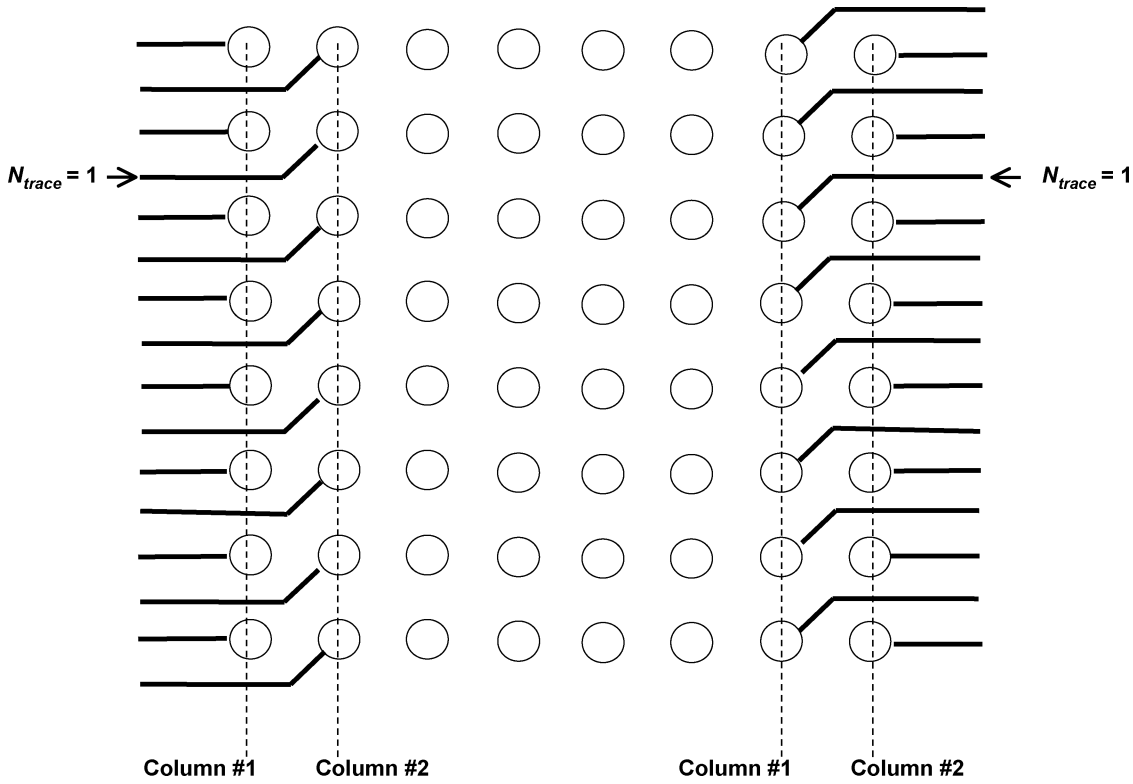

Fig. 4. Interconnection substrate with two columns of signal pads routed on a single layer.

The use of an optical I/O-compatible wafer-level packaging technology allows for a reduction in costs associated with optoelectronic device manufacture through the ability to fabricate electrical and optical I/O interconnect concurrently and the ability to perform wafer-level back-end-of-line testing. Coupling optical signals from external fiber first to the substrate rather than directly to the III–V die allows for both backside-of-die and through-substrate thermal management techniques to be employed in managing excessive power densities in high-performance optoelectronic transmitter and receiver arrays operating at ultrahigh bit rates. To facilitate optical I/O, volume grating couplers are incorporated that provide preferential-order, high-efficiency optical coupling to and from guided-wave regions. Volume diffraction grating couplers offer process compatibility with wafer-level packaging and provide efficient grating-to-grating optical coupling in traversing the board-to-chip propagation path.

Packaging costs associated with gigascale fiber-to-the-chip communication can be expressed in terms of the increase in chip-level electrical I/O and substrate-level interconnect layers associated with a specific chip-to-chip communication bandwidth. The revised number of chip-level signal I/O $N_{\mathrm{I/O,Elec'}}$ and substrate-level electrical interconnect wiring layers $N_{\mathrm{layers'}}$ are defined as the number of microprocessor signal I/O and substrate interconnect layers required for fiber-to-the-chip communication in conjunction with the electrical bandwidth requirements projected by the ITRS for a particular technology generation. Assuming B_{total} is the aggregate, bidirectional communication bandwidth, $N_{\mathrm{I/O,Opt}}$ optical fibers each operating at a bit rate B_{opt} are connected to an interconnection substrate (Fig. 1) where $N_{\mathrm{I/O,Opt}} = B_{\mathrm{total}}/B_{\mathrm{opt}}$. Communication between each optical fiber and III–V transmitter or receiver is facilitated by optical waveguides terminated by volume diffraction grating couplers (Fig. 1). Upon reaching a detector, for example, the optical data stream would be converted to electrons and demultiplexed into several electrical channels on each compound semiconductor die to reflect any differences between B_{opt} and the substrate-level chip-to-chip electrical interconnect bit rate, B_{Elec}. The revised number of microprocessor signal I/O $N_{\mathrm{I/O,Elec'}}$ can therefore be calculated from $N_{\mathrm{I/O,Elec'}} = N_{\mathrm{I/O,Elec}} + B_{\mathrm{total}}/B_{\mathrm{Elec}}$, where $N_{\mathrm{I/O,Elec}}$ is the ITRS-projected signal I/O count per generation.

The required number of electrical interconnect layers depends on the allowed lateral interconnect density between electrical escape vias [25], [26]. Fig. 3 illustrates electrical interconnect escape traces routed between escape vias on an interconnection substrate, while Fig. 4 illustrates a single layer of a two-layer interconnection substrate. In Fig. 3, r_{via} is the escape via radius, d_{trace} is the edge-to-edge spacing between escape traces, and W_{trace} is the trace width. Via pitch is equal to the I/O pad pitch, $d_{\mathrm{I/O}}$, which is given by

$$d_{\mathrm{I/O}} = \sqrt{\frac{A_{\mathrm{chip}}}{N_{\mathrm{I/O}}}}. \tag{1}$$

The trace width required to operate an electrical interconnect trace at a specific bit rate is given by

$$W_{\mathrm{trace}} = L_{\mathrm{trace}}\sqrt{\frac{B_{\mathrm{Elec}}}{K_o}} \tag{2}$$

where L_{trace} is the trace length and $K_o = 6.152 \times 10^7$ Gb/s [26]. This value for K_o assumes the presence of power/ground

TABLE II
ELECTRICAL I/O, INTERCONNECT LAYER REQUIREMENTS FOR 1–10 Tb/s GIGASCALE FIBER-TO-THE-CHIP (GF2C) COMMUNICATION

B_{Total} (Tb/s)	$N_{I/O}$,	% ↑, $N_{I/O}$, vs. $N_{I/O}$	N_{trace},	N_{layer},	% ↑, N_{layer} vs. N_{layer},
1	1827	3	5	3	0
2	1876	5	5	3	0
3	1924	8	5	3	0
4	1973	11	4	4	33
5	2022	14	4	4	33
6	2071	16	4	4	33
7	2119	19	4	4	33
8	2168	22	4	4	33
9	2217	25	4	4	33
10	2266	27	4	4	33

planes surrounding each signal plane to which all power/ground pins are connected. The number of escape traces that can be routed between adjacent vias, N_{trace}, and the resulting number of interconnect layers N_{layers} are given by

$$N_{\text{trace}} = \text{int}\left[\frac{d_{\text{I/O}} - 2r_{\text{via}} - d_{\text{trace}}}{W_{\text{trace}} + d_{\text{trace}}}\right] \tag{3}$$

and

$$N_{\text{layers}} = \text{int}\left[\frac{\sqrt{N_{\text{I/O}}}}{2(N_{\text{trace}} + 1)}\right] \tag{4}$$

respectively [25].

