Implementation of a compact, four-stage, scalable optical interconnect for photonic backplane applications

Frederic Lacroix, Eric Bernier, Michael H. Ayliffe, Frank A. P. Tooley, David V. Plant, and Andrew G. Kirk

We report on the implementation of a dense 512-beam free-space optical interconnect linking four optoelectronic VLSI chips at the backplane level. The system presented maximizes the positioning tolerances of the components by use of slow f-number ($f/16$) Gaussian beams and oversized apertures. A beam-clustering scheme whereby a 4×4 array of beams is transmitted by each minilens is used to provide a high channel density. A modular approach is used to decrease the number of degrees of freedom in the system and achieve passive alignment of the modules in the final integration phase. A design overview as well as assembly and experimental results are presented.

OCIS codes: 200.4880, 200.4650, 200.2610.

1. Introduction

Free-space optical interconnects have been proposed to alleviate the limitations encountered in present electrical interconnection technology. The Semiconductor Industry Association (SIA) roadmap[1] forecasts dramatic increases in chip size, density, and output pin number in the coming years. For example, clock rates for new high-performance processors are forecast to reach 1600 MHz by 2002 and 3000 MHz by 2011. These rates show a future need to move large amounts of data at high speeds across processing elements. Interconnects have been identified most likely to limit the performance of future electronic processing systems.[2] Although metal-based interconnects are relatively cheap for short distances and are constantly being improved, it is not clear at this point whether they will support the SIA-predicted throughputs and data rates at competitive costs. Physical limitations inherent to the technology such as signal integrity issues (e.g., frequency-dependent cross talk and attenuation) limit the data transfer rates attainable with the use of metal interconnect lines.[3] Free-space optics represents an interesting alternative technology for implementing dense arrays of high-speed interconnections.

The most important obstacle preventing the commercial acceptance of free-space optical interconnects is the issue of misalignment.[4] Optical components routinely have to be precisely positioned in the micrometer range laterally and tenths of a degree angularly to achieve high coupling efficiencies between the source and the detector. The optomechanical packaging required to house the components tends to be both expensive and delicate. Although some research effort has been spent on alleviating the effects of misalignment through the use of redundant arrays of emitters and receivers,[5] active alignment techniques,[6] and error correction codes,[7] the problem has rarely been tackled at a more fundamental level; that is, few systems have been designed to maximize the intrinsic positioning tolerances of the components.

This last point touches on some important design considerations. A prerequisite to designing an optical interconnect is the ability to accurately model the propagation of coherent optical beams through a sometimes complex train of refractive and diffractive micro-optical components. Accurate modeling of optical power propagation is necessary to calculate proper mechanical positioning tolerances for the optical components as these tolerances are based on ensuring a certain minimum amount of power being

When this research was performed, F. Lacroix (e-mail address: fredl@photonics.ece.mcgill.ca), E. Bernier, M. H. Ayliffe, D. V. Plant, and A. G. Kirk were with the Department of Electrical and Computer Engineering, McGill University, 3480 rue Université, Montréal, Québec, Canada H3A 2A7. F. Lacroix is now with ONI Systems Corporation, 166 Baypointe Parkway, San Jose, California 95134. F. A. P. Tooley is with the Department of Physics, Heriot-Watt University, Riccarton, Edinburgh, Scotland.

Received 15 March 2001; revised manuscript received 26 November 2001.

0003-6935/02/081541-15$15.00/0

coupled from the source to the detector. However, this is no trivial task. Rigorous vector or scalar diffraction theory must be used to properly account for clipping and diffraction at apertures; however, this is often impractical when designing actual systems owing to the computationally intensive nature of these methods. Because the complexity of the calculations scales according to the square of the number of sampling points times the number of surfaces, the simulation and tolerancing of a complex optical train composed of many optical surfaces could require weeks or months of simulation for each design iteration. In practice, simpler and less accurate modeling methods such as ray tracing and Gaussian beam propagation are used. For further discussion of this topic, see Ref. 8. The misalignment and optomechanical problem is therefore compounded by a modeling problem: Accurate optomechanical supports must be designed to guarantee alignment; however, the misalignment tolerances are often not known precisely owing to modeling inaccuracy. Also, because so few free-space optical interconnect systems have been built, there exists no engineering rules of thumb to guide us in the tolerancing process; i.e., the safety margins that must be incorporated into a design to ensure a reasonable probability of success are not known. The way in which this problem was circumvented was to use a very severe tolerancing metric (1% power falloff with respect to the maximum) to set the tolerances on the components. It is hoped that the experience of this demonstrator serves the purpose of developing such engineering rules of thumb for free-space optical interconnects.

In addition to alignment precision, alignment stability also constitutes a challenge. Alignment must be maintained despite temperature variations, vibrations, and impacts. Field replaceability of the active and passive components is an issue. Optoelectronic (OE) chips are susceptible to failure and must be replaceable without disruption of system operation.

This paper reports on the implementation of a four-stage clustered optical interconnect for optical backplane applications. It is meant to provide a broad overview of issues encountered during the implementation phase of this system. This implementation seeks to test hypotheses and strategies of design, alignment, and packaging and demonstrate a series of desirable optomechanical properties such as passive alignment, stability, and insertion repeatability of the key components. The demonstration of these properties would constitute an important step in proving the practicality of these systems. The interconnect presented in this paper seeks to maximize misalignment tolerances at the design level by use of arrays of large (800 μm on the side square) and long focal length (8.5 mm) minilenses. The interconnection density is high owing to the use of a beam-clustering scheme in which many optical beams are grouped around the axis of each minilens.

The optical system is described first. The optomechanical strategies employed in the implementation of the system are then reviewed. Measurements of the experimental alignment accuracy and repeatability are also presented. Those results reflect the precision to which the system can be aligned and give an appreciation of the system performance when the OE chips are integrated into the system.

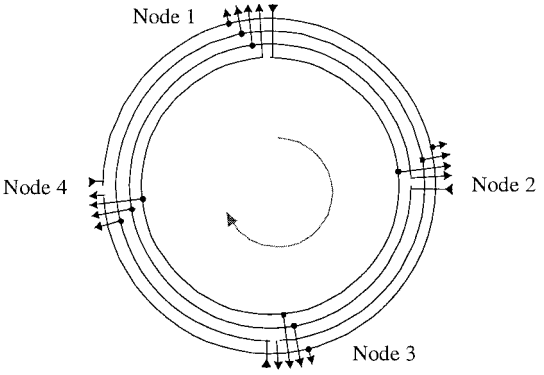

Fig. 1. Schematic diagram of the four-stage system.

2. System Overview

The system presented in this paper is a photonic backplane demonstrator implementing the hyperplane architecture.[9] It seeks to provide a large quantity of interconnections among printed circuit boards (PCBs) for packet-switching networks and massively parallel processing systems. Specifically, parallel electrical data in packet form is converted into parallel optical data and placed onto specific channels in the optical backplane. The optical data reach their destination (or destinations) and are then converted back into electrical data. Because address headers are incorporated into the packet, address recognition is performed at each OE chip, and the packet is routed to the correct PCB.

