SAN095-28456
CONF-960153--1

Low-Power Approaches for Parallel, Free-Space Photonic Interconnects

Richard F. Carson, Michael L. Lovejoy, Kevin L. Lear, Mial E. Warren,
Pamela K. Seigal, David C. Craft, Sean P. Kilcoyne, Gary A. Patrizi, and Olga Blum

Sandia National Laboratories
Center for Compound Semiconductor Technology
P. O. Box 5800 MS 0603
Albuquerque, NM 87185-0603

RECEIVED

DEC 27 1995

OSTI

Abstract — Future advances in the application of photonic interconnects will involve the insertion of parallel-channel links into Multi-Chip Modules (MCMs) and board-level parallel connections. Such applications will drive photonic link components into more compact forms that consume far less power than traditional telecommunication data links. These will make use of new device-level technologies such as vertical cavity surface-emitting lasers and special low-power parallel photoreceiver circuits. Depending on the application, these device technologies will often be monolithically integrated to reduce the amount of board or module real estate required by the photonics. Highly parallel MCM and board-level applications will also require simplified drive circuitry, lower cost, and higher reliability than has been demonstrated in photonic and optoelectronic technologies. An example is found in two-dimensional point-to-point array interconnects for MCM stacking. These interconnects are based on high-efficiency Vertical Cavity Surface Emitting Lasers (VCSELs), Heterojunction Bipolar Transistor (HBT) photoreceivers, integrated micro-optics, and MCM-compatible packaging techniques. Individual channels have been demonstrated at 100 Mb/s, operating with a direct 3.3V CMOS electronic interface while using 45 mW of electrical power. These results demonstrate how optoelectronic device technologies can be optimized for low-power parallel link applications.

KEYWORDS: Packaging, Interconnect, Three-Dimensional Stacking, Multi-Chip Module, VCSEL, HBT Photoreceiver, Microlens, Binary Optic Lens

1. INTRODUCTION

Clock speed, word width, and number of I/O per chip are growing rapidly in digital microelectronics and especially for microprocessor chip technology. The Semiconductor Industry Association (SIA) road map, for example, projects up to 2000 interconnects per chip or module (at 100-200 Mb/s) by the year 2001.¹ Though this set of requirements represents the highest levels of performance that might be anticipated, specific interconnect throughput levels may be pushed by applications such as distributed computing, advanced image processing and pattern recognition, mass memory access, and display technology.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED ^{DLC}

As a result of these advances in technologies and applications, Multi-Chip Modules (MCMs) are being pushed to provide larger chip counts, higher densities of chips on a given module, and constantly increasing numbers and densities of signal Inputs and Outputs (I/O). Studies indicate that current MCM technology could support digital clock rates on electrical interconnections as high as 1 GHz for line lengths compatible with a 2 inch square module.² Thus, the basic length-bandwidth limitations of MCM electrical interconnects have not yet been reached. The ability to provide sufficient numbers and densities of high-frequency connections from module to board or module to module, however, may become limiting factors in the performance of MCM technologies. Module-to-module connection will be especially important to specialized imbedded computing applications where package size and circuit footprint become critical issues.³

Present two-dimensional MCM technology has done much to enable shorter electrical interconnect paths for reduced volume and higher performance. Cross-talk and inductance problems can be minimized for high densities of *permanent* electrical lines with metallurgical bonds, which currently provide the large connectivity required for microprocessor chips and advanced MCMs.⁴ A logical extension is stacking of MCMs, which would allow the advantages of compact packaging to be extended into a third dimension, thus further enabling interconnect-intensive applications by providing additional short path-length connections between chips.^{5,6} Various examples of the stacking approach use permanently-connected three-dimensional stacks of chips and modules. One method includes arrays of thinned chips, permanently edge-connected to a mother board.⁷ In another case, chips are attached to both sides of Ball-Grid Array (BGA) substrates that are stacked using interposer layers. Peripherally-arranged solder balls or conductive adhesives are then used to provide vertical connection between the substrates.⁸ Other alternatives use advanced wire bonding and bumping to realize stacked modules.⁹

In many cases, a stacked MCM architecture would benefit from the ability of layers to be easily separated. This allows each layer in a high-value stack to be separately constructed, tested, and reworked as needed. In addition, the full advantage of increased connectivity in a stack might only be achieved if two-dimensional (area) arrays of interconnects are used between layers. This would be especially beneficial in applications such as pipelined parallel processor architectures, where a central bus might not be used. In the pipeline architecture, each processor successively acts on its input and feeds the processed data directly into the next processor.¹⁰ This would allow maximum throughput in specialized computing applications such as signal processing and pattern recognition for Synthetic Aperture Radar (SAR) images.

Three-dimensional MCM approaches to pipeline processor architectures are presently limited by the lack of an area array of high-density, low-crosstalk, separable interconnects that can run at 100 MHz (or higher) clock speeds with acceptable signal integrity. This problem occurs because, unlike the metallurgical contacts afforded by permanent electrical interconnect arrays, separable interconnects depend on the pressure of the contact between layers. Possible alternatives using direct separable electrical connections include "fuzz-buttons"⁵ or dendrite connection pads.¹¹ The pitch of the connectors in these approaches is presently about 1 mm, and further size reduction may be limited by the need for shielding against cross-talk.¹² Small-pitch spring and plunger connectors are also presently limited to about 1 mm pitch due to the need for enough area and force to

accommodate a wiping or spring displacement.¹³⁻¹⁵ Temporary connections have been proposed and implemented for test and burn-in using wires suspended in elastomers¹⁶ or metal bumps on a polymer membrane¹⁷ which also require a wiping force. Another alternative for separable electrical interconnection is the use of suspended particles, which do not need a wiping force.¹⁸ Though anisotropic conductive adhesives have shown high-performance operation and small pitch¹⁹, scaling of separable particle-based interconnects to small pitch may be limited due to needed pressures, particle concentrations and resulting probabilities of short circuits and cross-talk. Capacitive coupling for separable interconnection has also been proposed²⁰, but would preclude Non-Return-to-Zero (NRZ) operation.

2. CONSIDERATIONS FOR PHOTONIC INTERCONNECTION OF STACKED MCMs

Advances in photonic device technologies have opened up many new possibilities for high-density two-dimensional arrays of sources and detectors that can be applied to MCM stacking. This trend has been paced by the development of Vertical Cavity Surface Emitting Lasers (VCSELs), which exhibit low-power operation, circularly-symmetric output beams, and can be easily fabricated in two-dimensional arrays directly on the surface of a semiconductor substrate.²¹ Such laser arrays are much more compatible with the needs of stacked MCMs than traditional edge-emitting laser devices, which are typically limited to one-dimensional arrays that emit an elliptical light pattern parallel to the plane of the wafer.

Similarly, optoelectronic photoreceiver technologies and circuits have been developed that are compatible with two-dimensional array approaches to optical interconnection. Though a wide variety of circuit and device options exist for such arrays²², the specific power and speed requirements of a given system will dictate device and circuit choices.