An estimate of the minimum number of electrical interconnect layers required for a particular technology generation both with and without gigascale fiber-to-the-chip I/O communication is found by setting the trace width equal to the maximum of 1) that predicted by (3) and 2) the minimum linewidth projected by the ITRS. For all calculations, it is assumed that $r_{\text{via}} = d_{\text{trace}} = W_{\text{trace}}$, $W_{\text{trace}} = 30$ μm. It should be noted that constraints imposed by IR drop and simultaneous switching noise are assumed to dictate the required number of power and ground I/O, and as such, power and ground I/O are not included in the calculations for $N'_{\text{I/O}}$. In a similar fashion, the number of layers N_{layers} represents signal layers only, and does not include power and ground planes. For example, the 2003 ITRS projects a chip area of 310 mm^2 and chip-to-board speed of 20.5 GHz for the 32-nm technology generation [27]. Assuming a trace length $L_{\text{trace}} = 3\sqrt{A_{\text{chip}}}$ is the maximum length over which electrical communication occurs between neighboring optoelectronic and silicon CMOS die, the minimum trace width supporting a bit rate $B_{\text{Elec}} = 20.5$ Gb/s is found from (2) to be $W_{\text{trace}} = 30$ μm.

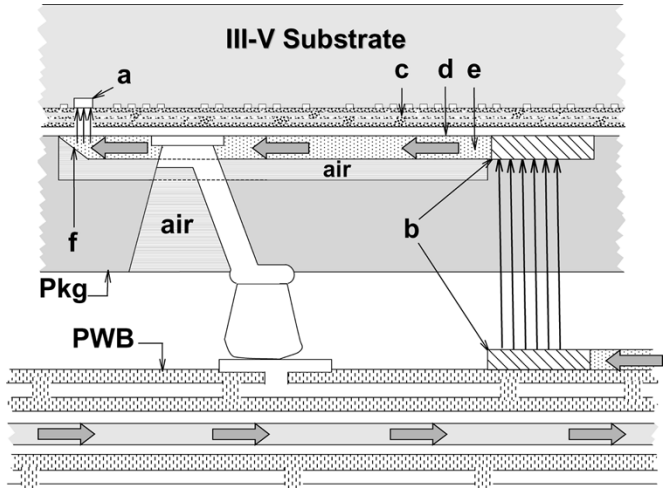

Fig. 5. Cross section of wafer-level batch package incorporating compatible electrical and optical I/O. Labeled are (a) chip-level detector, (b) volume grating coupler, (c) back-end-of-line metallization, (d) passivation, (e) waveguide core, and (f) reflection mirror.

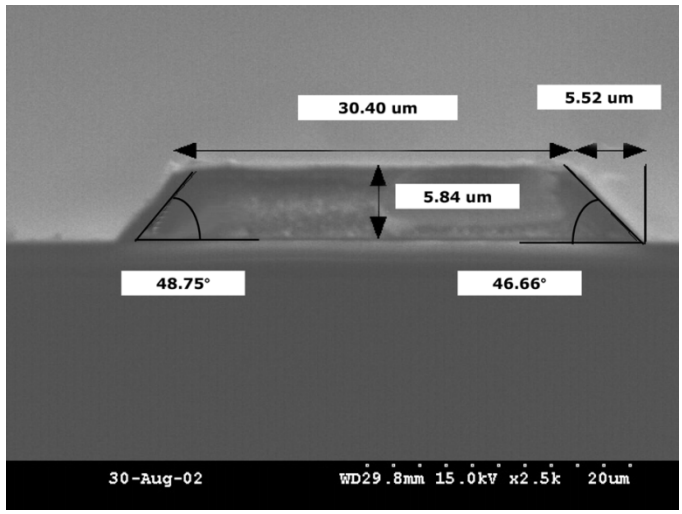

Fig. 6. Reflection mirror exhibiting near-45° mirror facets fabricated using plasma etching.

This value surpasses the minimum-resolvable trace width projected by the ITRS of $W_{\text{trace}} = 18$ μm for the 32-nm generation and, therefore, defines the value of W_{trace} in calculating N_{layers}.

Table II lists the aggregate communication bandwidth, B_{total}, revised total microprocessor I/O count, $N'_{\text{I/O}}$, percent-increase in I/O required for gigascale fiber-to-the-chip communication, the revised number of electrical interconnect traces $N_{\text{trace}'}$ and substrate-level interconnect layers $N_{\text{layer}'}$, and the percent-increase in substrate-level interconnect layers required for fiber-to-the-chip communication. Using (3) and (4), the number of allowed traces N_{trace} and substrate interconnect layers N_{layer} implied by ITRS projections for 32-nm technology are $N_{\text{trace}} = 5$ and $N_{\text{layer}} = 3$, for example. As seen in Table II, the percent-increase in microprocessor electrical I/O required to additionally incorporate an aggregate fiber-to-the-chip communication bandwidth $B_{\text{total}} = 10$ Tb/s for this technology generation is 27% over ITRS-projected

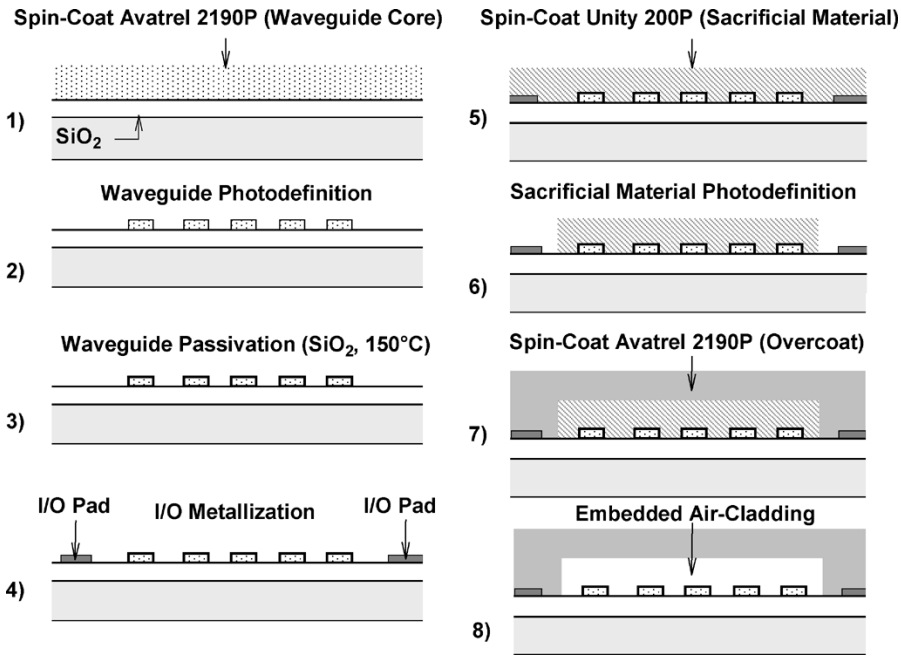

Fig. 7. Fabrication sequence for optical waveguides in embedded air-cladding regions.

values. The percent-increase in the number of substrate-level electrical interconnect layers for this case is 33%, or one signal layer in addition to the number of layers implied by ITRS projections.

