Design and Testing of a Kinematic Package Supporting a 32 × 32 Array of GaAs MQW Modulators Flip-Chip Bonded to a CMOS Chip

Michael H. Ayliffe, David R. Rolston, *Member, IEEE, Member, OSA*, Alan E. L. Chuah, *Student Member, IEEE*, Eric Bernier, *Member, IEEE*, Feras S. J. Michael, David Kabal, *Member, IEEE*, Andrew G. Kirk, *Member, IEEE*, and David V. Plant, *Member, IEEE*

***Abstract*—Innovative approaches to the design and packaging of a high-performance module supporting a 32 × 32 array of GaAs multiple quantum-well (MQW) modulators flip-chip bonded to a 9 × 9 mm^2 complementary metal–oxide–semiconductor (CMOS) chip are described. The module integrates a minilens array, a copper heat spreader, a thermoelectric cooler (TEC) and an aluminum heatsink. The minilens array is aligned and packaged with the chip using a novel six degrees of freedom (DOFs) alignment technique. The kinematic design allows for the manual insertion of the module into a free-space optical system with no need for further adjustments. The chip is mounted directly on a flexible printed circuit board (PCB) using a chip-on-board approach, providing over 200 bond pad connections to the chip. Impedance-controlled lines were operated at 1.0 Gb/s with a crosstalk of 4.0% between nearest neighbor lines. The junction-to-TEC thermal resistance is 0.4 °C/W, allowing for the use of a single-stage TEC to regulate the chip at an operating temperature of 40 °C under a maximum thermal load of 13.1 W.**

***Index Terms*—Kinematic design, optical interconnections, optoelectronic device arrays, packaging.**

I. Introduction

TWO-DIMENSIONAL (2-D) parallel optical interconnects (POIs) which utilize either free-space or guided-wave solutions enable high-density, high-bandwidth communications between very large scale integration (VLSI) chips in computing and switching systems [1], [2]. 2-D-POI technologies have experienced significant progress in recent years, primarily due to the advances made in optoelectronic-VLSI (OE-VLSI) and microoptical device technologies. Large 2-D arrays of surface-normal devices (such as vertical cavity surface-emitting lasers and electroabsorption modulators) are now routinely flip-chipped to foundry complementary metal–oxide–semiconductor (CMOS) chips [3]–[6]. Similarly, highly efficient multifunction diffractive and refractive microoptical components are available commercially [7]. Despite these achievements, there still exists a significant challenge involved in the construction of 2-D-POI systems which has to do with the difficulty of interfacing 2-D arrays of optoelectronic and optical devices to one another in a repeatable and stable manner. Some even contend that research in this field has now reached a point where the complexity of implementing the interface technologies far exceeds any shortcomings in the performance of the optoelectronics or the optics [8]. It is now apparent that the commercial deployment of 2-D-POIs is contingent on the development of novel interface solutions compatible with low-cost and high-volume manufacturing. These new techniques must simplify the assembly process while satisfying the requirements of industrial installations in terms of mechanical alignment, electrical signal integrity and thermal dissipation. Successful interface solutions must consider all these design aspects throughout the development stage.

This paper addresses these issues by describing the design and performance of a package that supports a 32 × 32 array of GaAs multiple quantum-well (MQW) modulators (half of which are used as detectors) flip-chip bonded to a 9 mm × 9 mm^2 CMOS chip. This work builds on an earlier design developed for a different chip [10]. The OE-VLSI chip is used in a four-stage photonic backplane demonstrator whose functional layout is shown in Fig. 1.

The backplane supports parallel high-speed optical communication channels between electronic processors located on different printed circuit boards (PCBs) referred to as motherboards. Although only four motherboards are shown in Fig. 1, this scheme can scale up to an arbitrary number of boards. The OE-VLSI chips operate in one of three modes: i) electrical data from a motherboard can be injected onto the backplane (transmitting state), ii) optical data on the backplane can be extracted by a motherboard (receiving state), or iii) optical data can be regenerated by an OE-VLSI chip via an optical-to-electrical-to-optical conversion performed by the on-chip high-speed transceivers (transparent state). In the transparent state, optical signals can quickly "hop" from one stage to the next without having to go through CMOS output pads. This is a key feature of OE-VLSI technology; it allows for low-power, low-latency and high-speed parallel transmission of data across the backplane. A detailed description of

Manuscript received October 3, 2000; revised June 6, 2001. This research was supported by a Grant from the Canadian Institute for Telecommunications Research as part of the National Center of Excellence program of Canada and NSERC and FCAR postgraduate fellowships.

M. H. Ayliffe and D. Kabal are with Picolight Inc., Boulder, CO 80301 USA.
D. R. Rolston is with FCI Canada, Dorval, QC H9P 2X6, Canada.
A. E. L. Chuah, A. G. Kirk, and D. V. Plant are with the Department of Electrical and Computer Engineering, McGill University, Montréal, QC H3A 2A7, Canada.
F. S. J. Michael is with Xanoptix, Nashua, NH 03060 USA.
E. Bernier is with Nortel Networks, Kanata, ON K2K 2B5, Canada.
Publisher Item Identifier S 0733-8724(01)08181-6.

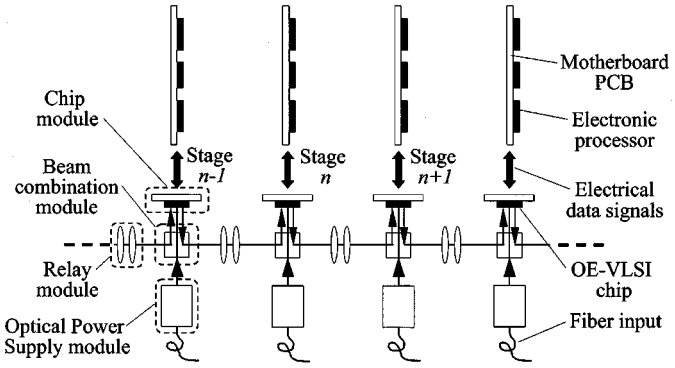

Fig. 1. Functional outline of a modulator-based photonic backplane demonstrator.

the architectural issues involved in the design of a photonic backplane can be found elsewhere [9].

To facilitate assembly and servicing, the system is divided into modular subassemblies. For the modulator-based system shown in Fig. 1, the interconnect is divided into four functional modules: the optical power supply generates a 2-D array of continuous wave (CW) read-off beams aligned onto the modulators, the beam combination module routes the receiving, read-off and transmitting beams through the interconnect, the relay module propagates the transmitting beams from one stage to the next and the chip module performs the electronic-to-optical and optical-to-electronic conversions between the motherboard electronics and the optical interconnect.

This paper focuses on the design, implementation, and experimental evaluation of the chip module. The intent is to provide a wide description of the multidisciplinary design considerations and tradeoffs involved in packaging 2-D OE-VLSI chips in the context of a highly parallel free-space 2-D-POI demonstrator. The paper is organized as follows. Section II discusses the optomechanical, electrical, and thermal considerations constraining the design of the chip module. Section III provides an overview of the mechanical design and a description of the assembly sequence. Section IV describes the alignment technique used to package the OE-VLSI chip with a minilens array element. The design and experimental evaluation of the optomechanical interface allowing for the 32×32 device array to be manually inserted and removed from the free-space backplane is presented in Section V. To the authors' knowledge, this constitutes the first demonstration of a replaceable OE-VLSI chip in a large-scale free-space demonstrator system. Electrical and thermal characterization of the module is presented in Sections VI and VII, respectively. Section VIII concludes the paper.

II. System Considerations and Design Choices

A. Physical Layout of System

A fundamental design consideration is whether or not to mount the OE-VLSI chip directly on the motherboard or to package the chip separately in a module that interfaces with the interconnect optics. Mounting the chip directly on the motherboard offers a number of advantages: it facilitates routing of the data paths, minimizes transmission latency, reduces signal integrity issues related to transmission line effects and electromagnetic interference, and removes the need for a high-bandwidth connector between the module and the motherboard. This approach, however, is difficult to implement in practice due to the alignment constraints it imposes on the motherboard PCBs. This problem is easily understood when we consider the alignment operations required for optical signals to be received and transmitted by the OE-VLSI chip at stage n: signals arriving from stage $n-1$ need to be aligned to the proper detectors and the optical signals transmitted by stage n must be aligned onto the proper detectors at stage $n+1$. Also, in the case of modulator-based systems, an additional alignment operation is required to position the read-off beams onto the modulator array. These operations are not trivial to perform because 2-D device arrays need to be aligned in all six degrees of freedom (DOFs): lateral (x, y), focus (z), rotational (θ_z), and tilt (θ_x, θ_y). For this reason, attempts at mounting the OE-VLSI chip directly on the motherboard have been limited to board-to-board demonstrators implementing a low number of optical channels [11]–[14] or systems employing some type of alignment compensation techniques such as liquid-crystal steering devices or array redundancy [15], [16]. In all cases, however, chips on adjacent motherboards are mounted face to face, a configuration that takes away the possibility of operating the OE-VLSI chip in the transparent state; data propagating through multiple stages across the backplane must go through CMOS pad drivers at every stage. Such approaches significantly increase power dissipation and transmission latency, which tend to limit system scalability.

The alternative option is to package the chip in a separate module and attach it to the motherboard via a flexible electronic interconnecting medium [10], [17]–[19]. This provides the required mechanical isolation of the module from the motherboard: it removes the alignment constraints placed on the motherboard PCB and allows for individual OE-VLSI chips to be precisely aligned to the optical system. Furthermore, OE-VLSI chips can now be oriented at 90° with respect to the motherboard PCBs, as depicted in Fig. 1, which facilitates the design and reduces the complexity of the optics used to interconnect one stage to the next. This results in a highly scalable system, which fully exploits the benefits of direct termination at the chip.

The disadvantages of the latter approach are the advantages of the former. The increased physical distance between the chip and the motherboard electronics can potentially aggravate signal integrity. The limited pin-count of the connector effectively puts an upper bound on the amount of data that come on and off the backplane. Furthermore, the footprint of the chip module becomes limited to the motherboard pitch and this may constrain routing of data lines to the chip.

This demonstrates how the location of the OE-VLSI chip relative to the motherboard affects several aspects of the system design. Although packaging the OE-VLSI chip in a separate module is certainly not optimal, it does lead to a practical approach to the construction of a scalable photonic backplane and it is the scheme adopted in this paper. It is understood that the use of a connector is undesirable and that future implementations will probably require the chip to be mounted directly on the motherboard. This will become possible with the development of novel interface technologies allowing motherboard PCBs to

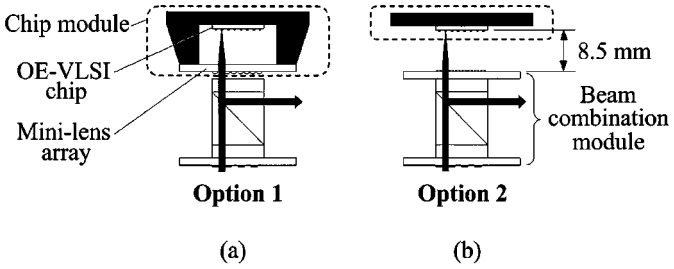

Fig. 2. Alternative approaches to the packaging of the minilens array. (a) In option 1, the minilens array is packaged with the OE-VLSI chip. (b) In option 2, the minilens array is integrated as part of the beam combination module.

be manually inserted and removed from the backplane without upsetting system alignment.