In addition to the advent of advanced photonic devices such as VCSELs, recent trends in photonic device packaging have also enabled the expanded use of photonics in MCM applications. The development of passive techniques for device-to-fiber and device-to-substrate alignment and mounting allows the photonic die to be treated much more like electronic chips in the MCM environment. In particular, the use of specially constructed alignment features with passive alignment by offset solder bumps²³ and/or robotic chip placement²⁴ have allowed edge-emitting laser die to be placed with micron to sub-micron accuracy with respect to optical fibers. Soldering techniques and flip-chip attachment are also important for placement of surface-normal components, including VCSELs.²⁵

To address the particular problem of vertical interconnection in MCM stacks, advanced photonic device and packaging technologies can be applied as in Fig. 1. Here, two-dimensional surface-normal photonic interconnect die that communicate between levels are placed on the MCM, thus forming a Z-Axis Photonic Interconnect (ZAPI). The photonic die are distributed among electronic signal processing and memory chips that are electrically interconnected on each level, such that all electrical line lengths are kept as short as possible. Arrays of photonic signals then pass vertically through optical via holes in the MCM substrate to provide interconnect paths between layers. It may also be

Figure 1. A photonic-interconnected MCM stack. Photonic interconnect chips communicate through via holes or at wavelengths transparent to the MCM substrate.

possible for the optical sources to operate at wavelengths where the substrate is transparent, thus eliminating the cooling and structural limitations that could be imposed by the need for via holes.²⁶

The ZAPI approach has a number of potential advantages over an area array of electrical interconnects for high-density, high-speed data transfer between layers in an MCM stack. The use of photonic paths can eliminate electrical parasitics. These parasitic resistances, inductances, and capacitances can be especially severe at the separable pads that would be required for high-density electrical interconnection of the stack. As pad size is reduced, the series inductance and resistance will tend to increase, thus causing additional high frequency signal loss and noise on the vertical interconnect channels in the MCM stack. A reduction in pitch will tend to increase capacitive coupling and thus produce more cross-talk between electrically-connected channels.

Because a ZAPI system will not require *physical* contact of each of the interconnect channels as separable electrical pads would, it is a high-density zero-insertion-force socket. The MCM stack would thus need less mechanical holding pressure than it would for large numbers of small-pitch electrical contacts, and high-speed performance would not be dependent on that pressure. In addition, optical signals may be tapped in order to test each of the MCM layers individually, thus enhancing yield. Another advantage is that the ZAPI approach will reduce the effect of ground bounce between the layers, thus allowing the photonics to act as a multi-channel high-performance optical isolator system.

3. A REPRESENTATIVE APPLICATION

The signal processing system that serves as a target application for the ZAPI implementation is based on processors using 3.3 V CMOS electronics. The stack will contain up to 2000 individual channels per layer because of the parallel pipeline architecture. Thus, there is no centralized data bus and the interconnect must operate at the processor clock speed for maximum performance. These processors will then transfer data at 80 to 100 Mb/s. The goal for electrical power consumption in the ZAPI is 20 W per layer (10 mW per channel), which represents a much lower value than for typical photonic data links. It is the low power requirement on this interconnect design that is the single most important feature driving the choice of photonic device technologies to be used.

Individual modules in the z-axis stack are to be based on one of several MCM technology choices. In the first option, they would be constructed as a "chips first" MCM technology such as General Electric's High Density Interconnect (HDI), where electrical interconnect layers are overlaid on top of the chips.⁴ As an alternative, the MCMs could be constructed "chips last" where the electronic and photonic die are flip-chip mounted onto a pre-patterned substrate (such as an MCM-D arrangement).

An important operational consideration for the z-axis interconnect system is that it convey data as accurately as an electrically-connected gate in a CMOS sequence. Because the parallel processors in the MCM stack are highly interconnected, timing margin and latency become important limiting factors in overall processor performance. The large number of parallel channels in the interconnect preclude the use of advanced error coding, multiplexing, and/or clock recovery circuits because of the additional complexity, power consumption, and latency that they imply. These considerations translate to a requirement that the photonic channels in the z-axis operate in parallel and that they carry synchronous NRZ data near the 100 Mb/s rate with rise and fall times of about 1.0 to 1.5 ns. These timing requirements, combined with the need for low-power, highly-parallel, low-crosstalk (nearly error-free) operation, call for a system based on photonic device technologies such as highly efficient Vertical Cavity Surface Emitting Lasers (VCSELs) and special photoreceiver circuits that feature low power consumption.

4. PHOTONIC APPROACHES TO HIGH DENSITY Z-AXIS INTERCONNECTION

The design of a low-power photonic interconnect stack was driven by the above requirements, in conjunction with the additional constraints implied by packaging compatibility. A number of possible approaches and wavelengths were considered for the baseline emitter design, including visible LEDs and/or VCSELs, 1300 nm LEDs, and 980 nm VCSELs. Packaging constraints were identified for each of the approaches, based on substrate transparency issues and output characteristics of the optical source. Expected power budgets were also developed, according to emitter device and receiver circuit power consumption characteristics and light collection efficiency in the optical path. In the designs that were considered, one goal was to balance electrical power consumption between the transmitter device and the photoreceiver circuit for maximum efficiency and

CMOS compatibility. In general, the more optical power that could be focused on the photoreceiver, the lower the amplification needed in the circuit, resulting in a lower overall receiver power consumption at a given bit rate. The two factors in the source device that strongly influenced the power budget were overall conversion efficiency from electric power to light, and the divergence of the output light beam to be focused on the photoreceiver. In both of these areas, VCSELs were found to be far superior to LEDs.

As demonstrated in an interconnect test module based on visible LEDs, even narrow-band resonant-cavity surface-emitting devices exhibited only a 2 percent power conversion efficiency and near-Lambertian beam profile when confined to small active output areas.²⁷ Thus, lenses with a Numerical Aperture (NA) of 0.4 (a full-angle acceptance cone of 48 degrees or f-number of 1.25) only collected 16 percent of the output light. This would lead to low overall efficiency in the power budget and high cross-talk, compared to a VCSEL-based system with 10% to 20% overall power conversion efficiency and an output divergence full angle of 12 to 20 degrees.

In addition to the power requirement, the most basic criteria influencing the choice of approach were the packaging constraints. The most restrictive of these was that the photonic devices must be compatible with a "chips first" packaging approach such as General Electric's High Density Interconnect (HDI).⁴ In this realization, the chips are all mounted in wells (as schematically shown in Fig. 1). The electrical interconnect is written adaptively on an overlay layer that is flush with the top surface of the chips. In such a configuration, it would become very difficult to mix the mounting mode such that some chips are mounted up (in the normal HDI fashion) while others are mounted down (thus needing electrical interconnects in the well of the MCM substrate). Mixed mounting, however, would help greatly with the need to communicate both to the layer above and the layer below in the MCM stack. With mixed mounting, a back-emitting VCSEL could be placed with its epitaxial layer on the top surface to communicate down or could be flip-chip mounted to communicate up.

The combination of chip mounting constraints and bi-directional layer-to-layer communication further influences the choice of photonic device technologies to combinations of wavelengths and materials that allow light to pass through the photonic device substrates. The transparent device substrates then become an optical platform for bi-directional communication as in the advanced stack cross-section of Fig 2. Here, the outputs of the transmitter and receiver devices are reflected and/or focused (through their transparent substrates) to achieve bi-directional communications.

For the VCSEL communicating downward in Fig. 2 or the photoreceiver receiving signals from below, the back of the photonic device substrate would simply contain a collimating lens structure. This could either be a diffractive or refractive optical lens, as will be discussed. In the case of a VCSEL communicating upward, a reflective grating or off-axis lens element must be used. This optical element would provide an angular offset of the VCSEL light beam, which would then be reflected back to the top surface of the substrate. Here, transmissive optics would deflect the beam to surface-normal propagation and collimate the output. In the case of a photoreceiver detecting a signal from above, the light would first pass through a small photodetector with a low-shadowing contact area. That portion of the beam overfilling the detector would then be

Figure 2. Cross-section of an advanced "chips-first" architecture for z-axis photonic MCM interconnects. Reflective and transmissive optics are integrated into the photonic device substrates to effect upward or downward communications.

reflected and re-focused onto the back of the detector by an optical element on the back of the substrate. This would be analogous to a Cassegrain telescope arrangement, and would allow the use of small area photodetectors with their corresponding advantages of low electrical capacitance for high speed and low power consumption.