A four-stage (4 PCBs), ring-based configuration (the last PCB is linked optically to the first) was selected. The configuration of the system is shown on Fig. 1. Using a ring configuration is advantageous because the data that are transmitted from any node can be sent to any other node through multiple hops, even if the system is not bidirectional. This simplified the optical design as only a unidirectional interconnect needed to be implemented.

Figure 2 presents a schematic diagram of the interconnect design. The high-level optical design of the interconnect is presented in detail in Ref. 10. It implements a high-density point-to-point link between the modulators located at stage 1 and the detectors at stage 2. The modulators and detectors are composed of optically sensitive GaAs/AlGaAs p-i-[multiple quantum well (MQW)]-n diodes hybridized onto standard complementary metal-oxide semiconductors (CMOSs). Modulator devices were used as large arrays of vertical-cavity surface-emitting lasers were not available at the time the system was designed. A total of 512 optical beams (256 channels in that the data are encoded by use of dual-rail logic) arranged in 32 4 $\times$ 4 clusters are relayed through 8 $\times$

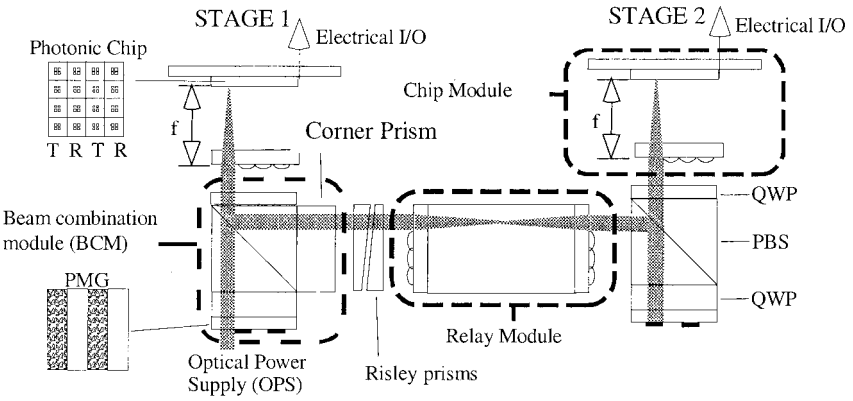

Fig. 2. Flattened layout of the optical system between the modulator and the detector stages.

8 arrays of large (800 μm) and long focal-length (8.5 mm) diffractive minilenses. Slow Gaussian beams ($f/16$) are used to maximize lateral and longitudinal tolerances. Telecentric imaging is used because it increases misalignment tolerances (defocus errors will not translate to lateral misalignments on the modulators or detectors because the chief ray is perpendicular to each plane). The interconnect uses a beam-combination method (BCM) based on the use of polarization optics. Polarization variations induced by quarter-wave (QWP) plates are used to route the beams between stages with the help of reflections off polarizing beam-splitter (PBS) cubes and mirrors located on the patterned-mirror–grating element (PMG) composed of alternating strips of diffractive fan-out grating and mirror.

The details of the optical power supply (OPS) used to generate the input 4×8 array of collimated beams are not shown here (see Ref. 11). A cascaded array-generation technique is used to produce the 512-beam array incident on the modulators: A 4×8 array is first generated by the diffractive fan-out present in the OPS, and each spot in this array is then fanned out by the PMG into a 4×4 subarray.

A total of 512 optical beams pass through the PBS–QWP assembly and are focused by a minilens array placed one focal length in front of the modulators. Those beams are then modulated and reflected back into the system where they pass through three minilens arrays arranged in a $4f$ relay configuration and reflect off stripes of mirror located on the PMG to be focused again on the detectors by another minilens array. Each Gaussian beam has a 13.1-μm $1/e^2$ focused spot radius at the modulator and detector planes. This 39.3-μm diameter should be well accommodated by the oversized 50-μm round modulators and the 70-μm square detectors. Note that the interconnect is actually of a three-dimensional nature: The function of the corner prism element is to deviate the beams by 90° with the help of total internal reflection. The Risley prisms are used for beam-steering purposes, to help correct possible minor misalignments of the beams traveling in the interconnect.

A clustered design employing arrays of diffractive minilenses in which many optical beams pass through the same minilens was employed to maximize misalignment tolerances while maintaining a high interconnection density. This scheme achieves both a high channel or window density (2500 windows/cm^2) and a moderate window size (40 μm). The smart-pixel arrays (also termed the OE-VLSI chips) are composed of optically sensitive GaAs/AlGaAs p-i-(MQW)-n diodes hybridized onto standard CMOS. The silicon CMOS chip incorporates logical circuitry permitting optical data to be transmitted, received, or retransmitted at each node.[12]

3. Modularization and Optomechanics

In implementing this system we sought to incorporate a number of desirable optomechanical properties listed below:

(i) Passive alignment of the optical modules within the supporting optomechanical structure. This helps to simplify system assembly.

(ii) Repeatable insertion and extraction of the modules, making them easily replaceable in case of failure (this is especially important for the OE chip module as it houses an active device).

(iii) Stability. The system stays aligned without needing adjustment once optimal alignment has been achieved.

The optical layer of the backplane is composed of four point-to-point optical interconnects (such as shown in Fig. 1) arranged in a unidirectional ring configuration. The optical system interconnecting the four OE-VLSI chips uses a total of 48 optical components, which occupy a 55 mm × 55 mm square area and are housed in a 7-cm side-square baseplate. All 48 components must be aligned relative to each other in all six degrees of freedom.

The alignment process can be simplified by grouping the components into pre-aligned modules. This reduces the number of degrees of freedom and relieves the tolerances on the optomechanics. A detailed tolerance analysis has been performed and the partitioning of the components into modules chosen

so as to maximize the system misalignment tolerances.

The misalignment tolerance analysis performed for this system is based on Gaussian beam propagation theory and therefore neglects diffraction and clipping effects at apertures. Gaussian beam propagation theory is often used for tolerancing free-space optical interconnects because of its simplicity. A criterion stating that a misalignment of an optical component in one degree of freedom (laterally, longitudinally, or in a roll, yaw, or pitch) must not cause more than a 1% power drop with respect to the maximum power throughput was used to set the tolerances for each degree of freedom for each component. A conservative metric of 1% was used because the effects of tolerance stackup are not known (the tolerances for each degree of freedom are not linearly independent as was assumed, so that coupling between multiple degrees of freedom might cause degradation in power throughput), and the effects of diffraction and clipping were neglected. Implementation constitutes a test of these tolerance-calculation hypotheses. The details of the misalignment-tolerance analysis can be found in Ref. 13.

The grouping of components is done such that the intramodule alignment requirements are more severe than the intermodule alignment requirements. However, note that severe intramodule alignment requirements are not critical as the modules can be assembled independently with active alignment techniques and then passively aligned in a baseplate. Modularization thus shifts the critical active alignment steps away from the system integration phase to the module assembly phase. In this way, it was intended to demonstrate the first optomechanical property (passive alignment).