B. Constraints Imposed by the Optical Interconnect Design

The optical design of the four-stage photonic backplane demonstrator [20] is based on the concept of optoelectronic device clustering, in which individual minilenses are used to image a small array of optical beams [21]. In this scheme, the surface-normal devices are not uniformly distributed across the surface of the OE-VLSI chip; instead, they are organized in clusters of devices, in which each cluster is centered onto its respective minilens. This approach offers a number of advantages, including higher system scalability and increase in tolerance to misalignment [22].

The features of the interconnect impacting the design of the chip module are as follows. The interconnect uses telecentric relays of 8×8 minilens arrays, each minilens having a diameter of 800 μm and a focal length of 8.5 mm, resulting in a spot size of 26.2 μm ($1/e^2$ diameter in intensity) at the device plane. The MQW modulator and detector devices are arranged in an 8×8 array of clusters, each cluster consisting of a 4×4 array of devices. To relax alignment tolerances, devices are oversized compared to the spot size; modulators are 50 μm in diameter and detectors are 65×65 μm^2. There are 1024 devices occupying a 6.4×6.4 mm^2 area on the chip, which results in a density of 2500 devices/cm^2. The physical distance separating two adjacent stages is 47 mm; this distance is critical because it puts an upper bound on the lateral dimensions of the chip module.

C. Misalignment-Tolerant Module

A critical design requirement was for the chip to be manually removable and replaceable from the interconnect optics. This is necessary because the OE-VLSI chip is the only active element in the photonic backplane; it is prone to failure and is likely to be upgraded with time. Replaceability of the chip module is an essential feature of a practical system and this calls for a kinematic alignment scheme at the chip module-to-optical backplane interface.

With this in mind, two different packaging choices were examined. In the first case, the minilens array is aligned to the chip and integrated as part of the chip module [see Fig. 2(a)]. In the second case, the minilens array is integrated in the beam combination module [see Fig. 2(b)]. At first glance, the latter option seems more practical because the minilens array is directly adjacent to the beam combination optics while there is a large gap (8.5 mm) distancing it from the chip—a larger gap typically complicates the alignment of out-of-plane components. Results of the alignment tolerance analysis comparing the two options are shown in Table I. The analysis was performed by calculating the allowed misalignment before the power throughput of an optical channel between two stages drops to 99% of its value for a perfectly aligned system. The analysis is performed by varying one DOF at a time. Only optical power losses caused by clipping effects are considered and the calculations are based on paraxial Gaussian beam propagation theory and geometrical optics for off-axis rays [23].

TABLE I
MISALIGNMENT TOLERANCE ANALYSIS

Degree of freedom	Option 1	Option 2
Lateral	± 26 μm	± 8 μm
Tilt	± 0.03°	± 0.12°
Rotational	± 0.36°	± 0.10°
Longitudinal	± 500 μm	± 125 μm

The results of Table I indicate that integrating the minilens array with the OE-VLSI chip relaxes the lateral, longitudinal, and rotational alignment tolerances by a factor of approximately four, but tightens the tilt alignment tolerances by the same factor. The improvement in lateral and rotational tolerances results from the surface-normal devices looking bigger when viewed through the minilens array. The decrease in tilt tolerance arises from the long focal length of the minilenses. The tradeoff between lateral and tilt tolerances is inherent to the alignment problem and can be linked to the principle of optical invariance [22].

By integrating the minilens array with the chip [see Fig. 2(a)], a kinematic design becomes possible mainly because this choice results in lateral and rotational alignment tolerances that can be achieved with standard machining capabilities, the accuracy of modern computer numerically controlled (CNC) machines being limited to about ± 10 μm [24]. This choice also gives rise to new design challenges. First, the minilens array must be accurately aligned to the OE-VLSI chip in all six DOFs. This is achieved by a novel alignment technique that uses off-axis diffractive elements placed on the CMOS chip along with metal alignment markers located on the minilens array substrate (Section IV). Second, the design of the interface mechanics must exhibit outstanding angular repeatability to satisfy the stringent tilt tolerance of the module ($\pm 0.03°$). This is accomplished by taking advantage of the optical flatness of the minilens array substrate and using this surface as the passive alignment plane that defines the tilt of the module (Section V).

D. Electrical Packaging Issues

The basic functions of the package are to provide stable mechanical support for the chip, distribute power, control and data lines to the CMOS chip, and offer an adequate means for removal of heat. The CMOS chip has a total of 232 bond pads, of which 207 are used during normal chip operation (25 unused pads are dedicated to device testing). In order to meet the spatial

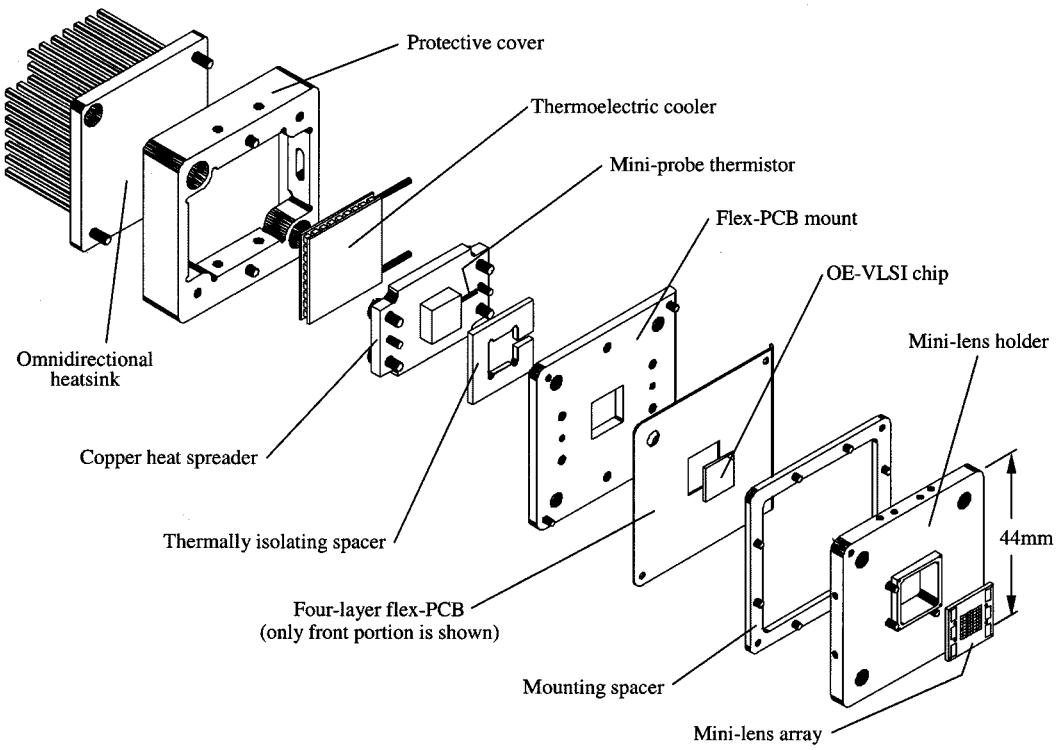

Fig. 3. Exploded view of the chip module design.

constraints imposed by the optics, a packaging scheme capable of providing the required connectivity in an area smaller than 47×47 mm^2 was needed. In addition, the design had to support a mechanically flexible connection to the motherboard PCB.

These objectives were realized using a chip-on-board (COB) packaging scheme in combination with flexible PCB (flex-PCB) technology with impedance-controlled microstrip lines [25]. In our approach, the OE-VLSI chip was mounted directly on a copper heat spreader that was subsequently inserted through an opening machined in the flex-PCB. The chip was wirebonded directly to the flex-PCB, thereby eliminating the need for an electronic chip carrier. This resulted in a small, simple, and low-cost method of chip packaging. This packaging scheme also improves high-frequency performance due to the absence of package lead inductance and the ability of placing surface mount components (i.e., decoupling capacitors and termination resistors) in close proximity to the chip bond pads.

E. Need for Temperature Stabilization

The basic principle behind the operation of MQW modulators is the wavelength shift of the excitonic absorption peak with applied voltage [26]. In 2-D-POI applications, MQW modulators are operated in reflection mode where a variation of the applied voltage translates into a change in reflectivity ΔR. For a given voltage swing, MQW modulators are usually designed to maximize ΔR for a given operating wavelength and temperature. Any temperature variation, however, results in a shift of the excitonic peak away from its optimal position. For MQW devices based on a GaAs/AlGaAs quantum well structure, this shift is approximately 0.27 nm/°C [27]. This dependence on temperature is undesirable because it limits the temperature range over which MQW modulators can operate efficiently under constant biasing conditions.

The useful temperature range of MQW modulators has been investigated by Venditti [28]. His measurements were performed on GaAs/AlGaAs modulators with a quantum well structure identical to the one used here. Assuming constant biasing conditions and a 5-V swing, his results show that the operating temperature must be within ± 5 °C of the designed temperature in order for ΔR not to fall below 90% of its optimal value. If the latter criterion is relaxed to 70%, the useful temperature range is effectively doubled. Thus, it can be concluded that without any means of controlling the operating temperature, it would be difficult, if not impossible, to operate MQW modulators efficiently. For this reason, a thermoelectric cooler (TEC) was incorporated into the chip module. This choice was motivated by the small size, the light weight, and the ability of TECs to achieve precise temperature stabilization when combined with a thermistor in a closed-loop configuration.

III. Chip Module Design Overview

A. Physical Dimensions

An exploded view of the chip module is shown in Fig. 3. The design integrates a minilens array, an OE-VLSI chip, a copper heat spreader, a TEC and a heatsink. This integration results in a compact module with a footprint of 44×44 mm^2, satisfying the physical constraints imposed by the optical design. The length of an assembled module, measured from the front of the minilens array to the back of the heatsink, is 45 mm, half of which is occupied by the heatsink fins.

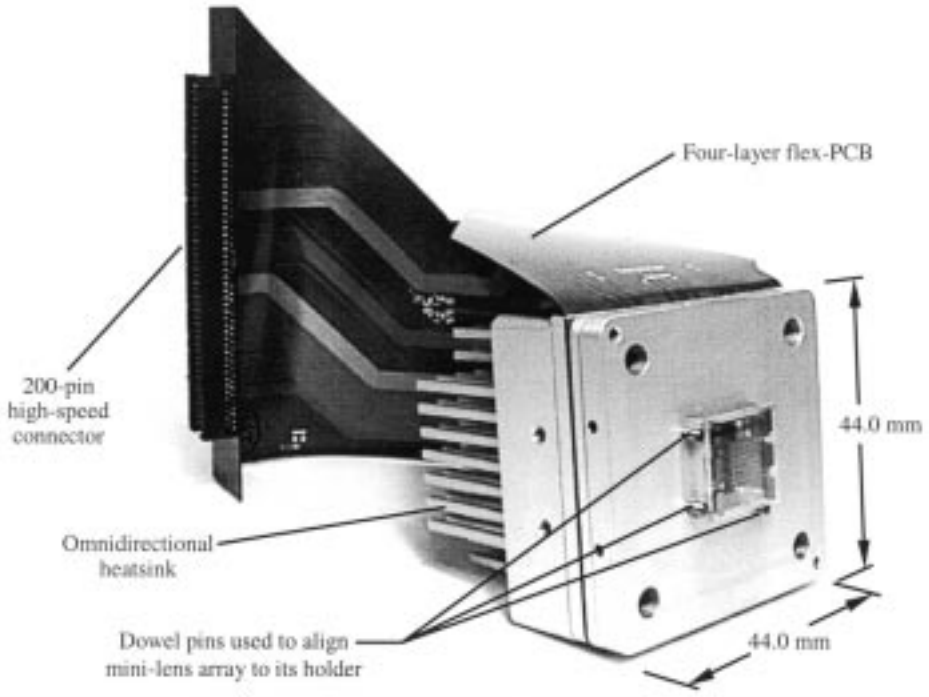

Fig. 4. Photograph of an assembled chip module.