5. VCSEL TRANSMITTER ARRAYS FOR LOW-POWER INTERCONNECTS

Because of their low-divergence beam characteristics, VCSEL-based transmitters enable optical links with high efficiency, low cross-talk, and (with integrated lenses) bi-directional transmission. In addition, VCSEL devices have shown high (14 GHz) small-signal modulation bandwidths.²⁸ It follows that the devices should easily meet the speed requirements of parallel links and have a substantial capacity for scaling to higher bit rates. VCSELs can also be designed for minimal sensitivity to ambient temperature excursions.²⁹ VCSEL arrays have been constructed that are compatible with the interconnect of Fig. 2. These are 980 nm back-emitting InGaAs devices on GaAs substrates. Two different device types are discussed.

In the first case, the devices are gain-guided, with the active region defined by a surrounding ion implant as in Fig. 3a. Here the confining region is 20 μm in diameter. Device die were constructed as 4 x 4 arrays of lasers on 500 μm centers as in Fig. 3b. The large contact pads around each laser were designed to assure continuity between the laser contact and the metal interconnect. DC light-current-voltage characteristics of a

typical device appear in Fig. 4a. Here, the power conversion efficiency is near 4 percent at an output power of 1.0 mW. Note that the drive voltage at this operating point is 2 volts, while the drive current at 1mW is 12 mA, thus enabling compatibility with direct 3.3V CMOS drivers.

Figure 3. Structure and layout for the implanted VCSEL array.

Time delays were measured for low duty-cycle operation in order to determine the effects of thermal lensing in these gain-guided devices. Gain-guided VCSELs operate with the aid of a thermally-induced index change that acts as a lens within the current flow path of the VCSEL structure.³⁰ This is not a major concern for high speed modulation in systems where the VCSEL can be pre-biased at its lasing threshold, thus pre-establishing the thermal lens. The effect is very important, however, for operation in a low-power interconnect with direct CMOS drive. Delay was found to be drive-level dependent for the 20 μm diameter device results shown in Fig. 4b. At drive levels of 12 mA (2x threshold with an optical output of 1 mW), the delay was as high as 2.5 ns. At higher drive values of 25 mA, the turn-on delay dropped asymptotically to 1.0 ns. For smaller laser diameters of 10 μm , the lasing characteristics became very dependent on the thermal lens. This was probably due to the increased effects of diffractive light loss attributed to evanescent field interaction with the un-pumped implant region. In this case, the laser threshold current became very dependent on the pulse duration and duty cycle that form a thermal history for the VCSEL. The resulting turn-on delays could be as high as several hundred nanoseconds, thus making the small-diameter implanted devices unsuitable for this low-power, CMOS-driven interconnect. Similar long-delay effects have been observed in implanted edge-emitting lasers.³¹

a. DC voltage and light output as a function of input current.

b. Pulsed response for the implanted VCSEL. Upper curves are input diode voltage. Lower curves are measured laser outputs. Note the delay on the output pulses.

Figure 4. DC and pulsed characteristics of implanted VCSEL devices, showing delay associated with operation at zero prebias.

Etched-post devices from the same epitaxially-grown wafer were also tested for compatibility with the ZAPI. Here, the current and light are confined by the etched post of Fig. 5a. Again, the lasers were arranged on $500\ \mu\text{m}$ centers as in Fig. 5b. In this pattern, the large bond pad was eliminated in favor of a small-area ring. Note that a common ground bus was used. This ground area was connected by ohmic contacting to the highly-doped n-type GaAs substrate, and was patterned to run close to each of the sixteen laser sites. Six pads provide for parallel external ground connections. This is particularly important for array operation, since any significant resistance in the path of the common ground will cause cross-talk, resulting from de-bias of the lasers as more are operated.

Figure 5. Structure and layout for the etched-post VCSEL array.

A typical light-current-voltage characteristic appears in Fig. 6a for a device diameter of 10 μm . The threshold current was slightly over 2 mA and the desired output of 1 mW was attained at 7 mA with a 1.8 V drop (a power efficiency of 8 percent). Because this device was not dependent on thermal lensing, the turn-on delays were typically less than 1 ns, as indicated by Fig. 6b. At a drive of 17 mA, the turn-on delay was about 200 ps.

Similar, more optimized gain-guided VCSELs have exhibited up to 21% efficiency.³² More recent devices using oxidized layers for lateral confinement have shown efficiencies higher than 50%.³³ Because the oxidized devices are index-guided,^{34,35} they should show small delays in turn-on, as do the etched post devices described above. Compatible oxidized VCSEL devices are now being fabricated for demonstration in the ZAPI.

6. HBT PHOTORECEIVER ARRAYS FOR LOW-POWER INTERCONNECTION

As indicated in the interconnect cross-section of Fig. 2, optical signals from the VCSEL transmitter channels are focused onto integrated photoreceiver arrays. The photoreceiver designs chosen for use in this interconnect were based on InGaAs/ InP Heterojunction Bipolar Transistor (HBT) circuits, constructed on InP substrates. This allows for efficient light collection in an integrated p-i-n photodetector that is fabricated from the base-collector junction of the vertical Npn HBT transistor structure.³⁶

a. DC voltage and light output as a function of input current.

b. Pulsed response for the etched post VCSEL. Upper curves are input diode voltage. Lower curves are measured laser outputs. Note the small delay on the output pulses.

Figure 6. DC and pulsed characteristics of etched-post VCSEL devices, showing small delays associated with operation at zero pre-bias.

Note in Fig. 2 that the receiver substrates must be transparent to allow lens integration and subsequent back-side detection in the p-i-n photodiode. The photoreceiver circuits are constructed on semi-insulating InP:Fe material. Though the basic InP material exhibits a sharp absorption edge at 960 nm, the iron doping can cause a tail effect that may induce significant absorption at 980 nm. This effect becomes especially pronounced at high levels of iron doping, and the levels can vary significantly over a single boule or from the outside of the wafer to the inside.^{37,38} Iron doping levels are often not tightly controlled by the material manufacturers, since they usually only need to be high enough to guarantee a certain level of resistivity in the substrate.

Because of these concerns for substrate transparency, InP substrate absorptions were measured over the wavelength range of interest for several temperatures. Absorption coefficients at various temperatures appear in Fig. 7. These, and measurements on substrates from other vendors, indicate that at our chosen substrate thickness of 350 μm , the room-temperature optical absorption in the substrate at 980 nm could range from 2%

Figure 7. Typical absorption curves for a Fe-doped semi-insulating InP substrate with thickness of 368 μm . Absorption tails varied greatly with sample.

to 20% . At higher temperatures such as 75 C, the combination of band-edge shift and tail effect could cause the absorption to range from 20% to 40%. These levels of attenuation could significantly degrade the power efficiency of the optical link. Thus, substrate pre-selection or die thinning and remounting (as on GaAs) may become important factors for efficient ZAPI realization. It may also be possible to build high-efficiency VCSELs that operate at wavelengths near 1000 nm.

The photoreceiver circuit schematic appears in Fig. 8. The circuit is optimized for large-signal operation with low power consumption and a relatively large input optical power. The receiver is designed for 3V CMOS-compatible outputs with a single 3.3 V power supply. It operates with minimal (<1 ns) turn-on delay when peak optical input powers are near 1 mW. The non-linear (large-signal) circuit response, which enhances data fidelity, requires the use of high speed transistors. The circuit was modeled using SPICE and was shown to switch (though with increased delay) at input photocurrents as low as 100 μA . Transient analysis for larger photocurrents (about 400 μA or 1 mW of optical power at a responsivity of 0.4 A/W) predicted <1.5 ns rise and fall times driving a typical CMOS 4 pF input load capacitance.