The mechanical tolerances for the various components were specified in accordance with the calculated tolerance budget. The calculated tolerance values for the various modules are given in Subsections 3.A to 3.C and Subsection 4.A. Where possible, these mechanical tolerances were characterized upon receiving the components from the manufacturer.

The system has been separated into four modules: OPS, BCM, OE-VLSI chip module, and relay module. The partitioning has been chosen so as to produce misalignment-tolerant, easy-to-assemble, and compact modules. The module partitioning is indicated on Fig. 1.

The OPS module is a continuous-wave (cw) spot-array generator and is not described here.

The BCM is a polarization-based unit designed to combine and route three arrays of beams: the cw spot-array of beams incoming from the OPS and directed to the modulators, the modulated spot array reflected from the modulators and directed into the relay module, and finally the spot array incoming from a previous stage and directed to the detectors. It is composed of five elements: a PMG, two QWPs, a PBS, and an interface plate (IP).

The OE-VLSI chip module packages a minilens array with an OE chip. The minilens array is positioned one focal length away from the active devices. A complete description of this module and its assembly method can be found in Ref. 14.

The relay module is composed of an optical spacer (which is simply a block of high-index glass that provides mechanical support) and two minilens arrays. The minilens arrays are separated by a distance of two focal lengths in order to constitute a telecentric relay.

Once assembled, the modules must be integrated within a baseplate. The baseplate acts as the support structure for the optical system and provides a mechanical reference to which the modules can be aligned. All components and modules are attached to the baseplate or mounted into cells that rest on the baseplate.

The baseplate was fabricated in aluminum alloy 6061. Aluminum is a good material to use, as it is one of the easiest to machine and has a high strength-to-weight ratio, is easy to procure, is inexpensive, and has a proven environmental resistance. A standard computer numerical control machine was used to perform the machining. Note that a high level of precision for hole and slot locations can be guaranteed as long as the precision features are all located on the same side of the part to be machined.

Figure 3 shows an exploded view of the baseplate. It is organized in the following manner: Four through holes are machined at the corners and are linked by rectangular slots milled into the baseplate. The BCM is aligned and mounted within the through hole. The OPS and chip modules are then aligned with respect to and connect to the BCM on either side of the baseplate.

Precision ground rods are positioned at the bottom of the rectangular slots milled into the baseplate. These rods are used as support structures to position the components and modules, such as the corner prism, the Risley prisms, and the relay module, mounted into the baseplate. They define mechanical reference lines to which modules and components can be semikinematically aligned. The tolerance on the rod diameter is extremely severe (± 5 μm). This ensures that a minimum amount of tilt will be introduced when the components and modules are positioned. Such a scheme also has the advantage that it smoothes any irregularities present in the milled baseplate groove.

Note that the baseplate was designed to fit in an industry-standard U3 frame (13.97-cm wide and 37.465-cm deep). This is a self-imposed design criterion designed to show that the optical layer can be incorporated inside an electrical backplane chassis and is thus backward compatible with existing technology. Most optical parts were clamped in the baseplate with brass clamps. These clamps were designed to be flexible enough such that they would not damage the parts but sufficiently rigid to hold them in place. When the system alignment was completed, the optical parts were locked into position with UV curable adhesive. The OPSs were solidly

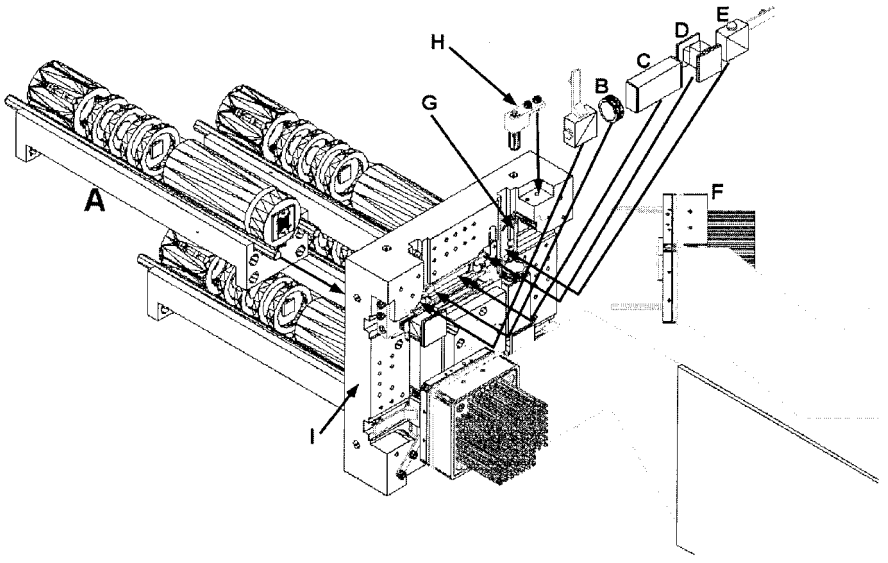

Fig. 3. Drawing of an exploded three-dimensional view of backplane. A, optical power supply; B, Risley prisms; C, relay module; D, BCM; E, corner prism; F, OE-VLSI chip module; G, hardened steel rod; H, adjustment screw; and I, baseplate.

fixed to the baseplate with two machine screws inserted in the back of the baseplate.

The modules' optomechanical alignment schemes are detailed below.

A. Beam-Combination Module

The BCM is housed in a fully kinematic mount; i.e., it should be possible to insert and remove the module repeatably to within or better than the required misalignment-tolerance limits (± 32 μm laterally and longitudinally, $\pm 0.04°$ angular). The angular tolerance is very severe. A method was required that would provide a precise angular reference relative to the baseplate. High-precision ruby ball lenses were used to provide a precise and uniform standoff distance with respect to the pads located at the bottom of the baseplate. These ball lenses provide contact points that define the height and tilt of the module with respect to the rest of the system. The diffractive optical elements located on the BCM are precisely diced and aligned with respect to dowel pins inserted into the baseplate to determine the lateral and angular position of the BCM. Figure 4 is a schematic diagram of the BCM mount.

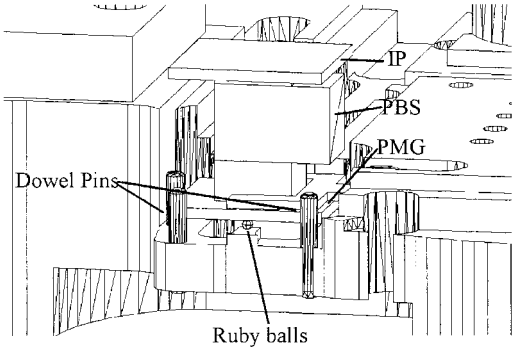

Fig. 4. Diagram of the BCM in a kinematic mount.