B. Materials Selection

All components were custom-designed and produced on a CNC machine, except for the TEC and heatsink, which were commercially available. The minilens holder, mounting spacer, flex-PCB mount, and protective cover were machined out of aluminum 6061, followed with a clear sulfuric anodization finish. The choice of this alloy was motivated by its ease of machining, lightness, low residual stress, low cost, and availability. In addition, the high electrical conductivity and nonsparking properties of aluminum make it an excellent electrical shielding material. The purpose of the anodization step was to electrically isolate the mechanical components to eliminate any possibility of electrical shorts with the nearby COB electronics. Note that sulfuric anodization was chosen instead of hard anodization to maintain the tight machining tolerances specified on some of the components; this is because sulfuric anodic coatings are typically 5–15 μm thick while hard anodizing coatings can be as much as 50 μm thick [24].

The finned heatsink was made of 6063 aluminum, followed by a clear sulfuric anodization finish. The use of 6063 aluminum for heatsinking applications is common and is mainly due to the fact that 6063 conducts 15% more heat than 6061 and can more easily be extruded into complex shapes. The heat spreader was made of a high copper alloy (C18500), a choice motivated by its high thermal conductivity, its satisfactory machinability rating, and its availability. High thermal conductivity was required to provide an efficient thermal path between the chip and the TEC. Also, the heat spreader was nickel-plated to prevent copper oxidation.

C. Chip Module Assembly

A photograph of an assembled chip module is shown in Fig. 4. The design of the chip module offers a high degree of modularity. Components are joined to one another using dowel pins and screws. A module can be disassembled and reassembled in a few minutes. The design also allows for an OE-VLSI chip to be removed and replaced (of course, this requires the wirebonds to be reworked). The assembly process was broken down in three main steps: i) mounting and wirebonding the chip to the flex-PCB, ii) mounting the minilens array onto its holder, and iii) alignment of the minilens array to the chip. What follows is a detailed description of this assembly sequence.

Referring to Fig. 3, the chip was passively aligned and fixed to the heat spreader. Chips were precision-diced to within 50 μm of the bond pad frame and the pedestal portion of the heat spreader was machined to within ± 10 μm of the chip's nominal dimensions. This allowed for the alignment of the chip to be performed manually using the sidewalls of the heat spreader pedestal as passive alignment references. The chip was attached using silver-filled conductive epoxy, providing an efficient electrical and thermal path to the heat spreader. Next, the chip was inserted through the opening of the flex-PCB, the latter having previously been glued onto the flex-PCB mount using high temperature thermal set lamination techniques. A 1.0 mm thick acetal spacer was placed between the flex-PCB mount and the heat spreader. The chip was then centered on the flex-PCB opening using a pair of acetal dowel pins and the heat spreader was fastened to the back of the flex-PCB mount using nylon screws. The use of plastic screws, pins and spacers is to avoid a "thermal short" between the heat spreader and the aluminum mount. The chip was then wirebonded to gold-on-nickel plated printed circuit using an aluminum wedge wirebonder. The TEC, protective cover and heatsink were then incorporated. The TEC was clamped in place by fastening the heatsink directly to the back of the flex-PCB mount using a pair of nylon screws. Thermally conductive, self-adhesive, elas-

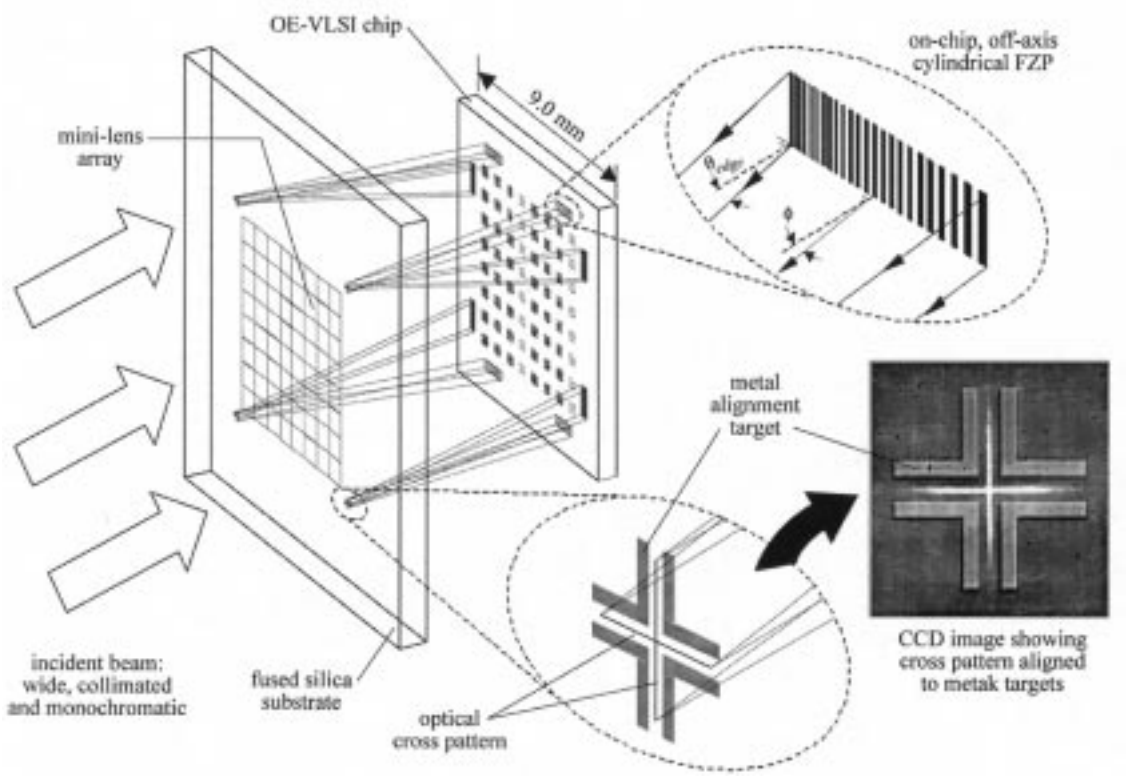

Fig. 5. Six-DOF alignment technique used for packaging the minilens array to the OE-VLSI chip (to scale).

tomer interface pads (250 μm thick) were placed on both sides of the TEC to ensure efficient heat flow across the interfaces.

Next, the minilens array was positioned and glued to its holder. The minilens substrate (fused silica, 1.0-mm thick) was precision-diced to within ± 25 μm of a lithographically defined metal frame. The positioning of the minilens array was done manually using a semikinematic passive alignment scheme. As shown in Fig. 4, the lateral position (x, y) and rotation (θ_z) of the minilens substrate are constrained by three press-fit precision dowel pins inserted in the minilens holder (dowel pin positional tolerance $= \pm 10$ μm). The remaining DOFs, tilt (θ_x, θ_y) and longitudinal (z), are constrained by the minilens holder pedestal (flatness tolerance $= \pm 20$ μm). The minilens array was fixed in place using ultraviolet curing optical epoxy.

The last step involves the alignment of the minilens array to the chip. First, an aluminum spacer (see Fig. 3) was glued onto the flex-PCB. The minilens array was then positioned in front of the chip and the two were brought into alignment using a technique described in Section IV. During the alignment process, the minilens holder was free to move in all six DOFs and did not come into physical contact with the aluminum spacer (there was a 400-μm gap separating the two components). Once alignment was completed, the minilens holder was fixed to the mounting spacer using room temperature epoxy.

IV. Novel Alignment Technique for Minilens Packaging

A. Need for Active Alignment

The benefits of integrating the minilens array with the chip module rely on the development of a technique capable of accurately aligning the minilens array to the chip. If the tolerance specifications associated with this alignment task are to follow the misalignment metric used in Section II (i.e., 1% optical power drop between adjacent stages), then the accuracy of the technique itself must satisfy the tolerance requirements listed in the last column of Table I.

Whenever possible, passive alignment techniques are preferred because of their greater manufacturability. For this reason, packaging lens arrays with OE-VLSI chips using a solder self-alignment method is being actively researched [29]. In the case at hand, however, the optical path length separating the minilens array from the chip is 8.5 mm, a distance that prohibits the use of this technology. Alternative passive alignment schemes were investigated; in all cases, the cumulative tolerances associated with the dicing, machining, and mounting operations far exceeded the required specifications, indicating the need for an active alignment technique.

B. A Novel Six-DOF Alignment Technique

The technique used to align the minilens array to the OE-VLSI chip is described in Fig. 5. A wide, collimated, and monochromatic beam of uniform intensity is incident normal to the minilens array substrate. The portion of the incident beam falling outside the region occupied by the 8×8 minilens array traverses the fused silica substrate and illuminates reflective diffractive elements located on the periphery of the OE-VLSI chip. In our implementation, the on-chip diffractive elements are cylindrical Fresnel zone plates (FZPs) that were fabricated using the top metal layer of the CMOS process (see Fig. 6). A cylindrical FZP focuses an incident plane wave into a line and thus a pair of such elements, operated off-axis and oriented at

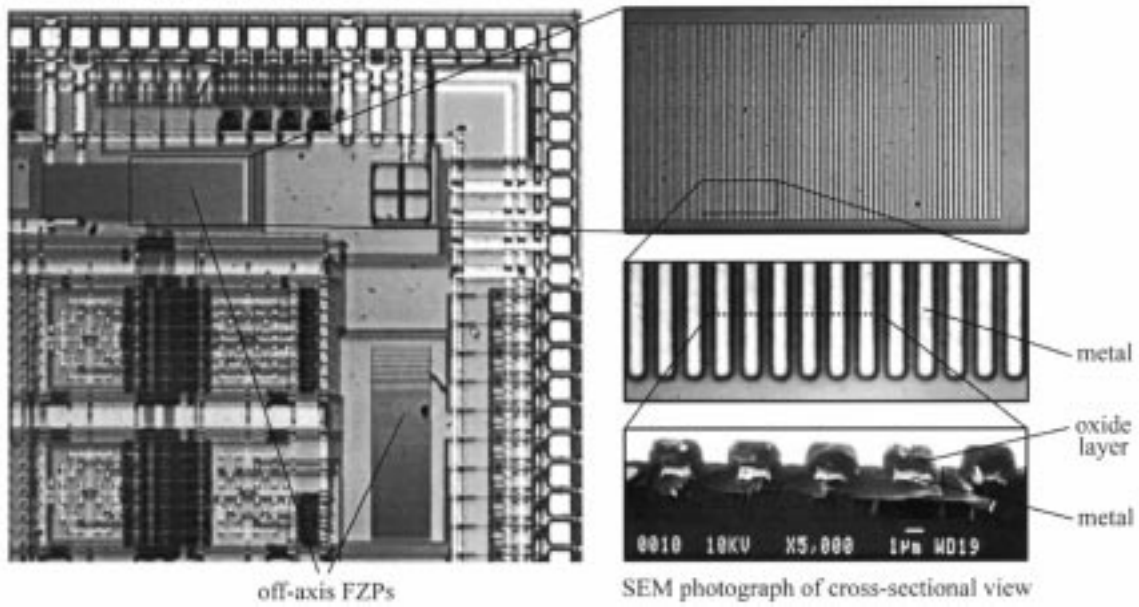

Fig. 6. Photographs of the off-axis FZPs fabricated using the top metal layer of the CMOS chip.