Figure 8. Circuit schematic for the HBT photoreceiver.

The array of photoreceiver circuits was constructed as in Fig. 9.³⁶ Note the round 50 μm diameter photodetector region in each circuit. Because of the lensed back-side illumination in each detector, a reflective contact was able to be placed on the top, and the detector efficiency was thus increased due to double-pass absorption. The detectors and associated circuits were arrayed into a 4 x 4 pattern on 500 μm centers to match the laser arrays of Figs. 3b and 4b. Note also in the pattern of Fig. 9 that separate power, ground and signal pads were provided for each circuit. DC operation matched that predicted, showing a switching threshold with an optical power input of 250 μW (100 μA of photocurrent), and a high-level output of 2.7 V.

Figure 9. Circuit layout for 4 x 4 HBT receiver array with circular 50 μm diameter detectors on 0.5 mm centers.

The circuits were tested at speed, using 980 nm laser stimulation. In this test, the laser was operating with a 200 MHz square wave input to simulate a 400 Mb/s data stream. The resulting output is shown in Fig. 10.³⁶ Here the photoreceiver tracks the rise and fall times of the pulses to sub-ns accuracy. The voltage swing is from 0.8 to 2.6 V, which is compatible with the drives needed for 3.3V CMOS. The optical input peak power was approximately 600 μW . Testing speed was limited by the frequency response of the available 980 nm laser source, as indicated by the long tail in the photoreceiver response of Fig. 10. This tail was also observed on fast photodiodes used to characterize the laser source.

Figure 10. Photoreceiver response driving a 4 pF load at an equivalent 400 Mb/s. The voltage swings between 0.8 and 2.6 volts. The horizontal scale is 2 ns/div and the vertical scale is 0.5 V/div. The slow decay of the low-level state is an artifact of the laser output.

7. OPTICAL DESIGN AND LENSES FOR A BOARD-LEVEL DEMONSTRATION

In order to demonstrate the basic elements of the ZAPI, a board-level laboratory prototype was constructed as in the cross-sectional view of Fig 11. This simplified two-layer test arrangement was used to test basic link characteristics. Unlike the more advanced stack of Fig. 2, it features flip-chip mounted laser and photoreceiver die that face each other over an adjustable distance. It thus does not demonstrate the reflective binary optical elements described in Fig. 2, but allows the basic collimation and re-focus lens technologies to be tested.

To accommodate VCSEL divergence, possible stacking misalignment, and re-focus of light to a $50\text{ }\mu\text{m}$ detector diameter, lenses are placed on the back sides of the VCSEL and HBT device substrates as in Fig. 11. An optical design for this demonstration link appears in Fig. 12. Here, the GaAs and InP device substrates have a thickness of $350\text{ }\mu\text{m}$ and are separated by up to $1200\text{ }\mu\text{m}$ to accommodate various possible thicknesses of the silicon MCM substrates that will eventually form the platform for the stack. The first lens, formed on the back of the VCSEL substrate, limits the optical beam diameter at the receiver substrate. Based on simple Gaussian beam approximations, this design accommodates lens-to-device misalignments of up to $3\text{ }\mu\text{m}$ and level-to-level misalignments of 5 to $10\text{ }\mu\text{m}$ by the use of a receiver lens that is larger than the diameter of the optical beam. Shorter layer separation distances and tighter alignment tolerances allow the receiver lens diameter to be reduced, thus increasing its f-number and making for easier receiver lens fabrication. In actual implementation, the receiver lens diameter was reduced to $100\text{ }\mu\text{m}$ and the focal length to $150\text{ }\mu\text{m}$, which was allowable for an $850\text{ }\mu\text{m}$ separation between source and receiver die. The alignment tolerance between VCSEL and

Figure 11. Breadboard test cross section. The lensed VCSEL and HBT can be brought into close proximity for alignment testing. The laser and HBT arrays are directly interfaced to 3.3V CMOS circuits.

lens is particularly important. Every $1\text{ }\mu\text{m}$ of misalignment between the VCSEL and the lens causes a $10\text{ }\mu\text{m}$ beam offset at the design distance.

In the demonstration board-level link of Fig 11, lenses were realized using both refractive and diffractive methods. A combination of both lens types will allow for maximum design latitude and efficiency in the advanced link cross-section of Fig. 2. In such an architecture, refractive lenses and/or gratings (or offset binary lenses) would be used on the VCSEL substrates to turn and collimate the beam, thus allowing upward transmission as described.

Figure 12. Optical design using collimating and refocus lenses for a simple stacking configuration.

a. Circular symmetry for collimation.

b. Asymmetric lens for 20 degree off-axis steering.

Figure 13. Binary lenses with 8 phase levels for VCSEL beam collimation and steering. The focal length is $110\text{ }\mu\text{m}$ at F/1.4.

The diffractive or binary optic lenses are etched into the back of the VCSEL substrate with an eight phase-level pattern as in Fig. 13. Alignment marks provided for front-to-back registration using a near-infrared imaging system and optical lithography. These lenses were then fabricated using electron-beam lithography and reactive ion beam etching.³⁹ These lenses, which were designed to reduce the beam diameter at 850 μm , were integrated with back-emitting implanted VCSEL arrays. Other work with binary lens integration on a VCSEL substrate has shown focusing at 105 μm .⁴⁰ For the etched-post VCSELs and for the larger (20 μm) implanted VCSELs needed for small turn-on delay, the beam shapes show multiple lateral modes. This condition resulted in lower efficiency in the binary lens and spreading of the output pattern in the quasi-collimated beam.

Figure 14. A comparison of light output with and without a binary lens on a bottom emitting VCSEL. The 100 μm diameter pinhole used for measurement was placed 800 μm from the laser.

An effective measure of lens efficiency was made by comparing the output VCSEL power (measured with a large-area detector) with that detected behind a 100 μm pinhole placed at the design distance of 850 μm from the VCSEL substrate. A typical result for an implanted VCSEL appears in Fig. 14a, where the output is somewhat reduced, showing an 80% efficiency for operation through the pinhole. Compare this to the result of Fig. 14b, which shows operation of a laser with no lens applied. Here, measurements were again made with and without the 100 μm pinhole. Note that the laser output through the pinhole ceases to rise even with increasing current injection. This is due to the additional beam spreading associated with the increasing high-order transverse mode effects of the laser at higher currents.³⁹

For the implanted VCSELs, the measured value of lens efficiency varied between 70% and 100%. This large variation in lens effectiveness was attributed to misalignment in fabrication, scratches on the substrate, widely varying mode profiles for the lasers, and shifts in wavelength over the VCSEL array. For the etched-post devices, the measured efficiency ranged between 30% and 60%. The lower efficiency of the etch-post devices

was due to a combination of processing difficulties and the higher mode orders associated with the strong optical guiding in the etched posts.