III. Polylithic Integration: Optical I/O-Enabled Wafer-Level Batch Package

The success of polylithic integration relies upon the ability to provide optical communication with semiconductor die in a compact, packagable manner commensurate with high-volume manufacturing. To meet this challenge, polylithic integration involves the use of an optical I/O-enabled wafer-level batch packaging technology to package optoelectronic die. A side-view of a wafer-level batch package integrating both electrical and optical input/output interconnection is depicted in Fig. 5. Compact packaging of the electrical/optical system is realized in a manner such that backside-of-die heat sink technologies can be simultaneously incorporated to combat high chip-level power dissipation. In addition, placing optical waveguides within the package eliminates via blockage issues with respect to waveguide routing, leaving blockage introduced by electrical I/O as the dominant routing constraint. Optical power is coupled into the wafer-level package from an off-chip source through preferential-order board and chip-level volume grating couplers. The integration of volume grating couplers within the package following back end-of-line metallization allows for the mitigation of alignment concerns with respect to board-to-package and package-to-chip coupling, as board-to-package input couplers can be sized to cover the range of expected deviations in input beam location.

The option of sizing chip-level receiving grating couplers rather than chip-level detectors represents a key advantage over coupling optical power directly from board-level gratings to chip-level detectors. Assuming the coupling configuration depicted in Fig. 5, chip-level detector dimensions can be reduced to those of the reflection mirror terminating each channel. For a mirror with a 45° slant, the longitudinal component of detector area (i.e., along the direction of light propagation) is on the order of the waveguide thickness. An example of a reflection mirror fabricated from a 6-μm-thick photopolymer film using plasma etching is depicted in Fig. 6, suggesting that a detector area of $6w_{wg}\mu\text{m}^2$ is possible, where w_{wg} is the waveguide width. The fact that waveguide-to-detector coupling mirrors can be made to reside directly above chip-level detectors relegates all alignment concerns to the grating-to-grating coupling interface. Assuming single-mode waveguides, w_{wg} would be on the order of w_{wg} =1–2μm, translating, therefore, into a total required detector area of 6–12 μm^2.

The routing of board and chip-level waveguides terminated with volume grating couplers and/or reflection mirrors can be performed with maximum flexibility by embedding each within an encapsulated air-gap region. An air-gap cladding surrounding each waveguides allows for a maximization in refractive index contrast, Δn between core and cladding regions, which in turn permits smaller bending radii, higher waveguide densities, and the incorporation of simplified tapered reflection mirrors for chip-level waveguide-to-detector coupling. Air-gap technology in wafer-level packaging has been previously developed for enhanced vertical compliance in electrical I/O [28]. The application of the sacrificial material technology presented here represents the first such application to optical waveguide routing, where air cavities span lengths on the order of a die edge and can be defined using low temperature ($\leq$ 160 °C) processing. For this application, photo-definable polycarbonates which decompose at a lower temperature than the thermal decomposition temperature of the waveguide polymer are used. Photo-definition of sacrificial material dramatically simplifies the processing and definition of sacrificial regions, thereby reducing cost and increasing yield.

Fig. 7 illustrates the fabrication sequence associated with the formation of embedded air-clad optical waveguides as part of

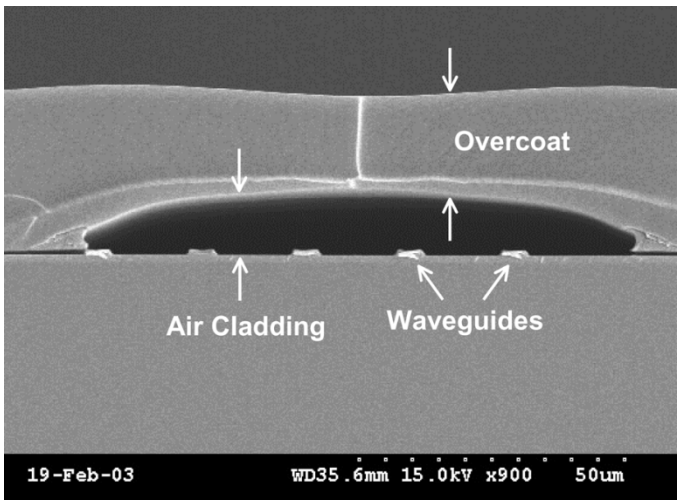

Fig. 8. SEM cross-sectional view of 1/5/25-μm-thick/width/pitch optical waveguides within an encapsulated air cavity. The encapsulated air-cladding region is $\sim$ 15-μm thick.

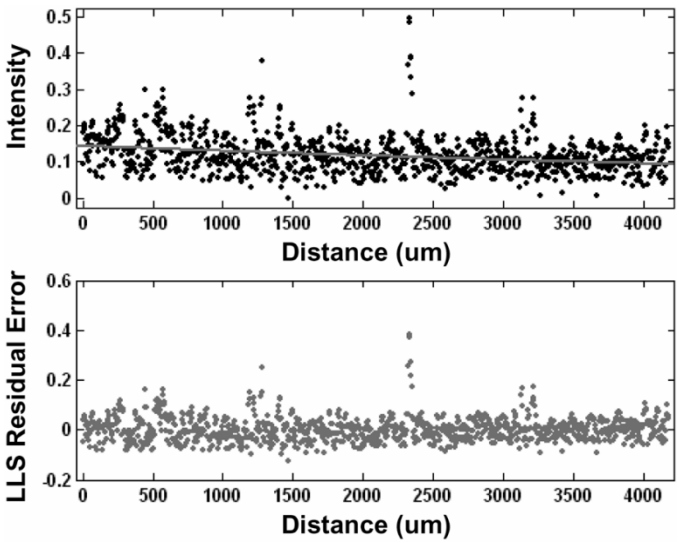

Fig. 9. Robust linear least square fit of loss measurement data for passivated Avatrel waveguide channel. An exposure time of 50 ms was used to record the image. The slope of the fit corresponds to a loss of $\alpha_{wg} = -4.38$ dB/cm. The plot of residual error reveals that the variance is constant along the chosen fit range.

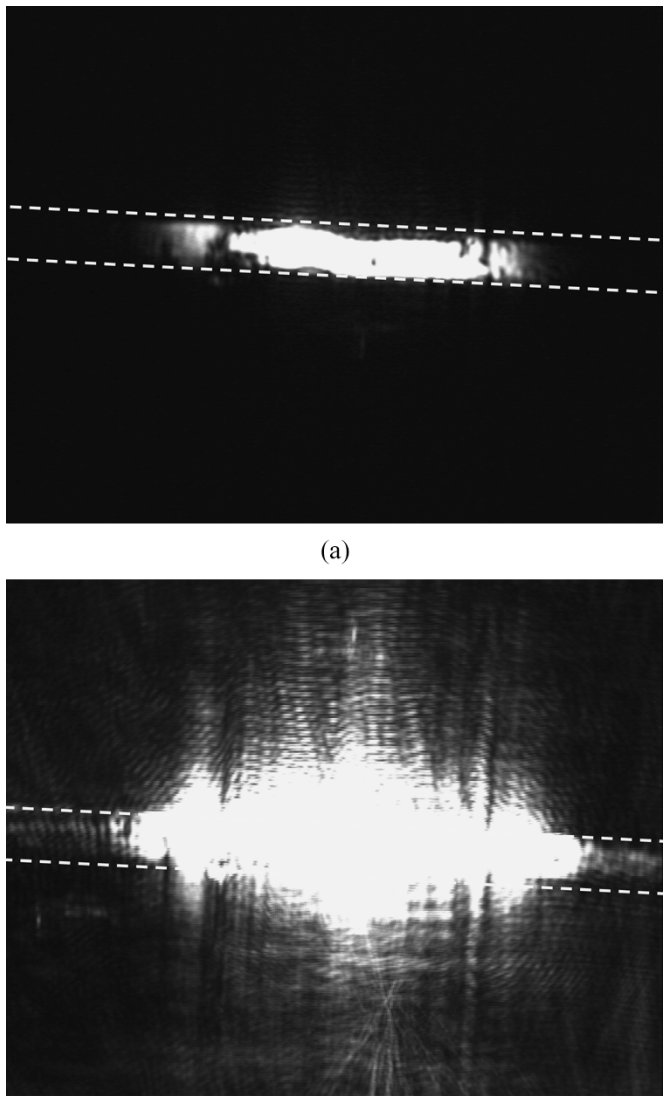

Fig. 10. View of out-coupled power looking longitudinally into end of excited index-defined waveguide under (a) low-power excitation and (b) high-power excitation.