90°, can be used to create a cross pattern in the focal plane. A pair of orthogonal cylindrical FZPs was placed in each corner of the chip. The focal length of the FZPs was 8.5 mm such that the reflected cross patterns were focused in the minilens array plane. Metal alignment targets were placed on the minilens array substrate; these were lithographically defined and fabricated at the same time as the minilenses. Alignment was achieved by positioning the OE-VLSI chip such that all focused cross patterns were properly registered with respect to the metal alignment targets.

Our technique is similar to the one developed by Boisset *et al.* [30] (where four on-chip circular FZPs were used), but with two significant improvements. First, our technique results in a more efficient usage of silicon area; this is because a pair of cylindrical FZPs can be made much smaller than a circular FZP with the same f number and focal length. In addition, cylindrical FZPs are easier to implement because rectangular geometries are readily compatible with commercial integrated circuit layout CAD tools. Second, the cylindrical FZPs were designed for off-axis operation, so that reflected beams were being focused at an angle. This means that the trajectory of the intersection of the focusing cross pattern does not follow a line perpendicular to the chip plane. Consequently, the intersection point shifts laterally as it is imaged at different planes along the optical axis, as shown in Fig. 7. The off-axis operation of the FZPs is the most important aspect of our implementation because it results in a technique sensitive to all six DOFs. Note that only three cylindrical FZP pairs are required to determine misalignments in all six DOFs. The fourth FZP pair is redundant and used as a backup.

C. FZP Design Considerations

Sensitivity to tilt and longitudinal misalignments is a function of the off-axis angle ϕ (see Fig. 5). The larger the off-axis angle, the larger the lateral shift of the cross as it is imaged in different planes along the optical axis and the higher the sensitivity. The off-axis angle can be maximized in two ways: by maximizing the diffraction angle on the outer edge of the lens and by removing low-order zones on the inside of the lens. The diffraction angle on the outer edge is limited by the minimum feature size allowable by the fabrication technology. In the present case, the limitation came from the CMOS design rules of the top level metal layer, which specified a minimum trace width and spacing of 2 μm, which translated into a maximum edge diffraction angle of 12.30°. The removal of low-order zones reduces the lens aperture; this increases the f-number and leads to a design compromise. This is because a larger f-number increases the diffraction-limited width of the crossing lines in the focal plane and wider lines increase the error involved in judging when the cross is correctly registered relative to the metal alignment targets.

The FZPs used in our implementation have a length of 1000 μm and a height of 300 μm. The design wavelength is 852 nm. This results in an off-axis propagation angle of 8.74° and a diffraction-limited line width of 14.5 μm. The error involved in judging the proper registration of the cross patterns depends on the optical magnification used and the width of the focused lines. By imaging the alignment targets on a charged-coupled device (CCD) camera with a $\times 20$ objective, it was determined that a ± 3 μm lateral shift of the cross relative to the alignment target could be clearly discriminated. Using this and the fact that alignment targets are placed 7150 μm apart on the minilens substrate, the alignment sensitivity for each DOF individually were found to be: lateral $= \pm 3.0$ μm, longitudinal $= \pm 20$ μm, rotational $= \pm 0.017°$, and tilt $\pm 0.11°$. These results satisfy the alignment requirements listed in the last column of Table I.

V. Optomechanical Interface Design and Testing

A. Optomechanical Design

Referring to the results of Table I, the optomechanical design must allow for the repeatable insertion of the chip module

(a)

(b)

(c)

Fig. 7. Cross alignment patterns imaged at three different planes along the optical axis. (a) 100 μm behind the focal plane. (b) At the focal plane. (c) 100 μm in front of the focal plane.

into the optical system with the following tolerances: lateral = ± 26 μm, rotational = $\pm 0.36°$ and tilt = $\pm 0.03°$. A simplified

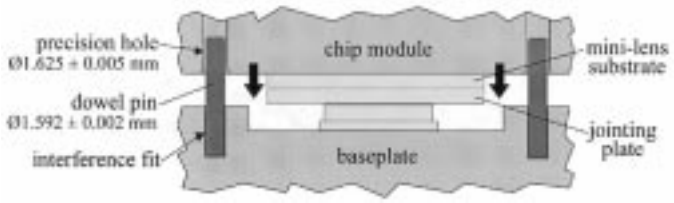

Fig. 8. Simplified representation of the semikinematic interface design.

drawing of the optomechanical design of the interface between the chip module and the beam combination module (BCM) is shown in Fig. 8. A complete description of the design and assembly of the BCM can be found in [31]. The design of the interface uses a pair of precision dowel pins (diameter tolerance of ± 2 μm) press-fit into the baseplate and a pair of precision holes (diameter tolerance of ± 5 μm) machined in the chip module. A three-dimensional (3-D) drawing showing the location of the pins and holes is shown in Fig. 9.

The strategy behind this design is to take advantage of the optical-grade flatness of the optical substrates at the interface and use them as passive alignment planes to control the tilt of the chip module relative to the BCM. Tilt misalignment is minimized when the minilens substrate comes in full contact with the first optical surface of the BCM (referred to as the jointing plate). The dowel pins and alignment holes are specified with tight tolerances, enough to provide adequate lateral and rotational alignment of the chip module relative to the BCM. Hence, the dowel pins constrain the lateral and rotational DOFs of the chip module while the flat optical substrates constrain the tilt and longitudinal DOFs. The alignment of the chip module is achieved simply by applying a force directed toward the BCM. This force is provided by a pair of spring-loaded screws inserted from the back of the chip module through the clamping holes shown in Fig. 9.

A potential problem of semikinematic designs, in general, is the possibility for components to be overconstrained. In the current design, there will always be a slight tilt misalignment between the jointing plate and the dowel pins. If the fit between the pins and the mating holes is too close, then it may become impossible for the minilens substrate to come in full contact with the jointing plate due to obstruction of dowel pins by the holes. Increasing the size of the holes increases the angular play of the chip module, but this can only be done at the expense of lateral and rotational precision. There is, thus, a tradeoff between lateral alignment and angular play.

The worst case lateral misalignment of the chip module relative to the BCM is the sum of the following tolerances: i) the locational accuracy of the dowel pins (± 10 μm, limited by CNC capability), ii) the alignment accuracy of the BCM relative to the dowel pins (± 20 μm, limited by mounting technique), iii) the locational accuracy of the precision holes (± 10 μm, limited by CNC capability), iv) the alignment accuracy of the minilens array relative to the precision holes (± 20 μm, limited by mounting technique), and v) the worst case misalignment of the dowel pin inside a precision hole (± 20 μm, corresponding to a minimum-size pin inserted in a maximum-size hole). The sum of the above tolerances leads to a worst case lateral misalignment of ± 80 μm, which far exceeds the allowances of the

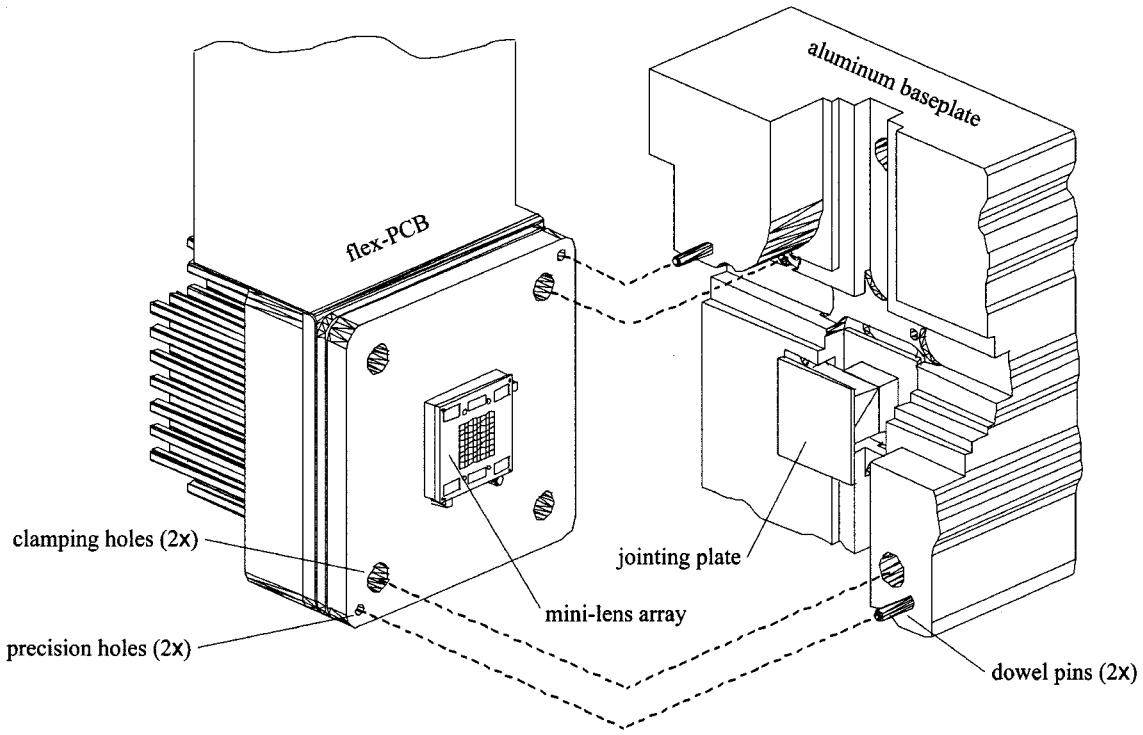

Fig. 9. Interface design showing the location of the dowel pins and precision holes.

misalignment budget (± 26 μm). To circumvent this problem, a pair of tilt plates were inserted in the optical power supply (OPS) module and used to adjust the lateral alignment of the beams incident on the chip module. In effect, the tilt plates allow for the first four tolerances in the above list to be compensated for by slightly shifting the beam array. Note that the tilt plates are required to be adjusted only once, during the first insertion of the chip module, in a manner similar to a calibration procedure. This results in subsequent insertions to have a worst case lateral misalignment determined solely by the pin-to-hole misalignment (± 20 μm). Using the fact that dowel pins are separated by a distance of 56.6 mm, the worst case rotational misalignment is calculated to be $\pm 0.04°$. Thus, the worst case lateral and rotational misalignments are within the allowances of the misalignment budget.

The last point to be considered concerns the angular play of the chip module. The design specified the penetration depth of the dowel pins in the precision holes to be 3.0 mm. This results in a minimum angular play of $\pm 0.5°$, which is plenty to ensure that the optical substrates come in full contact.

B. Experimental Evaluation of Interface Repeatability

To evaluate the repeatability of the optomechanical interface, a diagnostic chip module (DCM) was assembled using the alignment technique described previously in Section IV. A DCM is the equivalent of a chip module for which the heatsink, TEC, and heat spreader have been removed and the OE-VLSI chip replaced with a transparent fused silica substrate having lithographically defined metal targets replicating the location and size of the modulators and detectors on the chip. Thus, a DCM uses exactly the same optomechanics as a chip module and provides easy access to the back of the module, allowing for the spots' position in the device plane to be directly observed with a CCD camera. The positional information of the spots is used to quantify the repeatability of the chip module fixture.