Refractive lenses were constructed on the InP photoreceiver die by patterning and subsequent thermal flowing of polydimethylgluterimide (PMGI) photoresist at 290° C. A PMGI lens example appears as in Fig. 15. The actual receiver lenses had a focal length

Figure 15. Refractive lens formed in polymer material on an etched pedestal. Similar lenses with 100 μm diameter were used to focus light on the detectors in the photoreceiver array.

of 150 μm at a lens diameter of 100 μm (as previously described). These refractive lenses have also been integrated onto the back-emitting VCSELs, there collimating the beam to a divergence angle of 1 degree.⁴¹ Similar structures have been transfer-etched into GaAs for shorter focal length lenses of a given diameter.⁴² Refractive lenses may provide advantages over the binary lenses due to their ability to handle a wide range of wavelengths without loss of efficiency, their relative ease of fabrication, and their ability to be anti-reflection coated. Refractive lenses are less able to effect the angular offsets of Fig. 2, however, and their non-planar profiles can make handling of die difficult. In the case of the photoreceivers, the use of the refractive lenses eliminated the need for development of an anisotropic high-precision etch in InP (which would be needed for binary optics). InP is, however, amenable to various methods of refractive lens fabrication as previously shown.^{43,44}

8. BOARD-LEVEL DEMONSTRATION RESULTS

The board-level cross-section of Fig. 11 was realized using the lensed VCSELs and photoreceivers described above. These boards were designed and fabricated to allow for simultaneous testing of up to 16 channels of optically-exchanged data. The purposes of these tests were to 1) demonstrate and test performance with direct CMOS drive of the VCSELs and CMOS interface to their photoreceiver outputs, and 2) to explore the range of stack alignment tolerance afforded by the optical design of Fig. 12.

The laser and photoreceiver die were flip-chip mounted on silicon submounts to allow insertion into Leadless Chip Carriers (LCCs) for interchange within the board-level tests. The VCSEL die were mounted on the silicon subcarrier as in Fig. 16a. Here, the 350 μm thick GaAs VCSEL substrate was flip-chip bonded using Indium-alloy solder paste and the entire carrier was mounted in the well of the LCC. The 4 x 4 diffractive lens array with 0.5 mm pitch is visible on the back side of the die in Fig. 16a. Wire bonds connect the carrier to the pads of the LCC. As discussed, multiple ground connections were provided on the LCC in order to minimize common-path de-bias and signal cross-talk effects. For the same reason, ground lines on the submount were constructed with

additional width. The capacitance for each packaged laser at zero bias was 28 pF for the implanted devices of Fig. 3.

Figure 16. Transmitter and receiver die flip-chip bonded onto silicon submounts in leadless chip carrier packages.

The photoreceiver was mounted as in the top view of Fig. 16b. In this case, the InP substrate was thermosonically bonded onto the silicon submount using Au/Pd bumps. While thermosonic bonding does not allow for the self-alignment that make solder bumps advantageous²⁵, it is a simpler process, and might be acceptable for die placement on a MCM within the alignment tolerances afforded by advanced pick-and-place technology. The refractive lens array is visible in Fig. 16b. Separate V_{cc} lines for each channel on the device die were carried to the LCC. One ground line was provided for every two circuit channels. As shown in Fig. 16b, a 140 pF capacitor was placed between the power and ground for every two circuits.

Figure 17. Board-level test circuit schematic.

Each of the sixteen channels on the circuit board pair realized the circuit in the schematic of Fig. 17. Input from a word or signal generator was applied via a 50 Ohm line into the

SMA on the input side of the link. The signal then traveled down approximately 4 inches of microstrip line to a 50 Ohm termination resistor at the input to a CMOS driver. The drivers used were Integrated Devices 74FCT163244 CMOS buffer/driver circuits designed to operate at 3.3V. In order to keep untuned line lengths short and equidistant as possible, only eight of the sixteen possible channels were used on each of two driver ICs on the VCSEL board. A series resistor was used between each driver output and its corresponding VCSEL. The value of this resistor could be selected to adjust VCSEL input current and corresponding light output. The lensed output and corresponding photoreceiver response was fed into an identical CMOS buffer/driver on the second board. The CMOS output was transferred, via 50 Ohm lines, to a digital oscilloscope for signal analysis.

a. Board level test results: eye diagram with a 3.3 V power supply sourcing the VCSEL board.

b. Eye diagram for a 2.7 V supply voltage on the VCSEL board. Note degradation of timing margin and jitter.

Figure 18. Test results for the board level demonstration interconnect, showing operation at a supply level of 3.3 V (0.9 mW peak laser output) and 2.7 V (0.5 mW peak laser output).

The eye diagram of Fig. 18a was generated for a single channel on the test board pair, operated at the designed supply voltage of 3.3V with optimal laser-to-receiver alignment. The lasers used were implanted devices as in Fig. 3. A 24 ohm series resistor was used in as in Fig. 17, giving a nominal laser drive current of 20 mA. This level of drive corresponds to a peak laser output (into the 100 μ m lens diameter) of 0.9 mW, leading to a total demonstrated power consumption of 45 mW (66 mW peak for the laser, multiplied by a 50% average duty cycle, and 12 mW for the receiver). Here, 100 Mb/s NRZ pseudo-random data was used to simulate operation as in the multi-processor stack application. Note that the eye is very open. Timing margin is approximately 9 ns and jitter is on the order of 0.25 ns. The channel was connected to a bit error rate tester and was operated for several days. Greater than 10^{13} bits were passed through the link without error. In order to test the effects of laser power, the source voltage for the transmitter circuit was dropped to 2.7 V. In this case, the laser operated at a nominal 16 mA peak for an output of 0.53

mW. This lower laser power reduced the timing margin to 5 ns and increased jitter to 1 ns as in Fig. 18b.

In addition to the eye diagrams of Fig. 18, the timing curves of Fig. 19 were also generated using the digital oscilloscope. These allowed a more detailed study of the individual elements of the link and a measurement of delays within the circuit. The measurements were

obtained by using a single, 100 kOhm active probe on a 3GHz oscilloscope sampling head and storing the voltage traces from the measurements taken at various points within the circuit of Fig. 17. The design value of 3.3V for transmitter and receiver was again used as the supply voltage. Turn-on delays, as shown in the figure, indicated that the total delay from buffer-to-buffer for this link was 6.8 ns. Of this total, 3.3 ns of the delay was due to the laser and photoreceiver. The input and output CMOS buffer/drivers accounted for 1.7 and 1.5 ns, respectively. In order to evaluate the relative proportions of the 3.3 ns delay attributable to the laser and the photoreceiver, the output from the laser drive board was measured directly, using a fast photodiode. The result of that measurement appears in Fig. 20. Here, 1.7 ns of the total delay is due to the time needed for the drive signal to reach

Figure 19. Timing diagram for various points on the test board arrangement. Traces were taken using an active, high-impedance probe, and show the origin of delays within the interconnect.

Figure 20. Laser timing response when driven in the test board circuit. Note the combined effects of rise time and turn-on delay.

the laser threshold point and 1.0 ns is due to laser turn-on delay. Recall that these results are for the implanted laser, which shows more turn-on delay than the etch post laser. Further board-level testing will be done to verify the advantages of the etch-post laser technology in this application.

As indicated in Fig. 11, the relative positions of the transmitter and receiver boards could be manipulated to test alignment tolerances. This was done for several different separation distances as in Fig. 21. The quantity used to evaluate alignment was pulse width from the CMOS buffer on the receiver board. As indicated in Fig. 18b, the pulse width from the CMOS would fall as the laser power was decreased, thus reducing the effective timing margin for the link. This was due to a combination of effects, including increased delay in the photoreceivers for lower input photocurrents and reduced output amplitudes from the photoreceiver into the CMOS buffer. The latter effect can be understood in terms of output curve D in Fig. 19. Here, the output must reach a certain level to cause triggering in the CMOS output buffer. As input laser power is reduced, the overall peak level of the output drops. The threshold point then moves farther out along curve D, eventually transferring from the steep portion to the more gradually rising portion (which would also increase jitter, as in Fig. 18b). At lower powers, the entire output will fall below the threshold point of the CMOS, and the link will cease to function. This same effect occurs as the receiver and laser boards are offset in Fig. 21. Note that the collimation and re-focus of the optical system allows for operation with greater than 9 ns of width over a $\pm 25 \mu\text{m}$ alignment tolerance. As the misalignment becomes worse, the timing margin and jitter would degrade as in Fig. 18b. This experiment shows that timing margins can be maintained within reasonable stacking tolerance values for the z-axis MCM interconnect application.