B. Relay Module

The relay module is positioned in a semikinematic mount. The longitudinal position of the module did not need to be constrained as the longitudinal misalignment tolerance of the module is quite relaxed (± 500-μm longitudinal tolerance compared with ± 40-μm lateral misalignment tolerance and $\pm 0.1°$ angular tolerance). It was decided to use high-precision rods as a mechanical alignment reference to which to position the module at the bottom of the baseplate. The minilens arrays located at each end of the block of glass rest on the rods. The dicing and alignment precision of the minilens arrays thus define the lateral and angular alignment precision of the module. Finely threaded screws are used to vary the position of the corner prism laterally (thus displacing the beams laterally with respect to the nominal optical axis). Risley prisms are included to help compensate for possible misalignments of the beams (they can therefore be of use if passive alignment fails). They are mounted in cells held in place by magnets. This mounting mechanism allows rotation of the Risley prisms. Figure 5 is a schematic diagram of the relay module alignment scheme.

C. Optoelectronic-VLSI Chip Module

The chip module packages the OE-VLSI chip with a minilens array. The misalignment tolerances for the modules are ± 26-μm lateral misalignment tolerances, ± 500-μm longitudinal misalignment tolerance, and $\pm 0.04°$ angular misalignment tolerance. The angular tolerance is critical. It was decided to trade off lateral alignment precision for angular alignment precision in the following way: Two precision dowel pins inserted in the baseplate mate to two slightly over-large alignment holes machined in

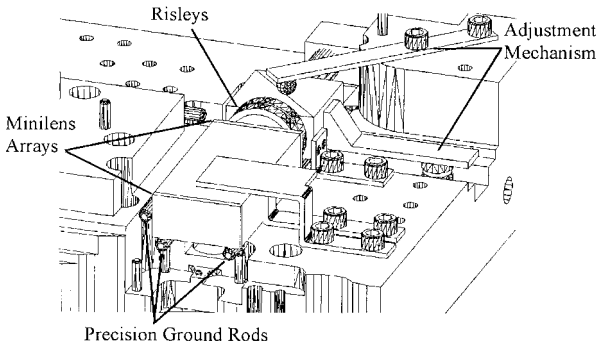

Fig. 5. Diagram of the relay module in a mount.

the module and thus define the lateral alignment precision. The OE-VLSI minilens array surface is brought into contact with the BCM IP surface, and the tilt of the chip module is adjusted perfectly to match the BCM by use of the play afforded by the oversized alignment holes on the modules. The scheme thus guarantees a high angular alignment precision (near perfect) with a reasonable lateral alignment precision (± 20 μm). Figure 6 is a schematic diagram of the optomechanical alignment scheme of the chip module.

4. Experimental Results

A. Module Assembly

The use of a modular assembly scheme meant that the critical alignment steps were shifted away from the system assembly phase to the module assembly phase. The module components' assembly tolerances thus tended to be extremely severe (e.g., approximately 5 μm laterally and a twentieth of a degree angularly for the BCM and 10 μm laterally for the relay module). Visual alignment under a microscope can achieve precisions of approximately 20 μm laterally, depending on the kind of alignment markers available, magnification, skill of the operator, etc. This is insufficient. Other methods had to be employed.

Two interferometric-based alignment techniques were employed to meet the severe lateral and angular specifications: One technique uses pairs of low-efficiency diffractive minilenses and gives a theoretical lateral alignment precision in the micrometer range, whereas the other technique exploits the fringes produced as a result of a Fabry–Perot effect between two uncoated angularly misaligned substrates and achieves an angular precision of as much as one wave over the area of the substrate (14 mm) or approximately 0.0035°. Interferometric techniques provide an excellent degree of alignment precision, far ahead of the precision usually attained through visual techniques. For example, a lateral precision of approximately 20 μm can be achieved through visual alignment compared with 5 μm or less through interferometric techniques.

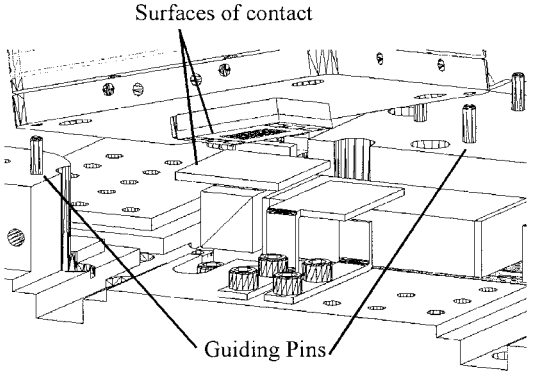

Fig. 6. Diagram of the VLSI chip module in a mount.

The first technique is described in Ref. 15. It uses pairs of extra-diffractive minilenses etched around the array of signal minilenses. These extra minilenses incorporate an etch depth error (they are designed to operate at 632.8 nm but are etched at 852 nm) to purposely render them inefficient at their design wavelength. This means that a significant undiffracted 0 order is present in the output beam. This 0 order recombines with the +1 order at the second minilens and forms an interference pattern. The number and shape of the fringes present give indications as to the type and amount of misalignment present.

The lateral resolution of this technique is calculated to be 6 μm for the BCM minilenses and 2.5 μm for the relay module minilenses (the resolution of this technique varies with the diameter of the minilens, which were different for the BCM and relay module). This is enough to meet the lateral tolerance specifications (± 18 μm). However, this technique possesses a longitudinal alignment resolution in the range of hundreds of micrometers. This is rather poor. This means that the technique cannot be used for tilt alignment by monitoring the differential longitudinal alignment between two minilenses. Another alignment technique had to be used.

Because the substrates are not antireflection coated, a beam propagating through them suffers multiple reflections between their respective surfaces, giving rise to Fabry–Perot fringes. Any tilt between the substrates will then result in a set of linear fringes appearing across the area of the substrate because the optical path length within the cavity varies spatially as a function of the tilt. Because the substrates are 14-mm square, the theoretical alignment resolution (one wave) by use of this technique is equal to 0.0035°. This is more than an order of magnitude below the required precision (0.05°).

The BCM, PMG, and IP had to be aligned and glued in-house to meet the severe misalignment tolerances (± 32 μm laterally and ± 0.05° angularly). The modules were assembled to precisions of 5 μm laterally and 0.03° angularly. Figure 7 is a picture of an assembled BCM. Note the compactness of the module.

The relay module was assembled following the same procedure used with the BCM. The relay modules were assembled to precisions of better than 5 μm

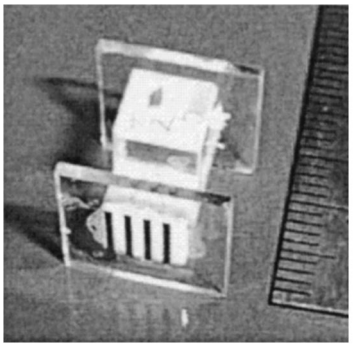

Fig. 7. Assembled BCM (scale in millimeters).

laterally. The tilt alignment precision of the minilenses was not critical in this case. Figure 8 is a picture of the assembled relay module.

In summary, the interferometric minilens alignment technique was found to be a simple and effective method to meet the severe lateral alignment specifications. However, this technique does not provide tilt alignment. This alignment has to be performed by monitoring the fringes present on the substrates that are due to Fabry–Perot effects. This technique is extremely sensitive as it possesses a resolution of one wave over the area of the substrates. A disadvantage is that the contrast between the fringes is poor owing to the low reflectivity of the substrates (4%). Alignment is thus slightly more delicate as it is more difficult to identify the fringes on the substrate. Increasing the reflectivity of the substrates is not desirable as it will induce greater optical losses.