a wafer-level batch packaging technology. Optical waveguides are defined using Avatrel 2190P as the waveguide core material (steps 1 and 2, Fig. 7). Unity 200P, formulated at MiRC, Georgia Tech, using a polycarbonate and photoactive additive dissolved in a suitable solvent, is used as the sacrificial material. Cured films of Avatrel 2190P are soluble in the solvent of the sacrificial material, thereby requiring the deposition of a thin passivation layer (1200 Å) of plasma-enhanced chemical vapor deposition (PECVD) SiO_2 grown at 150 °C prior to sacrificial material deposition (step 3, Fig. 7). After waveguide passivation, I/O pads are defined by depositing 1000/10 000 Å of Ti/Au using a Unifilm dc sputterer (step 4, Fig. 7). Following pad definition, Unity 200P is spin-coated onto the sample (step 5, Fig. 7), after which the wafer is soft-baked on a hotplate at 110 °C for a short time. A deep UV exposure at $\lambda_o = 240$ nm is performed using a Karl Suss MJB 3 mask aligner in conjunction with a positive photomask to define desired air-gap channel regions. Following exposure, the film is baked briefly at 110 °C to decompose exposed areas (step 6, Fig. 7). The sample is subjected to a short (5 s) agitated development using isopropanol and dried with an N_2 gun. A plasma descum using the same process conditions as those used for the I/O photoresist patterning step is performed for 30–60 s to remove residual sacrificial material. To encapsulate the air-gap regions (step 7, Fig. 7), Avatrel 2190P is spin-coated to produce a 23-μm–thick film. The sample is then soft-baked at 80 °C, given a 1 J/cm^2 dose of 365-nm UV radiation using an EVG mask aligner, and baked at 110 °C in a vacuum wire-rack oven. Decomposition of air-gap regions begins during the latter bake step, and is completed with a ramped cure performed in a nitrogen-purged oven, where the oven is ramped from 25 °C to 120 °C at 10°C/min, held for 12 min at 120 °C, then ramped from 120 °C to 150 °C at 1°C/min, held for 30 min, and finally cooled to room temperature at 5°C/min (step 8, Fig. 7).

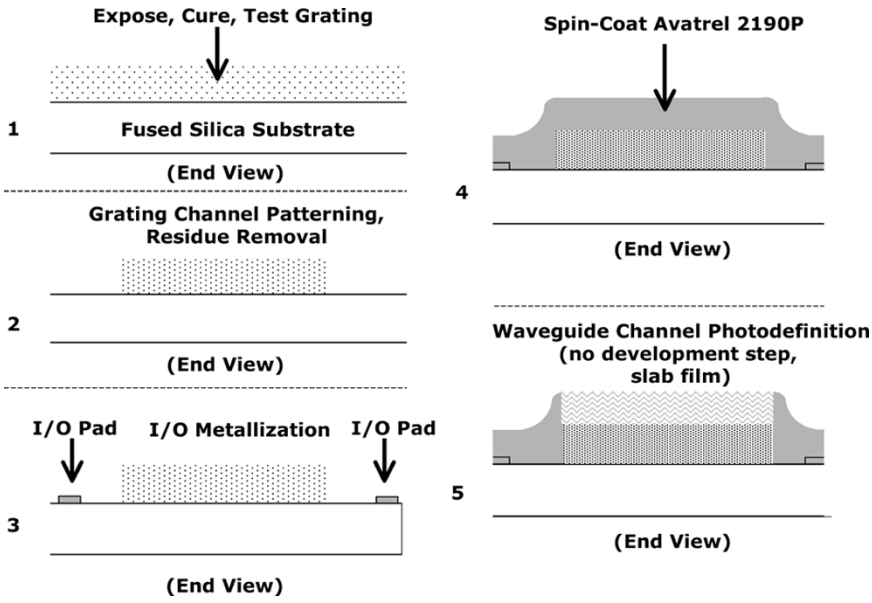

Fig. 11. Fabrication sequence for optical layer of wafer-level batch package with index-defined waveguides and volume grating coupler-based optical I/O.

An example of the successful creation of encapsulated air-clad waveguide channels is depicted in the cross-sectional SEM micrograph of Fig. 8. Propagation loss measurements performed on un-passivated and passivated air-clad 5/25-μm width/pitch Avatrel waveguides using the image capture method indicate loss coefficients $\alpha_{wg} = -$ (3.4-6.8) dB/cm and $-$(2.3–6.6) dB/cm, respectively, revealing no impact on propagation losses due to the presence of the passivation layer. Fig. 9 illustrates the results of a robust bisquare linear least squares fit to a set of measurement data collected from a passivated waveguide channel, where the lack of any trend exhibited by the residual error indicates the goodness of fit.

Experimental realization of the wafer-level batch package with optical I/O depicted in Fig. 5 requires the fabrication of optical interconnects with low-loss waveguides and high-performance gratings as part of the packaging process. Volume grating couplers are diffractive optical structures capable of high-efficiency coupling (e.g., 95% output coupling efficiency [29]) and low-cost, interferometric fabrication. The availability of photopolymer materials suitable for constructing such high-performance couplers, however, is limited. In the case where optical interconnects composed of waveguide channels terminated by volume grating couplers are desired, additional constraints are imposed on the choice of interconnect material by the need for low propagation loss. For example, although the photopolymer Omnidex HRF 600 from Dupont is capable of submicron resolution, low shrinkage, low moisture absorption, and high index modulation for high-performance grating formation, the need for plasma-definition of waveguide channels prohibits the formation of sub-1 dB/cm propagation-loss channel waveguides. Avatrel 2190P from Promerus, LLC, is a spinnable, negative-tone photopolymer capable of producing low-loss, raised-strip waveguide channels. Both polymer materials cure at 150 °C, and cured Omnidex films exhibit negligible solvent-related degradation when encapsulated in mesitylene-based Avatrel. Given the small Δn ($\Delta \mathrm{n} = 9 \times 10^{-3}$) between Avatrel and Omnidex materials, the

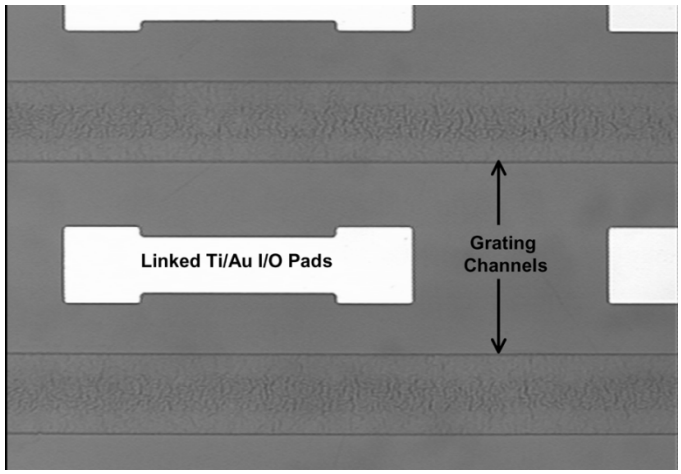

Fig. 12. Patterned grating channels encapsulated in Avatrel situated between linked Ti/Au I/O pads.

construction of an Avatrel/Omnidex waveguide/grating interconnect with negligible reflection loss at the material interface is possible. The creation of such an interconnect requires a degree of adhesion at the material interface sufficient to withstand the wet development step required for raised-strip, air-clad waveguide channels. The adhesion between blanket films of each material in the presence of a wet-development step and absence of additional preparatory processing, however, results in the delamination of Avatrel from Omnidex photopolymer.