Prior to the repeatability measurements, the alignment of the OPS beam array must be performed. To do this, a DCM is inserted in the system and secured in place. The lateral alignment of the beams is adjusted using the OPS tilt plates; this is done by imaging the minilens array from the back of the DCM and centering the beams on their respective minilenses. The tilt alignment of the beams is adjusted using Risley prisms inserted in the OPS (see [32]); this is done by imaging the device plane from the back of the DCM and centering the spots on their respective metal targets. The result of this alignment step is shown in Fig. 10, showing all 512 spots aligned to the DCM metal targets.

Repeatability measurements were performed by removing the chip module completely, inserting it back in the system, and securing it in place. This removal/insertion operation was repeated 50 times. For each removal and insertion cycle, the lateral misalignments of the two spots located at the opposite corners of the array were recorded. The upper left and lower right corner spots are separated by a distance of 7.17 mm across the chip. The misalignment data is plotted in Fig. 11. The resulting standard deviation in spot misalignment is $\sigma = 2.2$ μm. Assuming a random process and a normal distribution, this means that each spot in the array will fall within a circle of diameter 6.6 μm, centered on the OE devices, with a probability of $>99\%$. Thus, by using modulators and detectors that are larger than the 3ω spot diameter by the latter amount, the chip module can be manually inserted into or removed from the optical system with negligible insertion losses and no need for further adjustments. In our system, the OE devices

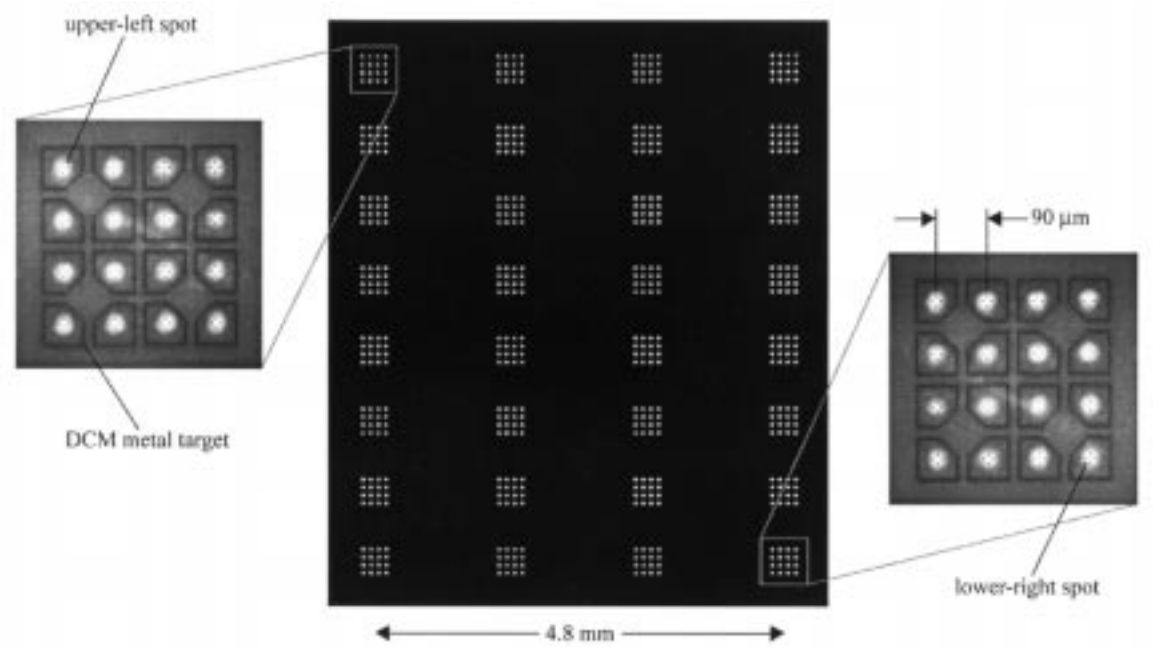

Fig. 10. CCD images showing alignment of spots on the DCM metal targets.

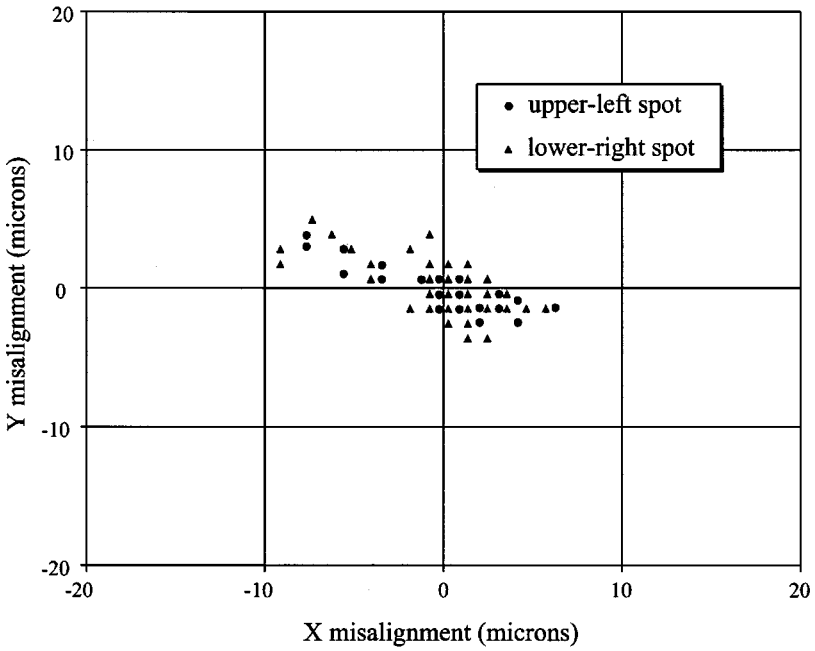

Fig. 11. Characterization of the kinematic design. Repeatability data for 50 removal and insertion cycles.

were generously oversized relative to the spot diameter: the 3ω spot diameter was 39 μm, the modulators were 52.5 μm in diameter, and the detectors were 66 μm $\times$ 66 μm. To the authors' knowledge, this constitutes the first demonstration of a replaceable 2-D OE-VLSI chip in a large-scale (>1000 devices per cm^2) free-space demonstrator.

VI. Electrical Packaging and High-Speed Testing

A. Electrical Packaging Design

A photograph of the COB packaging is shown in Fig. 12. The flex-PCB provides connectivity to 207 bond pads on the CMOS chip, 64 of which are high-speed signal lines. A total of 48 pads are dedicated to ground or power connections. To minimize routing area, power and ground wirebonding rings are placed around the periphery of the chip. The rings are directly connected to their respective copper planes through multiple vias. The use of ground and power rings reduces the total number of wirebonding fingers down to 159, allowing for all signal traces to be routed using a four-layer flex-PCB in an area smaller than 44 $\times$ 44 mm^2.

All four copper layers (signal, power, ground, signal) are 0.5 oz/ft^2 (17 μm) in thickness and are separated by 3 mils (76 μm) of kapton ($\varepsilon = 3.7$). The outside copper layers are

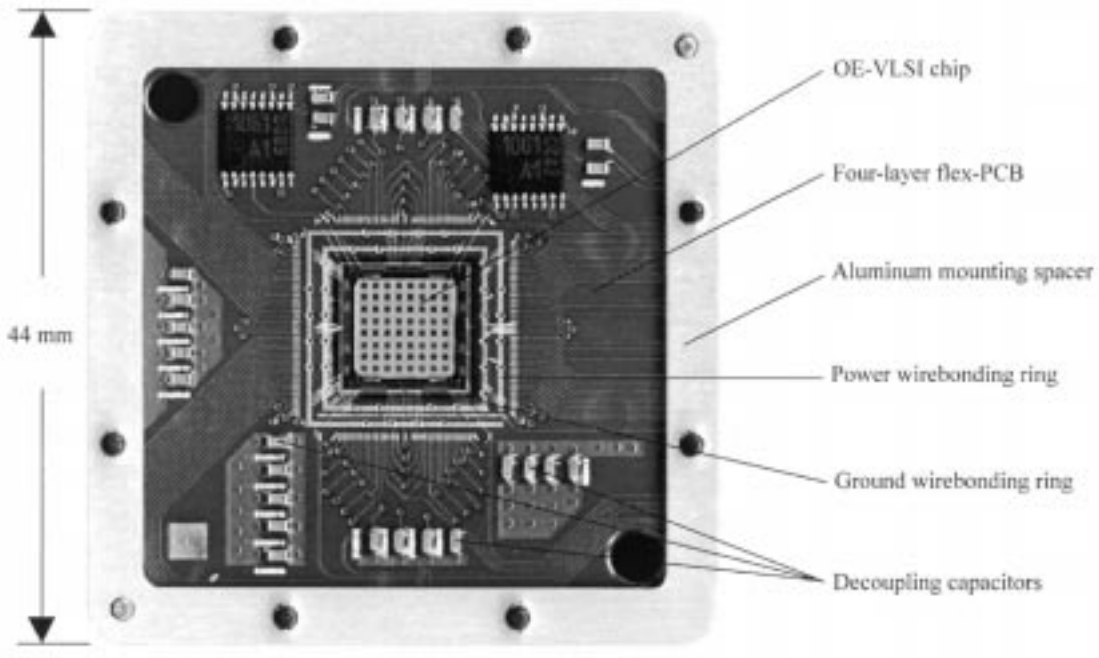

Fig. 12. Photograph of the COB assembly. Shown is the front view of chip module with the minilens array removed.

coated with 2 mils (50.8 μm) of kapton; this results in a total board thickness of 15.8 mils (400 μm). Signal traces are 5.5 mils (140 μm) in width and are placed at a minimum pitch of 10.5 mils (267 μm). Given nominal material thicknesses and dielectric constant, the calculated impedance for this coated microstrip stackup is 47.3 Ω and the propagation delay is 5.2 ps/mm [33].

The electrical packaging has been designed for high-speed operation ($>$100 Mb/s), with fall and rise times shorter than 2 ns. The nominal physical length of a signal trace between the chip and the motherboard is 180 mm. This corresponds to about half the effective length of a 2-ns rising edge; this means the flex-PCB must be treated as a distributed system requiring proper line terminations for both input and output signals. The choice of the line termination scheme is an important design issue. To minimize signal reflections and maximize switching speed, one would like to terminate high-speed signal lines with 50-Ω resistors. For our implementation, however, load termination was unsuitable for the following reasons. Input signals to the chip could not be load terminated because this required a large number of external 50-Ω resistors to be placed close to the chip, resulting in a prohibitively large footprint. Output signals from the chip could not be terminated with 50-Ω loads because this would have exceeded the current driving capability of the CMOS output pad drivers, not to mention the thermal problems associated with every output driver generating an average power of 0.25 W. For these reasons, series terminations were used for both inputs and outputs to the chip. Inputs to the chip were series terminated by placing a 33-Ω series resistance at the output of the motherboard driving electronics, this value corresponding to the difference between the line impedance and the nominal output resistance of the driving electronics. On the other hand, outputs from the chip did not require the use of external matching resistors because the output resistance of the CMOS output driver was already slightly larger than the line impedance.