Figure 21. Effects of misalignment on output pulse width and link operation for three laser-to-photoreceiver separation distances.

9. MCM PACKAGING

In "chips first" modules such as General Electric's HDI, wells are cut or etched into the MCM substrate material.⁴ In this arrangement, materials such as alumina or aluminum nitride are typically used for the MCM substrate, though silicon may also be used.^{4,45} In a stacked arrangement, silicon is a likely choice for the MCM substrate material, since it exhibits reasonably good thermal properties compared to other candidate substrate materials. The thermal conductivity of silicon is $150 \text{ W/(m}\cdot\text{K)}$ versus $170 \text{ W/(m}\cdot\text{K)}$ for aluminum nitride and $17 \text{ W/(m}\cdot\text{K)}$ for alumina.⁴⁶ Also, silicon is inexpensive and easily patterned using lithography, deposition, and etching.

The effectiveness of v-groove etching in silicon for alignment of vertical photonic communication channels was demonstrated for a simple LED-based stack as in Fig. 22. A 4×4 CMOS photodetector array was attached to a silicon submount as in Fig. 22a. Front and back patterning was used with anisotropic KOH etches to fabricate v-grooves on both sides of self-aligning spacers as shown.²⁷ Precision glass rods (such as optical fibers) were used to key the spacers to the submount layer and a 4×4 microlens array was mounted with the CMOS photodetectors at the lens foci. A corresponding submount, spacer set, and lens array was fabricated for an LED, and the two submounts were stacked (along with two sets of externally-mounted microlenses) into a PGA package for testing as in Fig. 22b. This stack was able to maintain needed alignment at the expected optical source power for operation with a pair of 0.3 mm collimating and focusing lenses, thus indicating that the stacking accuracy will be within that needed for the ZAPI application.

a. Partially-assembled Si submount.

b. Assembled demonstration stack.

Figure 22. Demonstration and test of v-groove alignment features for stacking of MCM layers. Externally-lensed LEDs and photodetectors were optically connected in the two-layer stack.

As discussed above, requirements on device power, interconnect speed, and package operating temperature drove this ZAPI design to 980 nm. Because this is a wavelength at which silicon is opaque, it follows that the MCM substrates must have optical via holes as shown in Fig. 2. In the silicon MCM substrates, these via holes can be constructed by the use of high-aspect-ratio laser-drilled holes.⁴⁷ Fig. 23

Figure 23. Needed high-aspect ratio vias holes can be laser-drilled into the silicon MCM substrate.

shows the cross-section of 2 mil (50 μm) holes that have been laser-drilled into a standard 25 mil silicon substrate. This is well in excess of the aspect ratio required for the 150 μm aperture in the optical design of Fig. 12. Laser drilling, when combined with other advanced packaging techniques such as stack alignment features and self-aligning solder, will enable the advanced stack of Fig. 2 to be realized. Laser drilling of large arrays, however, has disadvantages in that the structural integrity and thermal capability of the MCM could be compromised. It would then be desirable to use long wavelength (>1200 nm) sources with silicon MCMs (though such sources are not presently available in high-efficiency VCSEL form).

10. A MODULAR MCM DEMONSTRATION STACK

Prior to a realization involving full multi-level stacking, the packaging techniques described above will be combined with the device technologies and lens designs previously discussed to realize bi-directional communication in the two-layer MCM stack of Fig. 24. This stack is a follow-on to the board-level test arrangement of Fig. 11. As shown in the cross-section of Fig. 24a, the photonic die will be mounted in a mixed mode, where the electrical contacts are made on the bottom side of the die on the first layer of the stack. Here, the laser and photoreceiver would thus be flip-chip mounted using bump bonds for electrical connection as in Fig. 16. On the second layer of the stack, the bump bonds would be applied to back of the die and would thus only be used for mechanical alignment as in Fig. 22. Electrical connections would then be made by wire bonding. Note the use of laser drilling and etched stack alignment features.

The laser-drilled silicon MCM substrates that form the demonstration stack cross-section of Fig. 24a are to be assembled for testing as in the exploded view of Fig. 24b. Here, the CMOS drivers, buffers, decoupling capacitors, and photonic devices will be assembled onto the two-layer stack, which will fit into large cavity PGA or LCC package for operation in a standard digital tester, thus exercising all the elements needed for the advanced stack of Fig 2 (with the exception of the reflective optical elements, and

a. Cross-section of the two-layer MCM stack.

b. Exploded view, showing the VCSELs, photoreceivers, and drive electronics.

Figure 24. A prototype two-layer ZAPI MCM stack to demonstrate and test concepts using VCSELs, HBT photoreceivers, integrated lenses, and passive alignment methods.

handling of the cumulative thermal effects that would occur for a high density of stacked processor die).

11. APPLICATION TO OTHER INTERCONNECTS

The same low-power device technologies, direct electronic interface, and precision packaging approaches that are being developed for this system of "point-to-point" z-axis photonic links could be applied to a variety of advanced interconnect approaches. The low power, direct drive VCSELs and photoreceivers could be used in holographically-interconnected MCM substrate arrangements that have been previously realized using edge-emitting lasers and turning mirrors.⁴⁸ The integrated lenses demonstrated with these device technologies could also be used to pre-collimate and re-focus beams to ease constraints on board-level interconnections or for optical switching arrays.⁴⁹ The integrated lenses and direct electronic interface to be implemented in the two-layer stack of Fig. 24 are also applicable to fiber-array links in that the lenses may allow easier coupling and smaller photodetector areas, while the direct CMOS interface may allow for easier insertion into present electronic systems. If efficient VCSELs can be monolithically integrated with HBT and phototransistor technology, then a variety of programmable optical functions and resulting architectures may be realized.⁵⁰ Integrated laser drivers may also allow more efficient interfaces to be implemented for CMOS or other high speed electronic technologies. In addition, the use of more sensitive photoreceivers with the high efficiency VCSELs and MCM-based packaging describe here may further enable fan-out for 3-dimensional optically-interconnected computers. Such approaches make use of the full parallelism and global addressing afforded by photonics⁵¹, thus providing greater functionality than the point-to-point link demonstrated here.

12. CONCLUSION

This photonic interconnection link technology is compatible with the system-level requirements of stacked MCMs for distributed signal processing. It can be optimized to meet stringent power budgets, and could be expanded to produce large numbers of synchronous, uncoded, parallel photonic channels operating at the processor clock speed. The interconnect could thus solve the problem of providing separable, high density, two-dimensional interconnects for stacking of MCMs.

Transmitters are constructed from two-dimensional arrays of high-efficiency VCSELs that are driven directly by CMOS, thus minimizing external interface circuitry and electrical power. For this application, the VCSELs are designed to operate at a wavelength of 980 nm, which presently offers maximum performance and highest efficiency. The GaAs VCSEL die are then also transparent to the light, thus allowing micro-optics to be integrated into the device substrates for beam collimation and/or redirection. The VCSEL devices have been demonstrated to operate at 100 Mb/s without pre-bias. Turn-on delays may be optimized by the use of index-guided devices.

Photonic data channels defined by the transmitter die and lenses are completed by the use of corresponding monolithically-integrated InGaAs/InP HBT photoreceivers, built on InP substrates. This choice of materials allows for efficient absorption of the 980 nm photons in a p-i-n photodetector structure that is vertically integrated into the based-

collector junction of the HBT layers. Like the VCSEL, this photoreceiver design makes use of substrate transparency, though iron doping in the semi-insulating InP may cause some optical losses. Collection micro-optics are integrated into the device substrates to focus the light beam and thus allow for small photodetector areas without excessive loss of misalignment tolerance. The photoreceiver circuits are optimized for low power consumption, and have been demonstrated to operate in excess of 200 Mb/s. Digital data fidelity is achieved by matching the large optical power provided by the VCSEL-based transmitter with the large-signal saturating receiver allowed by the very high speed HBTs.