Two solutions to the adhesion issue involve 1) skipping the wet development step and relying on the index contrast created through a *positive-tone* photo-definition of Avatrel channels for light confinement leading into grating regions, and 2) passivating Omnidex regions first with a thin layer of SiO_2 prior to wet-development during the creation of raised-strip channels. The degree of adhesion between waveguide and grating regions in either case is confirmed by measuring the scattered/diffracted intensity profile at the material interface of both *index-defined*

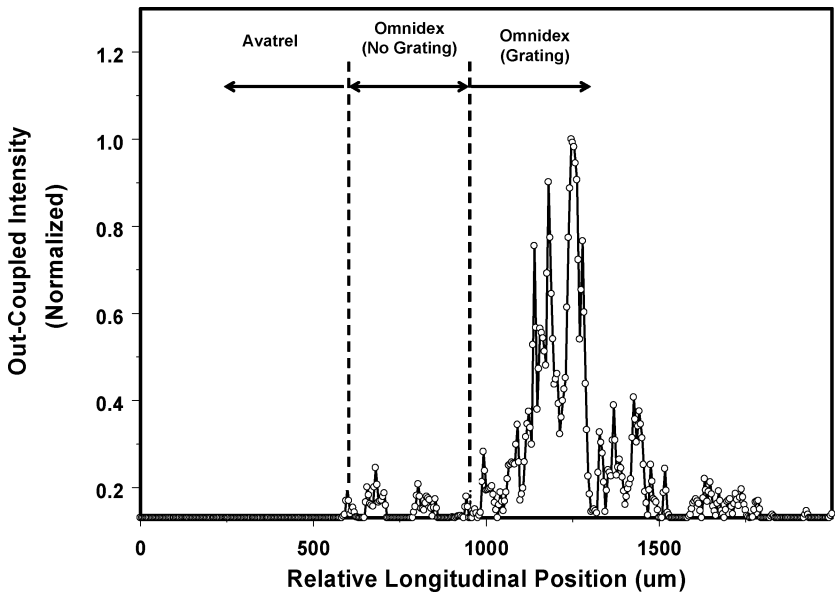

Fig. 13. Scattered/diffracted intensity profile for two-material grating-in-the-waveguide optical interconnect at the waveguide/grating interface. The waveguide region is represented by an index-defined waveguide composed of Avatrel, and the grating region is represented by Omnidex photopolymer.

waveguides (former solution) and air-clad, raised-strip waveguides with glass passivation (latter solution) and noting the lack of reflected power. Figs. 10–13 depict the formation of an optical interconnect layer with index-defined waveguides and volume grating coupler-based optical I/O. In Fig. 10, the end-view of an index-defined waveguide is shown under 1) low and 2) high-power excitation conditions, illustrating the degree of lateral power confinement resulting from the Δn created using a 1 J/cm^2 photo-exposure. Ellipsometer measurements of blanket Avatrel films indicate $\Delta n = 4.5 \times 10^{-3}$ between core and cladding regions for this exposure dose. The fabrication sequence associated with the incorporation of the optical layer as part of a wafer-level batch package is depicted in Fig. 11. Following volume grating coupler design, fabrication, and test atop a fused-silica substrate (step 1, Fig. 11) [29], the (initially 10-mm-wide slab) volume grating coupler is patterned into isolated channels using Au hard mask and O_2-based reactive ion etch processing. Following channel definition and hard mask removal, the sample is soaked briefly in a buffer oxide etch solution to undercut and remove Omnidex residual photopolymer (step 2, Fig. 11). Titanium/Gold I/O pads are defined using the same process described above for air-clad encapsulated waveguides (step 3, Fig. 11). In contrast to the process of Fig. 7, an additional step associated with the process of Fig. 11 is the removal of the SiO_2 passivation layer used to protect diffraction grating channels during I/O pad formation prior to Avatrel encapsulation (step 4, Fig. 11). Once Avatrel has been spin-coated atop patterned Omnidex channels, a 1 J/cm^2 positive photodefinition of waveguide channels leading into grating regions is performed prior to film cure (step 5, Fig. 11). Fig. 12 illustrates a top-view micrograph of patterned volume grating channels situated between alternating rows of linked Ti/Au I/O pads, where the channels and I/O pads are both encapsulated in Avatrel, while Fig. 13 provides an image capture of the scattered and diffracted light associated with a waveguide/grating interface of Fig. 12. The interface between Avatrel and Omnidex regions exhibits no signs of scattering, and grating functionality is reflected in the intensity profile corresponding to the grating region. Image capture-based loss measurements confirm propagation losses of $\alpha_{wg} = -\{5.4–8.8\}$ dB/cm for index-defined channels fabricated in this manner. In addition to index-defined channels, air-clad, raised-strip waveguide channels can be fabricated using the same process as Fig. 11, where the SiO_2 passivation layer is left in place prior to Avatrel deposition, and a *negative-tone* photo-exposure coupled with a wet-development step define waveguide channels. Fig. 14 depicts a micrograph of an optically-excited air-clad, raised-strip, two-material grating-in-the-waveguide interconnect in conjunction with an image capture of the scattered and diffracted intensities of the waveguide and grating region, respectively. As with Fig. 13, no reflection is evident at the waveguide/grating interface. Propagation loss metrics of $\alpha_{wg} = -\{0.5–3.1\backslash\}$ dB/cmhave been recorded for waveguides leading into grating channels located on a single substrate.