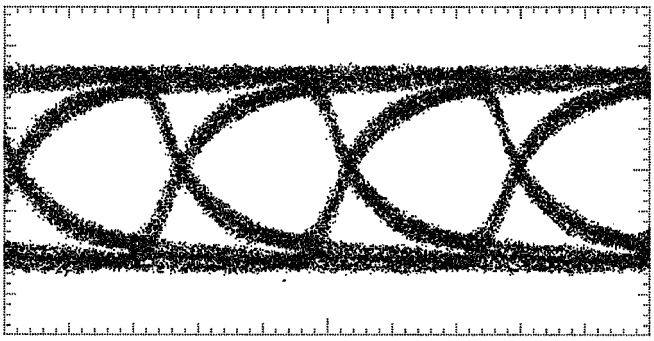

Fig. 13. Eye diagram at 1.0 Gb/s modulation with (pseudorandom binary sequence $2^{23}-1$).

An important signal integrity issue is related to the noise created by high supply current switching transients, often referred to as simultaneous switching noise (SSN) [34]. This noise is caused by a rapid change in current consumption of the circuit (due to many CMOS output drivers switching simultaneously), combined with the effective serial inductance of the power distribution networks. For instance, 10 CMOS output pads driving 10-pF loads with a 5-V swing and 1-ns rising edge will generate over 1 V of SSN across a typical 2-nH wirebond [35]. SNN can be minimized by reducing the effective serial inductance between the output drivers and the power and ground planes. This was done in the following ways: 1) by maintaining a large number of ground and power bond pads on the chip, 2) by using multiple vias to connect the ground and power wirebonding rings to their respective copper planes, and 3) by placing many decoupling capacitors in close proximity to the chip, thereby providing a low impedance path for high-frequency switching

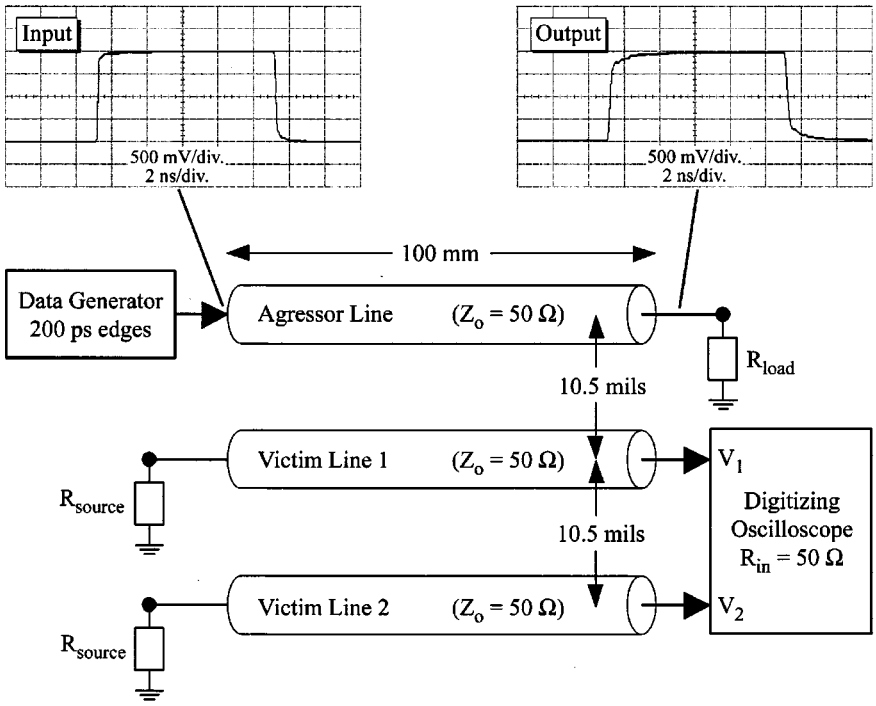

Fig. 14. Experimental setup used to characterize forward and reverse crosstalk between adjacent microstrip lines.

currents. The effective serial inductance of decoupling capacitors was kept to a minimum by selecting a small body size package and mounting them using multiple vias.

B. High-Speed Testing

Individual trace performance was measured using the time-domain reflectometry (TDR) feature of a digital oscilloscope. Results from TDR measurements performed on three traces show an average characteristic impedance of 53.1 Ω with a worst case peak deviation of 6.4%. The signal propagation delay, measured from the motherboard connector to the chip, was found to be 885 ps. This compares favorably with the calculated value of 936 ps (180 mm × 5.2 ps/mm).

To validate our design, a flex-PCB assembly was high-speed-tested as follows. First, two flex-PCB traces were shorted together by connecting their wirebonding fingers using a short wirebond. This arrangement allowed for a high-speed data signal coming from the motherboard to travel down one trace along the flex-PCB and make its way back to the motherboard on the other trace. The motherboard was equipped with a pair of impedance-matched microcoax connectors [voltage standing-wave ratio (VSWR) of 1.2 at 2.0 GHz]; one connected to a high-speed data source, the other connected to the 50 -Ω input of a digital oscilloscope. An eye diagram showing 500 mV nonreturn-to-zero (NRZ) modulation at 1.0 Gb/s is shown in Fig. 13. The long tail at the rising and falling edges are typical of a skin-effect limited transmission line [36].

C. Crosstalk Measurements

Considering the large number of closely packed high-speed lines on the flex-PCB, an important signal integrity issue is the amount of crosstalk between adjacent lines. Here, we define crosstalk as the ratio of the induced voltage amplitude to the driving voltage amplitude. As a voltage pulse propagates down a line, it generates both forward and reverse propagating crosstalk on adjacent lines. In general, for traces above a ground plane, the inductive and capacitive components of reverse crosstalk are approximately equal, have the same polarity and, therefore, reinforce. The amount of reverse crosstalk can be calculated using [33]

$$\text{reverse crosstalk} \approx \frac{1}{1+(D/H)^2} \tag{1}$$

where D is the center-to-center spacing between two lines and H is the trace height above the ground plane. For this design, the minimum spacing $D = 10.5$ mils and $H = 3$ mils, resulting in a worst case theoretical reverse crosstalk of 7.5% between nearest neighbor lines. Unlike reverse crosstalk, the inductive and capacitive components of forward crosstalk are of opposite polarity and, therefore, tend to cancel. For microstrip lines, however, most of the electric field travels through air instead of through the dielectric; this reduces the capacitive component of forward crosstalk and usually results in small negative forward crosstalk [33].

The amount of forward and reverse crosstalk was characterized using dedicated microstrip lines on the flex-PCB. As illustrated in Fig. 14, a 2-V pulse with 200-ps rising and falling edges was applied to the near end of an agressor line and the amplitudes of the induced voltages, V_1 and V_2, were measured at the far end of victim lines 1 and 2, respectively. Crosstalk for other remote lines was negligible. The length of the interfering microstrip lines was 100 mm. Measurements were performed under different termination conditions. First, the forward crosstalk component was determined by using 50-Ω terminations at both ends of the agressor and victim

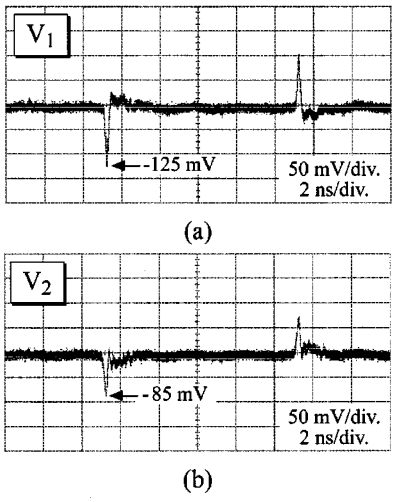

Fig. 15. Crosstalk response measured on both victim lines with $R_{\text{load}} = 50\ \Omega$ and $R_{\text{source}} = 50\ \Omega$. In this case, the response is mainly due to negative forward crosstalk.

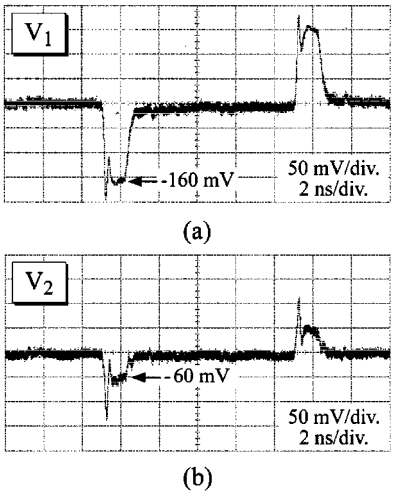

Fig. 16. Crosstalk response measured on (a) victim line 1 and (b) victim line 2. Line terminations are $R_{\text{load}} = 50\ \Omega$ and $R_{\text{source}} = 0\ \Omega$. In this case, the response is a superposition of negative forward crosstalk and inverted reverse crosstalk.

lines ($R_{\text{load}} = 50\ \Omega$ and $R_{\text{source}} = 50\ \Omega$). This arrangement eliminates the near-end reflection of reverse crosstalk. This way, the crosstalk response measured at the far end is mainly due to forward crosstalk. As shown in Fig. 15, forward crosstalk is negative; it appears as a narrow negative peak on a rising edge, the width of the peak being approximately equal to the rise time.

Next, the near end of the victim lines were shorted to ground ($R_{\text{load}} = 50\ \Omega$ and $R_{\text{source}} = 0\ \Omega$). This arrangement provides total reflection of the reverse crosstalk at the near end and results in the superposition of both forward and reverse crosstalk at the far end. As shown in Fig. 16, reverse crosstalk appears as a wide negative pulse at the far end, the width of the pulse being approximately equal to twice the delay of the microstrip line. Reverse crosstalk is positive; it appears as a negative voltage at the far end due to the negative reflection at the near end. Reverse crosstalk amplitudes of -160 and -60 mV were measured on victim lines 1 and 2, respectively, which translates into 8.0% and 3.0% reverse crosstalk, in good agreement with the predictions of (1).

Finally, the near end of the victim lines were 50-Ω terminated while the far end of the agressor line was open-circuited ($R_{\text{load}} =$ open-circuit and $R_{\text{source}} = 50\ \Omega$). This corresponds to the actual termination scheme used in our implementation. In this arrangement, a 1-V pulse is applied to the agressor line; the line is charged up to 2-V upon reflection at the far-end open circuit. In this case, the far-end crosstalk response will be a superposition of i) negative forward crosstalk generated by the forward-propagating 1-V pulse and ii) positive reverse crosstalk generated by the backward-propagating 1-V pulse. The result is shown in Fig. 17. Note that, in this case, the forward and reverse crosstalk components are of opposite sign and, therefore, cancel each other; this is why the forward crosstalk peak of Fig. 15 does not show up in the crosstalk response of Fig. 17. Reverse crosstalk amplitudes of $+80$ and $+30$ mV were measured on victim lines 1 and 2, which is consistent with the results of Fig. 16 and the fact that the driving pulse amplitude is now one half of 2 V.

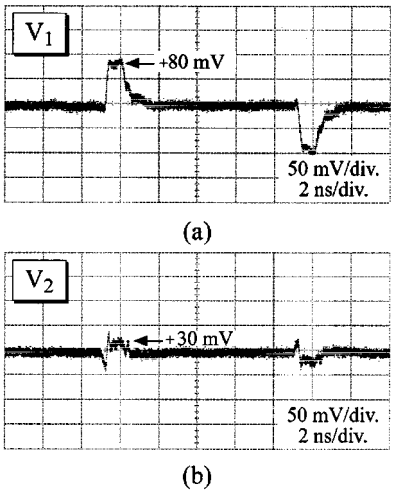

Fig. 17. Crosstalk response measured on (a) victim line 1 and (b) victim line 2. Line terminations are $R_{\text{load}} =$ open-circuit and $R_{\text{source}} = 50\ \Omega$. In this case, the response is a superposition of negative forward crosstalk and positive reverse crosstalk.