B. Diagnostic Modules

Once the modules are assembled, they must be integrated into the baseplate. Alignment must be monitored to guarantee that a maximum of the optical power emitted by the source falls on the appropriate detector. However, the experimental data that can be collected are often limited owing to the difficulty of performing the measurements. The main problem is often simply a question of space constraints: The necessary viewing ports cannot be accessed easily. For example:

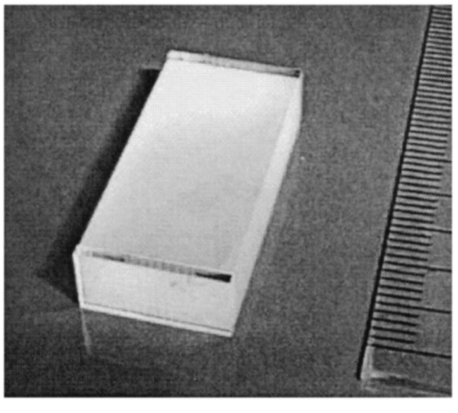

Fig. 8. Assembled relay module (scale in millimeters).

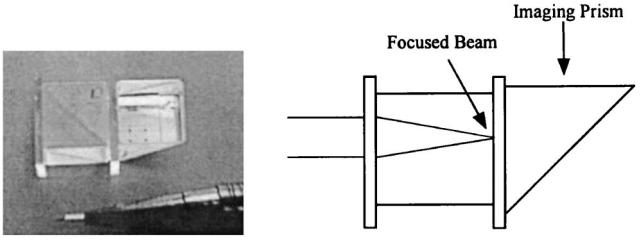

Fig. 9. Relay diagnostic module.

(1) The OE chips prevent direct access to the beams at the OE chip plane. This means that alignment and beam quality cannot be monitored at this critical plane.

(2) The system is very compact (~5 cm on the side). This means that an instrument such as a power meter or CCD camera is too bulky to position within the optical layer.

Ways have to be found to circumvent these problems. In particular, it is crucial to be able to image the OE chip plane to ensure that the focused spots actually fall on the modulators and detectors. Diagnostic testing thus constitutes a critical part of the design of an optical interconnect. Various diagnostic tools must be planned for and incorporated in the system design.

One way to overcome the first problem outlined above is to render the OE chip transparent so that through-the-chip imaging can be performed to monitor the spot alignment at the detector plane. A chip diagnostic module (CDM) combining a minilens array with a transparent substrate that contains alignment markers was constructed for this purpose. Of course, this module has to be interchangeable with the OE-VLSI chip module and its insertion repeatable for this scheme to function.

The second problem can be alleviated by inserting mirrors in the baseplate to allow imaging to be performed perpendicularly to the optical layer. A relay diagnostic module was constructed to perform this function.

The system could then be assembled and its alignment monitored without inserting the OE chips. Together these two modules can monitor the angular and lateral position of the beams at the OE chip plane and within the baseplate. They can then be removed when alignment is satisfactory and replaced by the appropriate optical modules.

The relay diagnostic module is shown in Fig. 9. It is designed to substitute for a relay module within the baseplate to diagnose the beam position and angle at the relay module position. The module is composed of a minilens array and of a patterned photomask that reproduces the features of the OE chip. The position of the focused beam on the patterned photomask allows the detection of any misalignment in the beams propagating in the interconnect. A beam splitter is added between the minilens and the photomask to act as a mechanical support. Note that

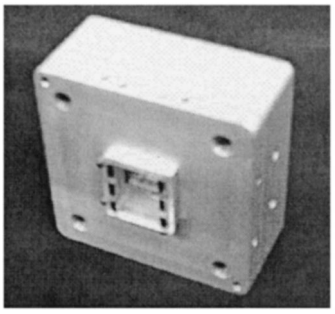

Fig. 10. Chip diagnostic module.

there is no special reason to use a beam splitter; it was used merely because it is a cube of high-quality optical glass of the appropriate dimensions and was readily available. A prism is glued at the back of the photomask to provide a viewing port so that the photomask surface can be imaged. The minilens array and the photomask are visually aligned relative to each other by monitoring of the position of the focused spots on the photomask surface. The alignment precision of the photomask relative to the minilens array is estimated to be 10 μm.

The CDM is shown in Fig. 10. The CDM is a replica of the chip module except that a transparent photomask replaces the OE chip. The photomask possesses thin chrome patterns mimicking detectors and modulators. An opening is present at the back of the module for imaging purposes. The CDM makes it easy to visualize where the optical signal arrives by imaging the photomask with the help of a CCD camera. If a misalignment is present, adjustments can be performed by use of the Risley prisms or the corner prism before the actual OE chip is inserted in the system. This module can also be used to evaluate other system performance metrics such as beam profile and power distribution.

C. System Assembly

System alignment was found to be more complicated than originally planned. One stage (two BCMs) was assembled in a purely passive manner, and alignment was monitored with the diagnostic modules. The focused spots at the detector plane of the second BCM were misaligned by more than 100 μm and were severely clipped, causing fringing to appear in the spots.

Passive strategies failed to provide satisfactory alignment. The first goal of this demonstrator (passive alignment) was therefore not achieved. Note that the Risley prisms could not be used to correct such a large misalignment. However, the alignment proved remarkably stable. No drift in the spot position could be observed in the course of several days.

To permit system assembly, active alignment strategies had to be employed. This meant varying the position of the modules through the use of pads to bring them into alignment.

It was suspected that the BCM reflecting plane was the origin of the problem. This module had been characterized with the help of a micrometer and its physical dimensions found to be within the specified tolerance range. A high-magnification microscope combined with a precision *XY* translation stage was used to measure various coordinates on the BCM. These coordinates were then used to calculate the error on the angular value of the PBS reflecting plane (this should be 45°). The effect of angular errors will be to induce lateral misalignments at the modulator–detector plane.

Table 1. Characterization Results for BCM Reflecting Plane

BCM Number	Angular Error (degrees)
1	0.055
2	0.094
3	0.114
4	0.940

Table 1 lists the measured reflecting plane angle values for the four BCMs used in the backplane demonstrator system. It can be seen that there is a wide variation in angular deviation of the PBS reflecting plane with respect to the nominal 45 deg (deviations range from 0.055° to 0.94°).

It was decided to compensate for angular errors by placing pads of the appropriate thickness at the contact points between the BCM and the baseplate to modify the tilt and height of the module.

Because passive alignment was not achieved, aligning one stage (two BCMs) now involved considerable work. The optical modules had to be actively aligned with the use of pads to modify their positioning within the baseplate. Analytical relationships were derived to calculate the pad thickness required to achieve a desired displacement of the spots on the modulator–detector plane. System assembly thus became a labor-intensive procedure comprising many delicate and time-consuming alignment steps.