The quasi-free space optical coupling between separate board and chip-level volume grating couplers depicted in Fig. 5 (and implied in Fig. 1) facilitates communication between off-chip optical fibers and on-chip sources and detectors. The grating-to-grating coupling efficiency for either source-to-board or board-to-chip optical coupling is defined as the input coupling efficiency of a volume grating coupler when excited by the diffracted intensity of a second volume grating coupler. Measurement of the coupling efficiency between two nonfocusing, nonoptimized, surface-normal grating couplers has been performed by attaching a volume grating coupler atop a transparent fused-silica substrate to a rotation stage and suspending the substrate above a second, optically-excited grating coupler. The optical power diffracted by the optically-excited grating is monitored as it is transmitted through the top substrate as a function of relative angle of rotation. Fig. 15 illustrates the transmitted intensity recorded through the top substrate grating coupler as a function of relative angular position, where the top substrate grating is encapsulated in a slab film of Avatrel. A peak grating-to-grating coupling

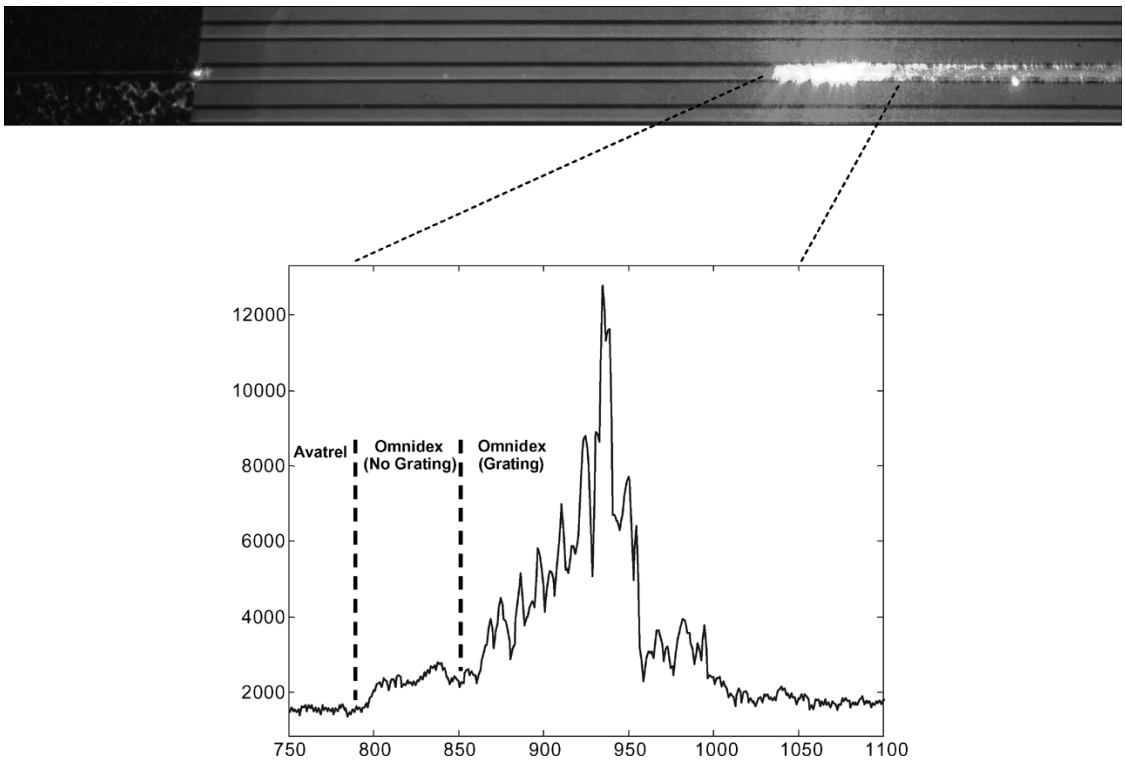

Fig. 14. Raised-strip, air-clad, two-material grating-in-the-waveguide interconnects under optical excitation. The micrograph of interconnect channels is captured under high-power excitation conditions, while the diffracted intensity profile is captured such that image pixel counts fall just below saturation.

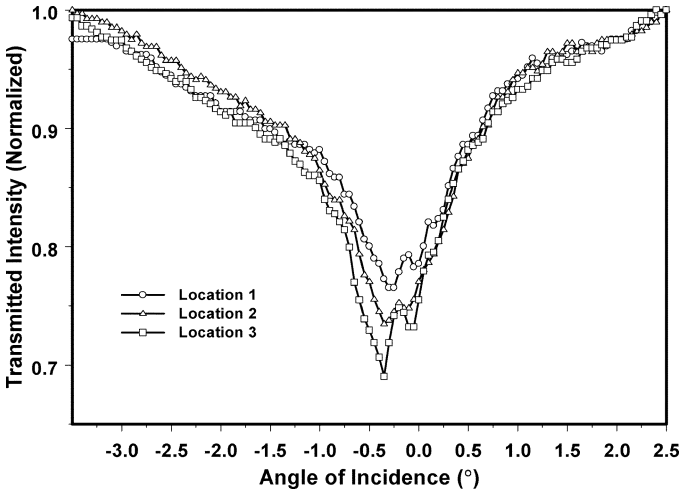

Fig. 15. Transmitted intensity versus relative angle of rotation as measured through a slab volume diffraction grating fabricated atop a transparent substrate and suspended above a second, optically-excited grating. The first (top) volume grating is encapsulated in Avatrel photopolymer to enhance the minima in transmitted intensity (i.e., input coupling efficiency).

efficiency of 31% is recorded over multiple lateral measurement locations [30].

IV. Conclusion

Polylithic integration of electrical and optical interconnect technologies for gigascale microelectronics is presented. In contrast to fiber-to-the-chip technologies implemented in low-cost fiber-optic transceivers, the design of a Tb/s fiber-to-the-chip communication system involves consideration for the performance requirements of gigascale microprocessors, including those associated with electrical I/O count and performance, package heat removal capacity commensurate with Tb/s aggregate bit rates of operation, and the simplification of optical I/O fabrication and packaging. A 27% increase in the electrical I/O projected by and 33% increase in the number of substrate-level electrical interconnect layers implied by 2003 ITRS projections for the 32-nm technology generation are required to facilitate a 10-Tb/s aggregate fiber-to-the-chip communication bandwidth. The fabrication and test of encapsulated, air-clad optical waveguides and two-material, index-defined, grating-in-the-waveguide optical interconnects are also demonstrated as part of a wafer-level batch packaging technology. Unpassivated and passivated air-clad waveguide channels and index-defined waveguide channels exhibit propagation loss metrics $\alpha_{wg} = -$ (3.4–6.8) dB/cm and $-$(2.3–6.6) dB/cm, respectively, while wafer-level package-compatible air-clad, two-material, grating-in-the-waveguide optical interconnects exhibit $\alpha_{wg} = -$ (0.5-3.1) dB/cm. Finally, the use of a grating-to-grating coupling path for source-to-board and board-to-chip optical signal coupling is proposed to enable fiber-to-the-chip communication. A grating-to-grating coupling efficiency of 31% is reported between two slab, nonoptimized, nonfocusing volume diffraction grating couplers. As the gratings are designed neither for outcoupling nor for incoupling, design optimization of the gratings is anticipated to improve substantially the grating-to-grating coupling efficiency.

REFERENCES

[1] J. W. Goodman, F. J. Leonberger, S. C. Kung, and R. A. Athale, "Optical interconnections for VLSI systems," *Proc. IEEE*, vol. 72, no. 7, pp. 850–866, Jul. 1984.

[2] S. M. Ting and E. A. Fitzgerald, "Metal-organic chemical vapor deposition of single domain GaAs on $Ge/Ge_xSi_{1-x}/Si$ and Ge substrates," *J. Appl. Phys.*, vol. 87, pp. 2618–2628, Mar. 2000.

[3] L. C. Kimerling, "Silicon microphotonics," *Appl. Surf. Sci.*, vol. 159–160, pp. 8–13, Jun. 2000.

[4] A. Liu, R. Jones, L. Liao, D. Samara-Rubio, D. Rubin, O. Cohen, R. Nicolaescu, and M. Paniccia, "A high-speed silicon optical modulator based on a metal-oxide capacitor," *Nature*, vol. 427, pp. 615–617, Feb. 2004.

[5] J. M. London, A. H. Loomis, J. F. Ahadian, and C. F. Fonstad, "Preparation of silicon-on-gallium arsenide wafers for monolithic optoelectronic integration," *IEEE Photon. Technol. Lett.*, vol. 11, no. 8, pp. 958–960, Aug. 1998.