These results indicate that 5-V NRZ data propagating down a line will generate a $+200$-mV crosstalk pulse on its nearest lines and $+75$ mV on the next, the width of the pulse being equal to twice the line delay. This translates into 4.0% and 1.5%, respectively, which is tolerable. In a worst case situation, however, crosstalk components originating from nearby lines add together. In this case, rising edges occurring simultaneously on multiple lines on both sides of a victim line can generate as much as 550 mV of crosstalk, which is significant.

VII. Thermal Design and Experimental Evaluation

The MQW modulators were designed to provide a maximum ΔR of 55% when operated at a wavelength of 852 nm and a temperature of 40 °C. For ΔR not to fall below 50%, the chip temperature must be stabilized to 40 ± 5 °C. Temperature stabilization is achieved with ± 0.01 °C precision by using a TEC in combination with a precision miniprobe thermistor in a closed-loop

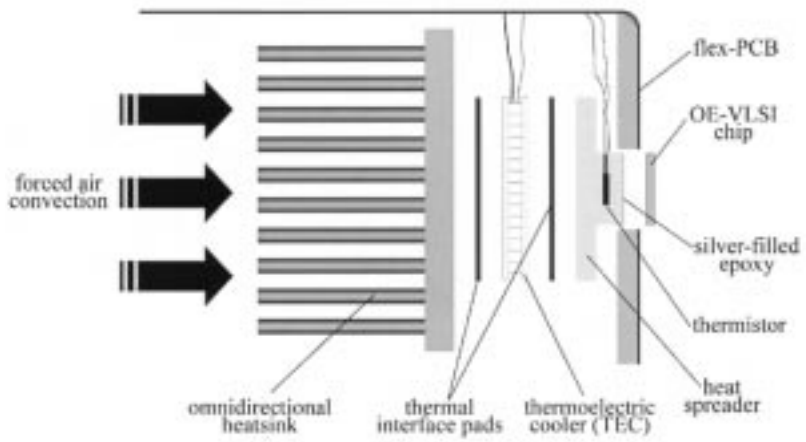

Fig. 18. Thermal components of the chip module.

feedback configuration. The thermistor probe (1.0 mm in diameter, 5.0 mm long) is inserted and fixed using thermal epoxy into a small hole drilled on the side of the copper heat spreader pedestal, just below the chip. The thermal resistance between the probe and the chip surface must be minimized to ensure that the thermistor temperature reading is close to the actual chip temperature.

HSPICE simulations were performed to estimate the worst case average power dissipation of the CMOS chip. Results indicate a worst case dissipation of about 7.5 W, of which 3.5 W is contributed solely from the bond pad drivers. Given this thermal load, minimizing the package thermal resistance is essential in order to avoid the need for multistage TECs or heatsinks that exceed the physical design constraints. The junction-to-TEC thermal resistance was minimized by mounting the chip directly onto a high copper alloy (C18500) heat spreader using a thin layer of silver-filled epoxy. This arrangement provides an excellent thermal path due to the high thermal conductivity of alloy C18500 (324 W/m/°C).

The thermal components of the chip module are shown in Fig. 18. A first-order thermal network model was developed by calculating the thermal resistance of each component in the thermal path. Results appear in Table II, except for the heatsink thermal resistance, which was not calculated as it depends on the fan speed and position. To validate the model, a calorimetric assembly was built by replacing the silicon chip with a 9×9 mm^2 uniform heating element and substituting the TEC with a component of identical size and known thermal resistance. Next, thermistor miniprobes were inserted into small holes drilled at specific locations in the assembly, allowing for simultaneous temperature measurements to be performed at multiple points along the thermal path. In addition, the front of the assembly was thermally insulated to ensure that all the heat would flow through the heat spreader and not be lost to the environment through natural convection. The heating element was set to a known amount of power dissipation and the steady-state temperature data at each thermistor was recorded. This step was repeated for several power settings of the heating element. For a given power setting, knowledge of the temperature drops across each component allows for each thermal resistance to be determined and the results appear in the second column of Table II. The measured thermal resistance of the heat spreader and silver-filled epoxy layer are in good agreement with theory. The thermal resistance of the interface pads, however, was significantly lower than what was specified on the manufacturer's data sheet. This discrepancy is due to the high contact pressure present when the heatsink is clamped to the heat spreader and the fact that thermal resistance decreases with increasing pressure. These results indicate that the junction-to-TEC thermal resistance is approximately equal to 0.4 °C/W.

TABLE II
MODULE THERMAL ANALYSIS AND EXPERIMENTAL CHARACTERIZATION

	R_{theory} (°C/W)	$R_{experimental}$ (°C/W)		Temperature (°C) Q_{chip} = 13.1 W[f]
$R_{heatsink}$	1.2[a]	N/A	$T_{TEC\text{-}hot}$	60.8
R_{pad}	0.17[b]	0.05	$T_{TEC\text{-}cold}$	37.7
$R_{spreader}$	0.24[c]	0.21	$T_{top\text{-}spreader}$	41.1
R_{epoxy}	0.06[d]	0.04	$T_{bottom\text{-}chip}$	41.6
R_{chip}	0.08[e]	N/A	$T_{top\text{-}chip}$	42.6

[a] From manufacturer's data sheet. Assumes 3 m/s (600 lfm) air flow.
[b] From manufacturer's data sheet. Assumes no contact pressure.
[c] Spreading resistance estimated using Kennedy's plots [25].
[d] Assumes uniform 10 μm epoxy layer and k_{epoxy} = 2.0 mW/°C.
[e] Silicon chip thickness = 675 μm and $k_{silicon}$ = 98 Wm/°C [36].
[f] A chip thermal load of Q_{chip} = 13.1 W required Q_{TEC} = 17.7 W.

Another important quantity to determine is the maximum thermal load that this design can dissipate under forced-air convection. To do this, the TEC was reinserted into the assembly and the temperature of the heat spreader thermistor was stabilized at 40 °C. The power generated by the heating element was slowly increased and eventually reached a point where the TEC was unable to pump additional heat out of the module while stabilizing the thermistor at the set temperature. This point was experimentally determined to be 13.1 W, and the corresponding temperature values along the thermal path are given in Table II. This result confirms that this design can easily withstand the worst case thermal load produced by the chip.

VIII. Conclusion

The design, implementation and testing of a high-performance chip module accommodating a 32×32 array of MQW modulators flip-chip bonded to a 9×9 mm^2 VLSI chip was described. The module integrates a minilens array, a copper heat spreader, a thermoelectric cooler, and an aluminum heatsink. The minilens array was aligned and packaged with the chip using a novel technique which provides high alignment sensitivity in all six DOFs. The key features of the design are as follows.

- The chip module assembly is simple and modular. Components are joined to one another using dowel pins and screws. A module can be assembled in a few minutes.
- The optomechanical design implements a semikinematic interface between the module and the interconnect optics, allowing for the module to be manually inserted into a free-space optical backplane with a standard deviation in spot misalignment of $\sigma = 2.2$ μm. In our demonstrator system, the OE-VLSI chips could be manually replaced with no need for further adjustments.
- The chip is mounted directly on a four-layer flex-PCB using a custom COB approach, providing 207 connections to the OE-VLSI chip in an area of 44×44 mm^2.
- The electrical design uses series terminations for both inputs and output data paths. Various techniques are used to minimize the amount of SSN.
- Impedance-controlled lines can be operated at 1.0 Gb/s with an open eye diagram. The measured crosstalk between nearest neighbor lines is 4.0%.
- The junction-to-TEC thermal resistance is 0.4 °C/W. This allows for the use of a single-stage TEC to regulate the chip at an operating temperature of 40 °C under a maximum thermal load of 13.1 W.

In summary, this work simultaneously addresses the issues of mechanical alignment, electrical signal integrity, and thermal dissipation. An integrated packaging solution is obtained by considering the large set of design issues throughout the development stage.

Acknowledgment

The authors would like to thank F. A. P. Tooley from Terahertz Photonics Ltd., Livingston, Scotland, U.K., for constructive comments and discussions and R. Adams from the Photonic Systems Group at McGill University, Montréal, QC, Canada, for his assistance in collecting the data for Fig. 11.

References

[1] J. W. Goodman, F. J. Leonberger, S.-Y. Kumg, and R. A. Athale, "Optical Interconnections for VLSI Systems," *Proc. IEEE*, vol. 72, pp. 850–866, 1984.

[2] F. A. P. Tooley, "Challenges in optically interconnecting electronics," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 3–13, Apr. 1996.

[3] A. V. Krishnamoorthy, L. M. F. Chirovsky, W. S. Hobson, R. E. Leibengath, S. P. Hui, G. J. Zydzik, 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, pp. 128–130, Jan. 1999.

[4] D. V. Plant, J. A. Trezza, M. B. Venditti, E. Laprise, J. Faucher, K. Razavi, M. Châteauneuf, A. G. Kirk, and W. Luo, "A 256 channel bi-directional optical interconnect using VCSEL's and photodiodes on CMOS," in *Proc. Optics Computing 2000 Conf.*, Québec, Canada, June 2000, pp. 1046–1054.

[5] A. V. Krishnamoorthy and K. W. Goossen, "Optoelectronic-VLSI: Photonics integrated with VLSI circuits," *IEEE J. Select. Topics Quantum Electron.*, vol. 4, pp. 899–912, Nov./Dec. 1998.

[6] J. A. Trezza, J. S. Powell, C. Garvin, K. Kang, and R. Stack, "Large format smart pixel arrays and their applications," in *Proc. 1998 IEEE Aerospace Conf.*, vol. 5, Aspen, CO, Mar. 1998, pp. 299–310.

[7] A. Erlich, "Micro-optical integration spurs mass production," *Laser Focus World*, vol. 34, pp. 77–81, 1998.

[8] F. A. P. Tooley, "Optical interconnects do not require improved optoelectronic devices," in *Proc. 1998 Optics Computing 1998 Conf.*, Brugge, Belgium, June 1998, pp. 14–17.

[9] T. H. Szymanski and H. S. Hinton, "A reconfigurable intelligent optical backplane for parallel computing and communications," *Appl. Opt.*, vol. 35, pp. 1253–1268, 1996.

[10] M. H. Ayliffe, D. Kabal, F. Lacroix, E. Bernier, P. Khurana, A. G. Kirk, F. A. P. Tooley, and D. V. Plant, "Electrical, thermal and optomechanical packaging of large 2D optoelectronic device arrays for free-space optical interconnects," *J. Opt. A: Pure Appl. Opt.*, vol. 1, pp. 267–271, 1999.

[11] D. Z. Tsang and T. J. Goblick, "Free-space optical interconnect technology in parallel processing systems," *Opt. Eng.*, vol. 33, pp. 1524–1531, 1994.

[12] T. Sakano, T. Matsumoto, and K. Noguchi, "Three-dimensional board-to-board free-space optical interconnects and their application to the prototype multiprocessor system: COSINE-III," *Appl. Opt.*, vol. 34, pp. 1815–1822, 1995.

[13] D. J. Goodwill, D. Kabal, and P. Palacharla, "Free-space optical interconnect at 1.25 Gb/s/channel using adaptive alignment," in *Proc. Optical Fiber Communication Conf.*, vol. 2, San Diego, CA, Feb. 1999, pp. 259–261.

[14] X. Zheng, P. J. March, D. Huang, O. Kibar, N. S. E. Ozkan, and S. C. Esener, "Optomechanical design and characterization of a printed-circuit-board-based free-space optical interconnect package," *Appl. Opt.*, vol. 38, pp. 5631–5640, 1999.