All four stages of the ring were assembled. Three to four days were necessary to align each stage once the alignment procedure was firmly established. Alignment was then deemed satisfactory for all the stages except the final one. The spots could not be brought on to the detectors at the last stage. The ring could not be closed. This is thought to be due to the accumulation of errors across the four stages of the ring, which led to a progressive elevation of the optical axis with respect to the mechanical axis of the baseplate. Figures 11 and 12 show pictures of the assembled system. A CDM has been removed to show the optical components underneath.

Each node in the system is numbered from 1 to 4 in a clockwise fashion starting from the upper-right corner. Note the compactness of the system. Figure 12 is a close-up of the optical layer. The BCM and the chip and relay modules can be clearly seen. Note the brass clamps that hold the modules in place.

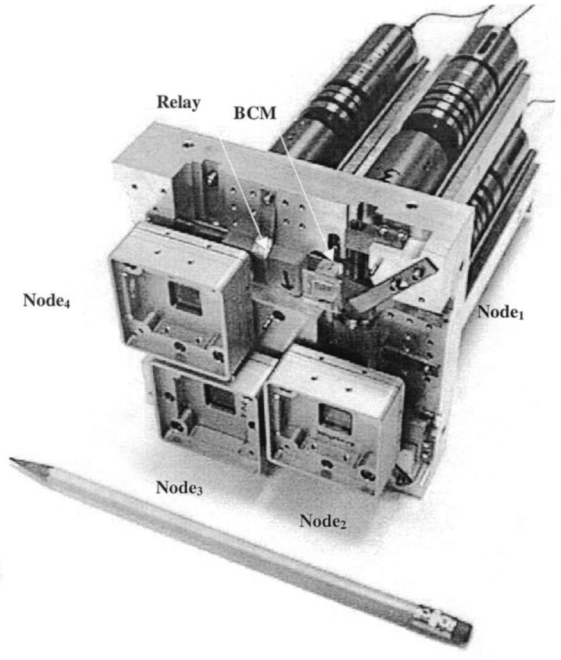

Fig. 11. Assembled system.

D. Characterization Results

Optical characterization results are useful to quantify the performance of the interconnect and the degree of alignment. The spots falling on the appropriate detector do not guarantee that alignment is perfect. Some degree of misalignment might still be present that is not detectable visually. Measurements such as throughput and uniformity are useful to verify this.

1. Optical System Power Throughput

The power throughput efficiency between the modulators at $node_2$ and the detectors at $node_3$ was measured. A 95% reflective mirror was placed at the modulator plane, and only a minilens array was placed in front of the detector plane. This was done to eliminate possible cavity effects in the uncoated photomask fused-silica substrate. The throughput efficiency from modulator to detector was found to be 23%, ±2% (−6.4 dB). This is quite good considering the number of interfaces (38) and minilenses (6) that the beam has to traverse when going from the modulator plane to the detector plane. The diffractive minilenses possess eight phase levels, giving them a theoretical diffraction efficiency of 95%. The real diffraction efficiency is lower when Fresnel losses are considered (90%). The beam must go through six minilenses before reaching the detectors (0.9^4 = 65%). The optical components are composed of various types of glass having different indices of refraction. For example, the PBS is made of SF56A having an index of 1.76 at 852 nm, whereas the diffractive components are made of fused silica having an index of 1.48 at 852 nm. The UV curable glue used to assemble the modules (Norland 61, index of 1.548 at 852 nm) does not perfectly index match to all these surfaces. The loss suffered by the beam when it goes from the PBS to the QWP is calculated to be 1.5%.

Considering all these sources of loss, the predicted theoretical throughput should lie between 21% and 36%. Note that this theoretical value does not include power losses that are due to clipping or misalignment. The experimentally measured value is within that range. This suggests that clipping does not contribute significantly to power losses. The system is well aligned.

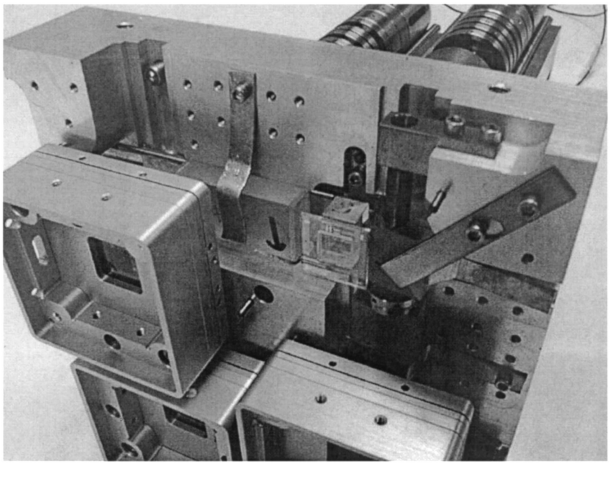

Fig. 12. Close-up of the assembled system.

2. Spot Quality

Spot quality is the degree to which the measured spot profile corresponds to the theoretical profile. It constitutes another way to verify both the imaging quality of the optical system and the degree of clipping that the beams suffer upon propagation through the system. Deviations from a Gaussian spot profile indicate that aberrations or clipping is present. In particular, the presence of fringes in the spots is a sign that optical power is being guided by multiple minilenses and recombining at the image plane. This is clearly undesirable.

Figure 13 shows a CCD picture of the array of 1024 spots at $node_3$. The columns coming in from $node_2$ are interlaced with the columns produced by the OPS module attached to $node_3$. The first column contains the spots generated by the second OPS module.

Note that the columns of clusters coming from the previous stage are indistinguishable from the ones produced by the second OPS in this low-resolution CCD image.

A 10× microscope objective mounted on a CCD camera was used to image clusters at the modulator and detector planes. The automatic gain control was disabled, and the camera was adjusted to respond linearly to variations of input power. Figure 14 shows an image of a 4×4 cluster at the modulator plane.

The power distribution in the cluster present at the modulator appears to be quite uniform. Note that

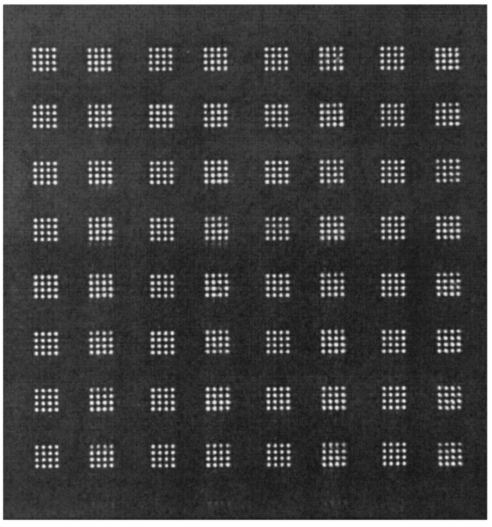

Fig. 13. 1024 spots at node$_3$.

the spots appear to be of good quality. No fringing can be discerned. Figure 15 shows a picture of the same cluster at the detector plane.

Note that the central spots appear to be slightly more intense than the outermost spots of the 4 × 4 array.

Figure 16 below shows the result of a multiple Gaussian fit performed on the first row of the 4 × 4 cluster at the modulator plane. The maximum peak intensity varies by approximately 15% across the row, and the spots can be seen to closely follow a Gaussian. Note the excellent fit to a Gaussian profile and the absence of fringes. Spot quality at the modulator plane is excellent.