[6] T. Nakahara, H. Tsuda, K. Tateno, S. Matsuo, and T. Kurokawa, "Hybrid integration of smart pixels by using polyimide bonding: demonstration of a GaAs p-i-n photodiode/CMOS receiver," *IEEE J. Selected Topics Quantum Electron.*, vol. 5, no. 2, pp. 209–216, Mar./Apr. 1999.

[7] 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, no. 4, pp. 360–362, Apr. 1995.

[8] A. V. Krishnamoorthy, L. M. F. Chirovsky, W. S. Hobson, R. E. Leibenguth, S. P. Hui, G. J. Zydnik, K. W. Goossen, J. D. Wynn, B. J. Tseng, J. Lopata, J. A. Walker, J. E. Cunningham, and L. A. D'asaro, "Vertical-cavity surface-emitting lasers flip-chip bonded to gigabit-per-second CMOS circuits," *IEEE Photon. Technol. Lett.*, vol. 11, no. 1, pp. 128–130, Jan. 1999.

[9] E. Yablonovitch, E. Kapon, T. J. Gmitter, C. P. Yun, and R. Bhat, "Double heterostructure GaAs/AlGaAs thin film diode lasers on glass substrates," *IEEE Photon. Technol. Lett.*, vol. 1, no. 2, pp. 41–42, Feb. 1989.

[10] C. Camperi-Ginestet, M. Hargis, N. M. Jokerst, and M. Allen, "Alignable epitaxial liftoff of GaAs materials with selective deposition using polyimide diaphragms," *IEEE Trans. Photon. Technol. Lett.*, vol. 3, no. 12, pp. 1123–1126, Dec. 1991.

[11] S. W. Bond, O. Vendier, M. Lee, S. Jung, M. Vrazel, A. Lopez-Lagunas, S. Chai, G. Dagnall, M. Brooke, N. M. Jokerst, D. Scott Wills, and A. S. Brown, "A three-layer 3-D silicon system using through-Si vertical optical interconnections and Si CMOS hybrid building blocks," *IEEE J. Select Topics Quantum Electron.*, vol. 5, no. 2, pp. 276–286, Mar./Apr. 1999.

[12] S. Seo, K. K. Lee, S. Kang, S. Huang, W. A. Doolittle, N. M. Jokerst, A. S. Brown, and M. A. Brooke, "The heterogeneous integration of GaN thin-film metal-semiconductor-metal photodetectors on silicon," *IEEE Photon. Technol. Lett.*, vol. 14, no. 2, pp. 185–187, Feb. 2002.

[13] M. S. Bakir, H. A. Reed, H. D. Thacker, C. S. Patel, P. A. Kohl, K. P. Martin, and J. D. Meindl, "Sea of Leads (SoL) ultrahigh density wafer-level chip input/output interconnections for Gigascale Integration (GSI)," *IEEE Trans. Electron Dev.*, vol. 50, no. 10, pp. 2039–2048, Oct. 2003.

[14] D. B. Schwartz, C. K. Y. Chun, B. M. Foley, D. H. Hartman, M. Lebby, H. C. Lee, C. L. Shieh, S. M. Kuo, S. G. Shook, and B. Webb, "A low-cost high-performance optical interconnect," *IEEE. Trans. Compon., Packag., Manuf. Technol. B*, vol. 19, no. 3, pp. 532–539, Aug. 1996.

[15] L. J. Norton, F. Carney, N. Choi, C. K. Y. Chun, R. K. Denton Jr., D. Diaz, J. Knapp, M. Meyering, C. Ngo, S. Planer, G. Raskin, E. Reyes, J. Sauvageau, D. B. Schwartz, S. G. Shook, J. Yoder, and Y. Wen, "Optobus I: a production parallel fiber optical interconnect," in *Proc. IEEE Electronic Components Technology Conf.*, San Jose, CA, May 1997, pp. 204–209.

[16] J. Grula, Optobus Technology, corporate presentation.

[17] Y.-M. Wong, D. J. Muehlner, C. C. Faudskar, D. B. Buchholz, M. Fishteyn, J. L. Brandner, W. J. Parzygnat, R. A. Morgan, T. Mullally, R. E. Leibenguth, G. D. Guth, M. W. Focht, K. G. Glogovsky, J. L. Zilko, J. V. Gates, P. J. Anthony, B. H. Tyrone, T. J. Ireland, D. H. Lewis, D. F. Smith, S. F. Nati, D. K. Lewis, D. L. Rogers, H. A. Aispain, S. M. Gowda, S. G. Walker, Y. H. Kwark, R. J. S. Bates, D. M. Kuchta, and J. D. Crow, "Technology development of a high-density 32-channel 16-Gb/s optical data link for optical interconnection applications for the Optoelectronic Technology Consortium (OETC)," *J. Lightw. Technol.*, vol. 13, pp. 995–1016, Jun. 1995.

[18] K. H. Hahn, K. S. Giboney, R. E. Wilson, J. Straznicky, E. G. Wong, M. R. Tan, R. T. Kaneshiro, D. W. Dolfi, E. H. Mueller, A. E. Plotts, D. D. Murray, J. E. Marchegiano, B. L. Booth, B. J. Sano, B. Madhaven, B. Raghavan, and A. F. J. Levi, "Gigabytes/s data communications with the POLO parallel optical link," in *Proc. IEEE Electronic Components Technology Conf.*, Orlando, FL, May 1996, pp. 28–31.

[19] K. H. Hahn, "POLO- parallel optical links for gigabyte data communications," in *Proc. IEEE Electronic Components Technology Conf.*, Las Vegas, NV, May 1995, pp. 21–24.

[20] L. Buckman, A. Yuen, K. Giboney, P. Rosenberg, J. Straznicky, K. Wu, and D. Dolfi, "Parallel optical interconnects," in *Proc. Lasers Electro-Optics Conf.*, San Francisco, CA, May 2000, pp. 535–536.

[21] Critical Design Guidelines for Successful Application of Parallel Fiber-Optic Modules *HFBR-712BP transmitter and HFBR-722BP receiver specifications sheet.* www.agilent.com [Online]

[22] K. Drogemuller, D. Kuhl, J. Blank, M. Ehlert, T. Kraeker, J. Hohn, D. Klix, V. Plickert, L. Melchior, I. Schmale, P. Hildebrandt, M. Heinemann, F. P. Schiefelbein, L. Leininger, H.-D. Wolf, T. Wipiejewski, and A. Ebberg, "Current progress of advanced high speed parallel optical links for computer clusters and switching systems," in *Proc. IEEE Electronic Components Technology Conf.*, Las Vegas, NV, May 2000, pp. 1227–1235.

[23] K. Katsura, M. Usui, N. Sato, A. Ohki, N. Tanaka, N. Matsuura, T. Kagawa, K. Tateno, M. Hikita, R. Yoshimura, and Y. Ando, "Packaging for a 40-channel parallel optical interconnection module with an over-25 Gbit/s throughput," *IEEE Trans. Adv. Packag.*, vol. 22, no. 4, pp. 551–560, Nov. 1999.

[24] M. Usui, N. Sato, A. Ohki, N. Matsuura, N. Tanaka, K. Enbutsu, M. Amano, M. Hikita, T. Kagawa, K. Katsura, and Y. Ando, "Parabit-1: 60 Gb/s-throughput parallel optical interconnect module," in *Proc. IEEE Electronic Components Technology Conf.*, Las Vegas, NV, May 2000, pp. 1252–1258.