[15] K. Hirabayashi, T. Yamamoto, S. Hino, Y. Kohama, and K. Tateno, "Optical beam direction compensating system for board-to-board free-space optical interconnection in high-capacity ATM switch," *J. Lightwave Technol.*, vol. 15, pp. 874–882, May 1997.

[16] D. V. Plant, E. Bernier, E. Bisaillon, M. Mony, M. Salzberg, T. Yamamoto, D. J. Goodwill, and A. G. Kirk, "A 5 Gb/s, 2-channel bi-directional adaptive redundant FSOI demonstrator system," in *Proc. Optics Computing 2000 Conf.*, Québec City, QC, Canada, June 2000, pp. 465–472.

[17] D. V. Plant, B. Robertson, H. S. Hinton, M. H. Ayliffe, G. C. Boisset, W. Hsiao, D. Kabal, N. H. Kim, Y. S. Liu, M. R. Otazo, D. Pavlasek, A. Z. Shang, J. Simmons, K. Song, D. A. Thompson, and W. M. Robertson, "A 4×4 VCSEL/MSM optical backplane demonstrator system," *Appl. Opt.*, vol. 35, pp. 6365–6368, 1996.

[18] G. C. Boisset, M. H. Ayliffe, B. Robertson, R. Iyer, Y. S. Liu, D. V. Plant, D. J. Goodwill, D. Kabal, and D. Pavlasek, "Optomechanics for a four-stage hybrid-self-electro-optic-device-based free-space optical backplane," *Appl. Opt.*, vol. 36, pp. 7341–7358, 1997.

[19] F. B. McCormick, T. J. Cloonan, A. L. Lentine, J. M. Sasian, R. L. Morrison, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, R. J. Crisci, R. A. Novotny, and H. S. Hinton, "Five-stage free-space optical switching network field-effect transistor self-electro-optic-effect-device smart-pixel-arrays," *Appl. Opt.*, vol. 33, pp. 1601–1618, 1994.

[20] B. Robertson, "Design of an optical interconnect for photonic backplane applications," *Appl. Opt.*, vol. 37, pp. 2974–2984, 1998.

[21] D. R. Rolston, B. Robertson, H. S. Hinton, and D. V. Plant, "Analysis of a microchannel interconnect based on the clustering of smart-pixel device windows," *Appl. Opt.*, pp. 1220–1233, 1996.

[22] M. H. Ayliffe, "Alignment and packaging techniques for two-dimensional free-space optical interconnects," Ph.D. dissertation, McGill Univ., Montréal, QC, Canada, 2001. Chapter 7.

[23] F. Lacroix, B. Robertson, M. H. Ayliffe, E. Bernier, F. A. P. Tooley, M. Châteauneuf, D. V. Plant, and A. G. Kirk, "Design and implementation of a four-stage clustered free-space optical interconnect," in *Proc. 1998 Optics Computing Conf.*, Brugge, Belgium, June 1998, pp. 107–110.

[24] E. Oberg, F. D. Jones, H. L. Horton, and H. H. Ryffel, *Machinery's Handbook*, 25th ed. New York: Industrial, 1996.

[25] R. R. Tummala and E. J. Rymaszewski, *Microelectronics Packaging Handbook*. New York: Van Nostrand Reinhold, 1989.

[26] D. A. B. Miller, D. S. Chemla, T. C. Damen, A. C. Gossard, W. Wiegmann, T. H. Wood, and C. A. Burrus, "Band-edge electroabsorption in quantum well structures: The quantum-confined Stark effect," *Phys. Rev. Lett.*, vol. 53, pp. 2173–2176, 1984.
[27] Information Service in Physics, Electrotechnology, and Control, *Properties of Gallium Arsenide*, 2nd ed, ser. EMIS Data Reviews. London, U.K.: INSPEC, 1990, vol. 2.
[28] M. B. Venditti, "Temperature dependence of QCSE modulator and detector efficiency for free-space optical interconnect applications," M. Eng. thesis, McGill Univ., Montréal, QC, Canada, Sept. 1999.
[29] N. R. Basavanhally, M. F. Brady, and D. B. Buchholz, "Optoelectronic packaging of two-dimensional surface active devices," *IEEE Trans. Comp. Packag. Manufact. Technol. B*, vol. 19, pp. 107–115, 1996.
[30] G. C. Boisset, B. Robertson, W. S. Hsiao, M. R. Taghizadeh, J. Simmons, K. Song, M. Matin, D. A. Thompson, and D. V. Plant, "On-die diffractive alignment structures for packaging of microlens arrays with 2-D optoelectronic device arrays," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 918–920, July 1996.
[31] F. Lacroix, "Design, analysis and implementation of free-space optical interconnects," Ph.D. dissertation, McGill Univ., Montréal, QC, Canada, 2000.
[32] D. F-Brosseau, F. Lacroix, M. H. Ayliffe, E. Bernier, B. Robertson, F. A. P. Tooley, D. V. Plant, and A. G. Kirk, "Design, implementation, and characterization of a kinematically aligned, cascaded spot-array generator for a modulator-based free-space optical interconnect," *Appl. Opt.*, vol. 39, pp. 733–745, 2000.
[33] H. W. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
[34] R. Senthinathan and J. L. Prince, "Simultaneous switching noise calculation for packaged CMOS devices," *IEEE J. Solid-State Circuits*, vol. 26, pp. 1724–1728, 1991.
[35] S. Gong, H. Hentzell, S.-T. Persson, H. Hesselbom, B. Lofstedt, and M. Hansen, "Packaging impacts on switching noise in high-speed digital systems," *IEE Proc.-Circuits Devices Syst.*, vol. 145, pp. 446–452, December 1998.
[36] R. L. Wigington and N. S. Nahman, "Transient Analysis of Coaxial Cables Considering Skin Effect," *Proc. IRE*, vol. 45, pp. 166–174, 1957.
[37] Y. C. Lee, H. T. Ghaffari, and J. M. Segelken, "Internal thermal resistance of a multi-chip packaging design for VLSI-based systems," in *Proc. 38th Electronics Computing Conf. (ECC)*, 1988, pp. 293–301.

Michael H. Ayliffe received the B.Eng. and Ph.D. degrees in electrical engineering from McGill University, Montréal, QC, Canada, in 1996 and 2001, respectively.

His Ph.D. dissertation focused on the development of alignment and packaging techniques for 2-D free-space optical interconnects. His current research activities are in the areas of microoptics, optoelectronic VLSI, high-speed packaging, and very short reach parallel optical interconnects. His interests also include all-optical networking and dense wavelength division multiplexing technologies. In 2000 and 2001, he provided industrial training to a variety of major photonics companies. He is now with Picolight Inc., Boulder, CO, where he works on optical design and packaging of low-cost multigigabit vertical cavity surface emitting laser-based transceivers.

David R. Rolston (S'92–M'99) was born in Ottawa, Canada. He received the B.Eng., M.Eng., and Ph.D. degrees in electrical engineering from McGill University, Montréal, QC, Canada.

His Master's thesis was an analysis of optical optoelectronic device placement in a 2-D optical interconnect and his Ph.D. dissertation was on the design and testing of optoelectronic VLSI chips for free-space optical interconnects. He currently works for FCI Canada Incorporated, Montréal, a division of FCI, a worldwide electronic connector company. His primary goal is to help develop a photonic research and development team within FCI in Montréal. His research interests include computing and switching architectures applied to 2-D parallel optical interconects, as well as the development of high-speed optical transceivers and high-speed optical switching. He is currently working on low-cost optical transceiver development and long-term research into optical backplane technology.

Dr. Rolston is a member of the Optical Society of America (OSA).

Alan E. L. Chuah (S'99) was born in Penang, Malaysia, in 1976. He received the B. Eng. degree in electrical engineering from McGill University, Montréal, QC, Canada, in 1999. He is currently pursuing the Master's degree in electrical engineering with the Photonic Systems Group at McGill University.

His current research is on the design and implementation of high-speed transmitters and receivers for optical input–output in CMOS technology. His research interests include high-speed ICs and PCB design, as well as signal integrity problems.

Eric Bernier (S'95–M'00) was born in Noranda, QC, Canada, in December 1974. He received the B. Eng. degree in electrical engineering and the M. Eng. degree in photonics from McGill University, Montréal, QC, Canada, in 1997 and 2000, respectively.

He is now with Nortel Network, Ottawa, ON, Canada, where he works on advanced optical technology. His research interests include all-optical network, dense wavelength division multiplexing optoelectronic technology, signal quality analysis, and advanced interconnect technology.

Feras S. J. Michael received the B. Eng. and M. Eng. degrees in electrical and computer engineering from McGill University, Montréal, QC, Canada, in 1998 and 2000, respectively.

His M. Eng. research investigated the feasibility of displacing a lenslet array using a novel multiaxis active alignment microelectromechanical system. He is currently a digital hardware and optical design engineer at Xanoptics Inc., Merrimack, NH.

David Kabal (S'96–M'97) was born in Toronto, ON, Canada, in 1972. He received the B. Eng. and M. Eng. degrees in electrical engineering from McGill University in Montréal, QC, Canada, in 1996 and 1998, respectively.

His M. Eng. thesis research was in the area of packaging of 2-D optoelectronic device arrays for free-space optical interconnects. In 1997, he joined Nortel Networks, Ottawa, ON, Canada, where he worked in optoelectronics and metropolitan area networks until September 2000. In October 2000, he joined Picolight Inc., Boulder, CO, where he is currently working as a Senior Technology Planner in strategic marketing for optoelectronic transceiver modules.

Andrew G. Kirk (M'97) received the B.Sc. degree in physics from the University of Bristol, Bristol, U.K., in 1989 and the Ph.D. degree in physics from the University of London (King's College), London, U.K., in 1992.

Since 1986, he has been a Professor and member of the Photonic Systems Group in the Department of Electrical and Computer Engineering at McGill University, Montréal, QC, Canada. For the last three years, he has provided industrial training to a variety of major photonics companies. His current research interests are in the area of microoptics, optical microelectromechanical systems, and free-space optical interconnects. He has published a number of refereed papers in these fields and has collaborated actively with a variety of companies.

Dr. Kirk was appointed as William Dawson Scholar at McGill University in 2001. He also recently received the annual Teaching Award of the Electrical and Computer Engineering Department.

David V. Plant (S'86–M'89) received the Ph.D. degree in electrical engineering from Brown University, Providence, RI, in 1989.

From 1989 to 1993, he was a Research Engineer in the Department of Electrical and Computer Engineering at the University of California, Los Angeles (UCLA). Since 1993, he has been a Professor and member of the Photonic Systems Group within the Department of Electrical and Computer Engineering at McGill University, Montréal, QC, Canada. Since joining McGill University, he has been studying optical interconnects for very short reach applications. He is the Major Project Leader of the Photonic Interconnection Systems major project within the Canadian Institute for Telecommunications Research (CITR) and a Project Co-Leader within the Canadian Institute for Photonic Innovation (CIPI). During the 2000–2001 academic year, he took a leave of absence from McGill University to become the Director of Optical Integration at Accelight Networks, a startup company based in Pittsburgh, PA, and Ottawa, ON, Canada, which will deploy optical internetworking equipment.

Dr. Plant received the Outstanding Departmental Teaching Award and the Faculty of Engineering Teaching Award in 1996. In recognition of his accomplishments in research and teaching, he was named a James McGill Professor in 2001.