The VCSEL transmitter and HBT photoreceiver have been assembled into a board-level test station to demonstrate direct CMOS interface to the laser and photoreceiver, lens integration, and alignment tolerance. Compatibility with MCM-based packaging is also important to the use of photonic interconnections. This z-axis interconnection approach will make use of emerging packaging techniques such as solder-bump bonding for alignment, laser drilling of via holes to allow for inter-layer communication within the stack, and etched features for stacking alignment. Since the optical lenses are integrated into the photonic die, assembly and alignment is greatly simplified, thus helping to make the photonics more compatible with emerging MCM package techniques. All of these advances in device, package, and interconnect technology will help to make photonics more attractive for use in a variety of interconnect applications, ranging from free-space data architectures to optical fiber-based data communication links.

Acknowledgments

The authors gratefully acknowledge the technical support of Denise R. Tibbets-Russell, Melissa A. Cavaliere, Glen Knauss, Terry Hardin, Florante Cajas, and Tony Carter. Contributions were also made by the related work of Tu Du, Joel Wendt, Alan Vawter, and Ben Rose. This paper is dedicated to the memory of W. Jeffrey Meyer for his leadership in Sandia's CCST. All work was performed for the United States Department of Energy under Contract DE-AC04-94AL85000

References

1. "National Technology Road Map for Semiconductors", From the SIA Semiconductor Workshop- Working Group Reports and Workshop Conclusions, Published by the Semiconductor Industry Association, Santa Clara, CA 1993.
2. T. R. Haller, B. S. Whitmore, P. J. Zabinski, and B. K. Gilbert, "High Frequency Performance of GE High Density Interconnect Modules," *IEEE Trans. on Components, Hybrids, and Manufacturing Tech.*, vol. 16, no. 1, pp. 21-27, February 1993.
3. B. Gilbert, R. Thompson, G. Fokken, W. McNeff, J. Prentice, G.-W. Pan, "Advanced Multichip Module Packaging and Interconnect Issues for GaAs Signal Processors Operating Above 1 GHz Clock Rates," *Proceedings of the SPIE*, vol. 1390, pp. 235-248, 1990.
4. M. Alder, "GE High Density Interconnect: A Solution to the System Interconnect Problem", *Proceedings of the SPIE*, vol. 1390, pp. 504-508, 1990.
5. T. J. Moravec, R. C. Eden, and D. A. Schaefer, "The Use of Diamond Substrates for Implementing 3-D MCMs," *Proceedings of the SPIE*, vol. 1986, pp. 86-91, 1993.
6. J. L. Arnold, "3D Silicon Technology," Presented at the 1995 IEEE Workshop on Interconnects within High Speed Digital Systems, Santa Fe, NM, May 16, 1995.

7. D. E. Ludwig and J. L. Arnold, "Unique 3 Dimensional Architecture for High Resolution Imaging," *Proceedings of the SPIE*, vol. 2226, pp. 219-226, 1994.
8. M. Mita, T. Kumakura, "Advanced TAB/BGA Multi-Chip Stacked Module for High Density LSI Packages," *Proc. of the 1994 IEEE Multi-Chip Module Conference MCMC-94*, pp. 68-76, March, 1994.
9. C. G. Massit, G. C. Nicolas, "High Performance 3D MCM Using Silicon Microtechnologies," *Proc. of the 45th Electronic Components and Technology Conference*, pp. 641-644, May, 1995.
10. H. W. Lawson, *Parallel Processing in Industrial Real-Time Applications*, Prentice-Hall, Englewood Cliffs, N.J., p. 72, 1992.
11. J. Campbell and J. H. A. Knight, "A High Density Edge Connector," *Proc. of the 45th Electronic Components and Technology Conference*, pp. 140-143, 1995.
12. G. Pan, X. Zhu, and B. Gilbert, "A Quasi-Static Analysis of Fuzz Button Interconnects," *Proc. of the 1993 IEEE Multi-Chip Module Conference MCMC-93*, pp. 85-91, 1993.
13. K. Wright, R. Rothenberger, and R. Mroczkowski, "A New Contact/Connector Design for High-Density, High-Performance Connectors," *Proc. of the 45th Electronic Components and Technology Conference*, pp. 134-139, 1995.
14. D. G. Grabbe, R. J. Pryputniewicz, H. Merkelo, "The AMPSTAR Connector Project," *Proceedings of the SPIE*, vol. 1986, pp. 347-357, 1993.
15. J. R. Trent and G. L. Westbrook, "Fine Pitch Pad Array Carrier Sockets for MultiChip Modules," *Proc. of the 1992 IEEE Multi-Chip Module Conference MCMC-92*, pp. 40-43, 1992.
16. D.-Y. Shih, B. Beaman, P. Lauro, K. Fogel, M. Norcott, G. F. Walker, J. Hedrick, L. Shi, F. Doany, and J. Shaw, "A Novel Elastomeric Connector for Packaging Interconnections, Testing, and Burn-In Applications," *Proc. of the 45th Electronic Components and Technology Conference*, pp. 126-133, 1995.
17. R. H. Parker, "Bare Die Test," *Proc. of the 1992 IEEE Multi-Chip Module Conference MCMC-92*, pp. 24-27, 1992.
18. L. DiFrancesco, C. Reynolds, "Socketing Chip-On-Board to a Multichip Module Using Particle Interconnect," *Proceedings of the SPIE*, vol. 1986, pp. 335-340, 1993.
19. A. M. Lyons, E. E. Hall, Y.-H. Wong, and G. Adams, "A New Approach to Using Anisotropically Conductive Adhesives for Flip Chip Assembly," *Proc. of the 45th Electronic Components and Technology Conference*, pp. 107-113, 1995.
20. D. Salzman and T. Knight, "Capacitive Coupling Solves the Known Good Die Problem," *Proc. of the 1994 IEEE Multi-Chip Module Conference MCMC-94*, pp. 95-100, 1994.
21. K. L. Lear, "Technologies for Highly Parallel Optoelectronic Integrated Circuits," *Proceedings of 1994 IEEE GaAs IC Symposium*, pp. 201-206, 1994.
22. R. A. Novotny, A. L. Lentine, L. M. F. Chirovsky, and T. K. Woodward "Receiver Design Issues for Parallel Optical Interconnections Fabricated in the FET-SEED Technology," *Proceedings of the SPIE*, vol. 2400, pp. 300-312, 1995.
23. K. P. Jackson, "High-Density, Array, Optical Interconnects for Multi-Chip Modules," *Proc. of the 1992 IEEE Multi-Chip Module Conference MCMC-92*, pp. 142-145, 1992.
24. C. A. Armiento, A. J. Negri, M. J. Tabasky, R. A. Boudreau, M. A. Rothman, T. W. Fitzgerald, and P. O. Haugsjaa, "Gigabit Transmitter Array Modules on Silicon Waferboard," *IEEE Trans. on Components Hybrids, and Manufacturing Tech.*, vol. 15, no. 6, pp. 1072-1079, 1992.
25. Y. C. Lee and N. Basavanahally, "Solder Engineering for Optoelectronic Packaging," *Journal of Metals*, vol. 46, no. 6, pp. 46-50, 1994.