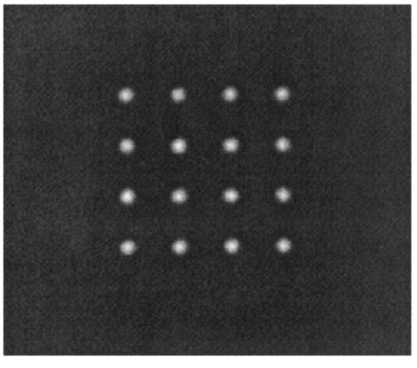

Fig. 14. Modulator cluster.

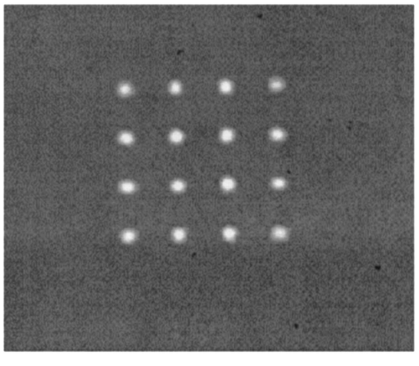

Fig. 15. Detector cluster.

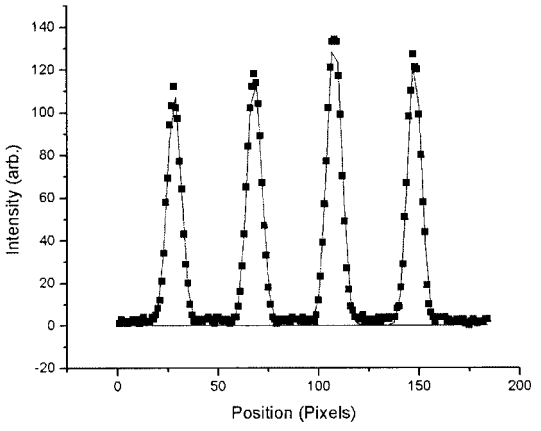

Fig. 16. Multiple Gaussian fit of spots at the modulator plane.

Figure 17 below shows the result of a multiple Gaussian fit performed on the first column of the 4 × 4 cluster at the detector plane. The maximum peak intensity varies by approximately 20% across the column. The spots can be seen to closely follow a Gaussian. Spot quality at the detector plane is also excellent.

3. *Array Uniformity*

A good method to perform uniformity measurements is the scanning pinhole method, which consists in scanning a pinhole across a grid and recording the power passing through the pinhole at each grid point. A plot of power versus position can then be obtained. This method is quite accurate provided that the pinhole is small enough compared with the size of the beams under measurement and that the laser source has a stable power output. In this case, the Spectra Diode Labs external cavity tunable laser source that was used has a power stability of 0.1%. A 5-μm pinhole was used to scan the 40-μm spots. The

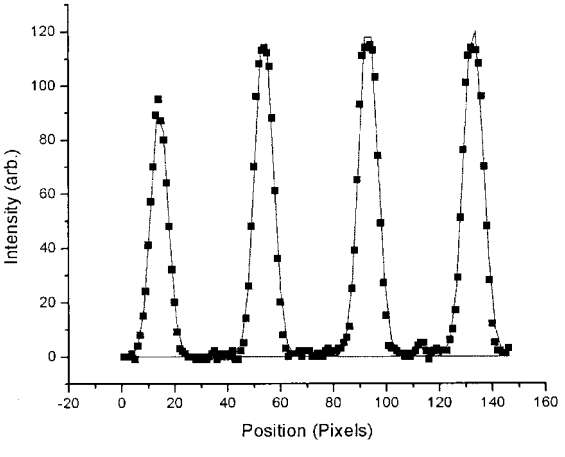

Fig. 17. Multiple Gaussian fit of spots at the detector plane.

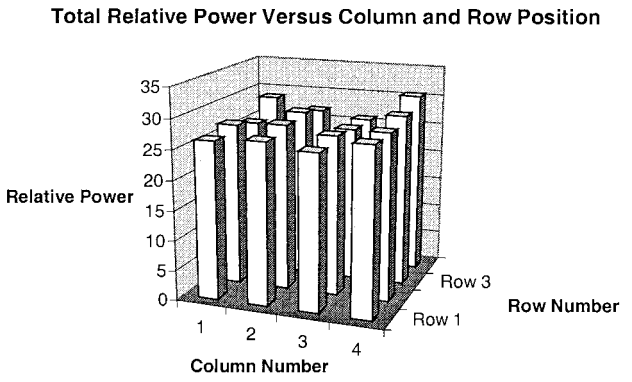

Fig. 18. Total relative power versus spot column and row position at the modulator plane.

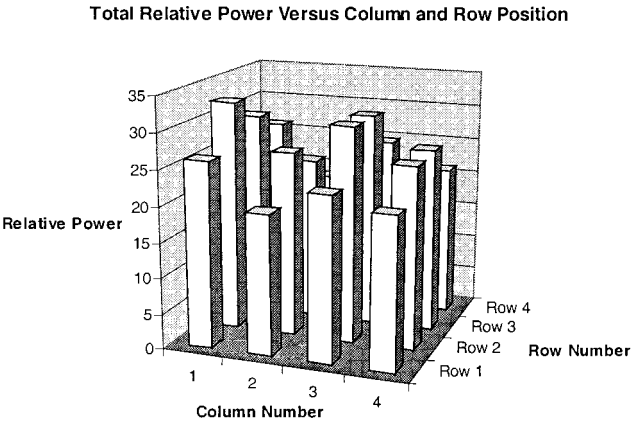

Fig. 19. Total relative power versus spot column and row position at the modulator plane.

scanning was done on a 5 μm $\times$ 5 μm square grid to eliminate convolution effects between successive readings. The large grid size used meant that each scan lasted several hours.

Note that uniformity constitutes a particularly critical measurement in this system as a dual-rail encoding scheme is used. Each bit is encoded in the power difference between two optical beams. There must therefore be enough power difference between the two channels for the receivers to switch. Any nonuniformity will degrade the power difference and affect the ability of the receivers to switch.

Figure 18 is a chart that helps to visualize the relative power uniformity in a cluster at the modulator plane. The total power in each spot was calculated by summing the relative power in each pixel of each spot.

The average value in the cluster is 27.23. The standard deviation is one way to characterize the spread in values in the data, assuming that they have a normal distribution. In this case, the standard deviation is equal to 5.07%. However, because a dual-encoding scheme is used, perhaps a more accurate way to characterize the uniformity is to look at the maximum power difference between pairs of spots. Here the maximum power difference takes place between the two bottom left spots, which possess relative powers of 25.74 and 28.75, respectively. This is an approximately 10% power difference. The rest of the spots are nearly identical in power. Uniformity is therefore quite good.

Figure 19 charts the total relative power present in each spot in the same cluster after it has traveled through the interconnect up to the detector plane. The central spots appear to be stronger in intensity than the peripheral spots, which points to a decrease in uniformity compared with the spots at the modulator plane.