[25] A. Naeemi, P. Zarkesh-Ha, C. S. Patel, and J. D. Meindl, "Performance improvement using on-board wires for on-chip interconnects," in *Proc. IEEE 9th Topical Meeting Electrical Performance Electronic Packaging*, Scottsdale, AZ, Oct. 2000, pp. 325–328.

[26] A. Naeemi, J. Xu, A. V. Mule', T. K. Gaylord, and J. D. Meindl, "Optical and electrical interconnect partition length based on chip-to-chip bandwidth maximization," *IEEE Photon. Technol. Lett.*, vol. 16, pp. 1221–1223, Apr. 2004.

[27] 2003 International Technology Roadmap for Semiconductors, 2003.

[28] H. A. Reed, M. S. Bakir, C. S. Patel, K. P. Martin, J. D. Meindl, and P. A. Kohl, "Compliant wafer-level package (CWLP) with embedded air-gaps for sea of leads interconnections," in *Proc. IEEE Int. Interconnect Technology Conf.*, San Francisco, CA, Jun. 2001, pp. 151–153.

[29] S. M. Schultz, E. N. Glytsis, and T. K. Gaylord, "Volume grating preferential-order focusing waveguide coupler," *Opt. Lett.*, vol. 24, pp. 1708–1710, Dec. 1999.

[30] A. V. Mule', R. Villalaz, T. K. Gaylord, and J. D. Meindl, "Quasifree-space optical coupling between diffraction grating couplers fabricated on independent substrates," *Appl. Opt.*, vol. 43, pp. 5468–5475, Oct. 2004.

Anthony V. Mule' (S'97–M'04) received the B.S.E.E. degree from the University of Illinois, Urbana-Champaign, in 1996 and the M.S.E.E. and Ph.D. degrees in electrical engineering from the Georgia Institute of Technology, Atlanta, in 2004. His graduate thesis work involved integrated optics, diffractive optics, optoelectronics, optoelectronic packaging, and optical materials.

He is currently a Senior Process Integration Engineer in Portland Technology Development, Intel Corporation, Hillsboro, OR, where his work involves back-end-of-line interconnect technology development for the 45-nm technology generation. He is an author of 16 technical publications and three U.S. patents.

Dr. Mule' is currently a Member of the Optical Society of America (OSA).

Ricardo A. Villalaz (M'04) received the B.Sc. degree in electrical engineering, and the M.Sc., and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 1998, 2000, and 2004, respectively.

His current research interests are in volume grating couplers for optical interconnections.

Dr. Villalaz is a Member of the Optical Society of America (OSA) and the Society of Hispanic Professional Engineers, Tau Beta Pi, and Eta Kappa Nu.

Paul Jayachandran Joseph received the Ph.D. degree in physical-organic chemistry from the University of Madras, Chennai, India, in 1997.

From 1994 to 1997, he was a Council of Scientific and Industrial Research—Senior Research Fellow at the Department of Physical Chemistry, University of Madras, where he has carried out an extensive research work in the field of phase transfer catalysis. From 1997 to 2000, he was a Fellow of the National Science Council of Taiwan, R.O.C. Currently, he is a Research Scientist at the School of Chemical and Biomolecular Engineering, Georgia Institute of Technology, Atlanta. His research is focused toward the development of new sacrificial polymeric materials and its application in the field of microelectronics, microfluidics, and microelectromechanical systems.

Azad Naeemi (S'99–M'04–SM'04) received the B.S. degree from Sharif University, Tehran, Iran, in 1994 and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology (Georgia Tech), Atlanta, in 2001 and 2003, respectively.

He joined the faculty at the Microelectronics Research Center at Georgia Tech as a Research Engineer in 2003. From 1994 to 1999, he taught at the Khazra'i Institute and worked as a Design Engineer at Partban and KCR Companies, Tehran. From 1999 to 2003, he worked at the Gigascale Integration (GSI) research group at Georgia Tech as a Graduate Research Assistant. His research interests are in the areas of modeling and optimization for GSI interconnects, exploring alternative interconnecting solutions, chip-package codesign methodologies, and thermal and power management for GSI chips. He has published over 16 papers in refereed journals and conferences and has contributed to a book chapter.

Dr. Naeemi received the Colonel Oscar Cleaver Prize, which recognized him as the outstanding graduate student in the school of Electrical and Computer Engineering at Georgia Tech.

Paul A. Kohl (A'92–M'03) received the B.S. degree from Bethany College, Bethany, WV, and the Ph.D. degree from the University of Texas, Austin, in 1974 and 1978, respectively.

He was with AT&T Bell Laboratories, Murray Hill, NJ, from 1978 to 1989. He is currently Regents' Professor of Chemical Engineering at the Georgia Institute of Technology, Atlanta. His research interests include materials and processing for microelectronic devices and electrochemical engineering.

Thomas K. Gaylord (F'83) received the B.S. degree in physics and the M.S. degreein electrical engineering from the University of Missouri-Rolla and the Ph.D. degree in electrical engineering from Rice University, Houston, TX.

Currently, he is with Georgia Institute of Technology, Atlanta, where he is Julius Brown Chair and Regents' Professor of Electrical and Computer Engineering. He is the author of some 350 technical publications and 25 patents in the areas of diffractive optics, optoelectronics, and semiconductor devices.

Dr. Gaylord received the Curtis W. McGraw Research Award from the American Society for Engineering Education, the IEEE Centennial Medal, the IEEE Graduate Teaching Award, the Georgia Tech Outstanding Teacher Award, and the Engineer of the Year Award from the Georgia Society of Professional Engineers. He is a Fellow of the Optical Society of America (OSA) and of the American Association for the Advancement of Science.

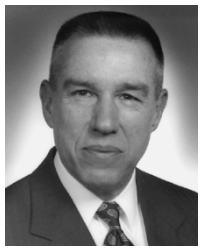

James D. Meindl (LF'97) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Carnegie Institute of Technology (now Carnegie Mellon University), Pittsburgh, PA.

He is the Director of the Joseph M. Pettit Microelectronics Research Center and the Joseph M. Pettit Chair Professor of Microelectronics at the Georgia Institute of Technology, Atlanta. He is also Director of the Interconnect Focus Center, a multiuniversity research effort managed jointly by the Microelectronics Advanced Research Corporation and the Defense Advanced Research Projects Agency for the Department of Defense. His current research interests focus on physical limits on gigascale integration. He has been Senior Vice President for academic affairs and Provost of Rensselaer Polytechnic Institute, Troy, NY, from 1986 to 1993. He was with Stanford University, Stanford, CA, from 1967 to 1986 as the John M. Fluke Professor Of Electrical Engineering, Associate Dean for research in the School of Engineering, Founding Director of the Center for Integrated Systems, Director of the Electronics Laboratories, and Founding Director of the Integrated Circuits Laboratory.

Dr. Meindl is a Fellow of the American Association for the Advancement of Science, and a Member of the American Academy of Arts and Sciences and the National Academy of Engineering and its Academic Advisory Board. He received the Benjamin Garver Lamme Medal from ASEE in 1991, the IEEE Education Medal in 1990, and the IEEE Solid-State Circuits Medal in 1989. He has also been awarded the IEEE Electron Devices Society's J.J. Ebers Award, the 1997 Hamerschlag Distinguished Alumnus Award from Carnegie-Mellon University, as well as five outstanding paper awards from the IEEE ISSCC. He also received the 1999 SIA University Research Award, the IEEE Third Millennium Medal, and most recently, the Georgia Institute of Technology 2001 Distinguished Professor Award.