26. K. H. Calhoun, C. B. Camperi-Ginestet, and N. M. Jokerst, "Vertical Optical Communication Through Stacked Silicon Wafers Using Hybrid Monolithic Thin Film InGaAsP Emitters and Detectors," *IEEE Photonics Technology Letters*, vol. 5, no. **, pp. 254-257, 1993.
27. R. F. Carson, P. K. Seigal, D. C. Craft, and M. L. Lovejoy, "Future Manufacturing Techniques for Stacked MCM Interconnections," *Journal of Metals*, vol. 46, no. 6, pp. 51-55 June 1994.
28. G. Stengel, H. Temkin, P. Brusenbach, T. Uchida, M. Kim, C. Parsons, W. E. Quinn, and S. E. Swirhun, "High-Speed Vertical-Cavity Surface Emitting Laser" *IEEE Photonics Technology Letters*, vol. 5, no. 12, pp. 1359-1362, 1993.
29. D. B. Young, J. W. Scott, F. H. Peters, M. G. Peters, M. L. Majewski, B. J. Thibeault, S. W. Corzine, and L. A. Coldren, "Enhanced Performance of Offset-Gain High-Barrier Vertical-Cavity Surface-Emitting Lasers," *IEEE J. of Quantum Electronics*, vol. 29, no. 6, pp. 2013-2021, 1993.
30. G. Hasnain, K. Tai, L. Yang, Y. H. Wang, R. J. Fischer, J. D. Wynn, B. Weir, N. K. Dutta, A. Y. Cho, "Performance of Gain-Guided Surface Emitting Lasers with Semiconductor Distributed Bragg Reflectors," *IEEE Journal of Quantum Electronics*, vol. 27, no. 6, pp. 1377-1385, 1991.
31. F. C. Prince, N. B. Patel, D. Kasemset, C. S. Hong, "Long Delay Time for Lasing in Very Narrow Graded Barrier Single-Quantum-Well Lasers," *Electronics Letters*, vol. 19, no. 12, pp. 435-437, 1983.
32. K. L. Lear, R. P. Schneider, K. D. Choquette, S. P. Kilcoyne, J. J. Figiel, and J. C. Zolper, "Vertical Cavity Surface Emitting Lasers with 21% Efficiency by Metalorganic Vapor Phase Epitaxy," *IEEE Photonic Technology Letters*, vol. 6, no. 9, pp. 1053-1055, 1994.
33. K. L. Lear, K. D. Choquette, R. P. Schneider, S. P. Kilcoyne, and K. M. Geib, "Selectively Oxidized Vertical Cavity Surface Emitting Lasers with 50% Power Conversion Efficiency," *Electronics Letters*, vol. 31, no. 3, pp. 208-209, 1995.
34. K. D. Choquette, K. L. Lear, R. P. Schneider, Jr., and K. M. Geib, "Cavity Characteristics of Selectively Oxidized Vertical-Cavity Lasers," *Applied Physics Letters* vol. 66, no. 25, pp. 3413-3415, 1995.
35. K. L. Lear, K. D. Choquette, R. P. Schneider, Jr., and S. P. Kilcoyne, "Modal Analysis of a Small Surface Emitting Laser with a Selectively Oxidized Waveguide," *Applied Physics Letters*, vol. 66, no. 20, pp. 2616-2618, 1995.
36. M. L. Lovejoy, G. A. Patrizi, P. M. Enquist, B. H. Rose, D. B. Slater, Jr., R. J. Shul, R. F. Carson, D. C. Craft, and J. A. Hutchby, "Low-Power, High Speed InGaAs/InP Photoreceiver for Highly-Parallel Optical Data Links," *GaAs IC Symposium Technical Digest 1995*, pp. 181-184, 1995.
37. R. Fornari and J. Kumar, "Infrared Absorption Spectra in Bulk Fe-Doped InP," *Appl. Phys. Lett.*, vol. 56, no. 7, pp. 638-640, 1990.
38. F. Mosel, A. Seidl, D. Hofmann, and G. Müller, "Infrared Absorption of n- and p-Type Fe-Doped InP and Mapping of the Fe Distribution", *Applied Surface Science*, vol. 50, pp. 364-368, 1991.
39. M. E. Warren, T. C. Du, J. R. Wendt, G. A. Vawter, R. F. Carson, K. L. Lear, S. P. Kilcoyne, R. P. Schneider, and J. C. Zolper, "Integration of Diffractive Lenses with Addressable Vertical-Cavity Laser Arrays," *Proceedings of the SPIE*, vol. 2398, pp. 12-20, 1995.
40. K. Rastani, M. Orenstein, E. Kapon, and A. C. Von Lehman, "Integration of Planar Fresnel Microlenses With Vertical-Cavity Surface-Emitting Laser Arrays," *Optics Letters*, vol. 16, no. 12, pp. 919-921, 1991.

41. O. Blum, S. Kilcoyne, M. E. Warren, T. C. Du, K. L. Lear, R. P. Schneider, R. F. Carson, G. Robinson, and F. H. Peters, "Vertical Cavity Surface-Emitting Lasers with Integrated Refractive Microlenses," *Electronics Letters*, vol. 31, no. 1, pp. 44-45, 1995.
42. E. M. Strzelecka, G. D. Robinson, M. G. Peters, F. H. Peters, and L. A. Coldren, "Monolithic Integration of Vertical-Cavity Laser Diodes With Refractive GaAs Microlenses," *Electronics Letters*, vol. 31, no. 9, 1995.
43. F. W. Ostermayer, Jr., P. A. Kohl, and R. H. Burton, "Photoelectrochemical Etching of Integral Lenses on InGaAsP/InP Light-Emitting Diodes," *Applied Physics Letters*, vol. 43, no. 7, pp. 642-644, 1983.
44. P. A. Kohl and D. B. Harris, "Photoelectrochemical Methods for Semiconductor Device Processing," *Electrochimica Acta*, vol. 38, no. 1, pp. 101-106, 1993.
45. M. Gdula, W. P. Kornrumpf, and B. K. Gilbert, "An Overlay Interconnect Technology for 1GHz and Above MCMs", *Proc. of the 1992 IEEE Multi-Chip Module Conference MCMC-92*, pp. 171-174, 1992.
46. T. Sudo, "Silicon-on-Silicon Technology for CMOS-Based Computer Systems", *Proc. of the 1992 IEEE Multi-Chip Module Conference MCMC-92*, pp. 8-11, 1992.
47. W. D. Miller, R. A. Gassman, and D. M. Keicher, "Laser Drilling of Vertical Vias in Silicon" Sandia National Laboratories Report SAND93-1772, Albuquerque, NM, 1993.
48. J. E. Morris, M. R. Feldman, W. H. Welch, M. Nakkar, H. Yang, J. Childers, Y. Raja, I. Turlik, G. Adema, P. McGill, and E. Young, "Prototype Optically Interconnected Multichip Module Based on Computer Generated Hologram Technology," *Proceedings of the SPIE*, vol. 1849, pp. 48-53, 1993.
49. N. C. Craft and A. Y. Feldblum, "Optical Interconnects Based on Arrays of Surface-Emitting Lasers and Lenslets," *Applied Optics*, vol. 31, no. 11, pp. 1735-1739, 1992.
50. J. Cheng, Y.-C. Lu, B. Lu, A. Alduino, G. Ortiz, J. C. Zolper, J. Klem, and G. A. Vawter, "High Speed Optical Interconnect Technology for a Reconfigurable Multi-Processor Network," *Proceedings of the SPIE*, vol. 2400, pp. 160-169, 1995.
51. P. S. Guilfoyle, R. V. Stone, J. M. Hessenbruch, and R. Kaminski, "Photonic Switching Demonstration Focused Towards HPOC Module Implementation," *Proceedings of the SPIE*, vol. 2400, pp. 16-23, 1995.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.