The average power value is equal to 25.25, and the standard deviation is equal to 31%. This is a 5.7 times increase over the standard deviation calculated at the modulator plane. The two pairs of center top spots possess a power difference of approximately 30%, whereas the bottom left spots differ by only 10%. The average power difference between pairs of spots is approximately 20%. This is a significant increase over what was seen at the modulator plane. The 4 $\times$ 4 cluster at the detector plane is significantly less uniform than at the modulator plane.

This decrease in uniformity as the beams travel through the optical system can be attributed to fabrication errors in the diffractive minilens arrays used for focusing and collimation properties. The minilenses possess eight phase levels. Fabricating such a diffractive minilens requires the use of three masks to define the patterns that will be etched in the fused silica. Any misalignment between these masks, combined with possible mask fabrication defects, leads to inaccurate fabrication of the fine features located at the edges of each minilens.[16] The result will be a decrease in the diffraction efficiency at the edges, which induces nonuniformities in the intensity of beams propagating through the minilens. This is precisely what is observed at the detector plane.

The effect of nonuniformity is to introduce a power difference floor that must be respected if the system is to function properly. Dual-rail receivers are known to be quite tolerant of variations of input optical power (receivers have been demonstrated having common-mode dynamic ranges of 16 dB at 622 Mb/s); however, the tolerance to variations of the relative power between the two beams encoding the data is much lower. There is little data available on this topic in the literature; however, one reference[17] quotes a 3-dB variation as being the tolerance limit at 622 Mb/s for transimpedance-type receivers. Although application of these data to our system directly is slightly dubious as we do not use the same receiver circuits, it can nonetheless be used as a guideline. This would then mean that a 50% differential power variation is the maximum that can be tolerated if the system is still to function properly. A 20% differential power decrease owing to nonuniformity should therefore be tolerable.

4. Module Insertion Repeatability

One of the objectives of the demonstrator system was to demonstrate field-replaceable components and modules. To ensure that parts are completely interchangeable, the insertion precision of the modules is important. The package module and the OPS insertion repeatability and precision was evaluated to validate the design.

The test on the OPS was performed on a prototype baseplate that simulated the interface between the OPS and the baseplate. The OPS was introduced and removed from this setup 30 times to verify the repeatability and precision of the OPS insertion. The position and angle of the OPS optical axis was measured with respect to the ideal position. It was found that the OPS alignment is repeatable to 0.018° angularly and 15 μm laterally.[11] However, as can be verified on the two graphs, the absolute alignment accuracy was not as good as expected. This was caused by a design error discovered on the prototype baseplate used to perform the measurements. The error was corrected in the final baseplate design.

We also evaluated the repeatability of the OE-VLSI chip module insertion. The measurements were performed in a way similar to that described for the OPS. The measurement method is described in detail in Ref. 14. The results obtained show that the worst case of repeatability for the chip module insertion is ±12 μm. From the repeatability measurement, it can be shown that the chip module can be removed and replaced by another module as long as the optically sensitive region of the chip was larger than the spot size by this amount. This is a significant result as it suggests that field-serviceable systems with more than 1000 optically active devices are practical.

E. Discussion of System Assembly Results

The difficulties encountered during the alignment of this system can largely be ascribed to improper fabrication tolerances on the PBS cube that constitutes the heart of the BCM. Note that the PBS manufacturer could not guarantee a tolerance as this parameter is not usually specified by customers and is not part of standard optical shop practice. This parameter was therefore specified on a best-effort basis, which meant that the manufacturer could not guarantee degree angular precision better than approximately 0.15° to 0.2°. Specifying a custom tolerance for the PBS was beyond the financial means of our university research laboratory; however, it proved of crucial importance during alignment. The crucial importance of the BCM is one of the main lessons to be retained from this project.

The reflecting planes of the PBSs were measured to possess deviations ranging from 0.055° to 0.94° with respect to 45°. These constitute small but extremely important fabrication errors. It is important to remember that the angular error in the beams will increase upon reflection from the PBS and that angular errors will translate to lateral errors at the modulator–detector plane. Figure 20 presents a diagram of the effect of an angular error on the PBS reflecting plane where N is the refractive index of the PBS and ε is the angular fabrication error.

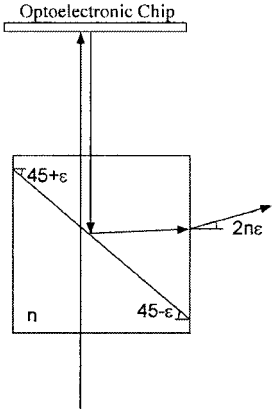

Fig. 20. Schematic diagram of the effect of the PBS angular error.

It was calculated that the deviation suffered by a beam upon propagating in a BCM possessing a fabrication error will be equal to twice the refractive index of the PBS times the angular error present in the PBS.[18] Knowing that the index of refraction of the PBS is 1.76 at 852 nm, it can be expressed as

$$\Delta = 3.52\varepsilon, \tag{1}$$

where ε is the angular error in the PBS. A small angular error will be magnified and will lead to a large deviation in the output beam. Note that no deviation will be induced as the beams propagate from an OPS to the modulator plane as no reflection on the PBS takes place. When a beam propagates from the modulator to the detector plane, it will be reflected once by the PBS in each BCM. The resulting total angular deviation at the output of the second BCM can be expressed as

$$\Delta = 3.52(\varepsilon_1 + \varepsilon_2), \tag{2}$$

where ε_1 is the angular error in the first PBS and ε_2 is the angular error in the second PBS. Calculations show that, depending on the combinations of BCMs used, the total beam deviation can vary from 0.5° to 3.7°. This translates to 77- and 290-μm lateral misalignments at the detector plane. As misalignments over 90 μm cannot be corrected by use of the Risley prisms, it was necessary to tilt the BCM with pads of unequal thicknesses to direct the spots onto the detector. Although this solves the tilt problem for the BCM under alignment, it will worsen it for the next BCM as the induced tilt will add to the fabrication error tilt. Progressively thicker pads and greater induced tilts will be necessary to align each successive node. For example, 500-μm pads were necessary to align node_1 to node_4. This meant that the optical axis of node_1 was 500 μm higher than the optical axis of Node_2, which prevented the ring from being closed.

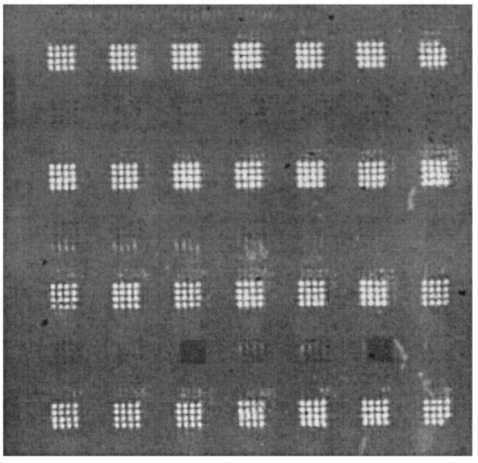

Fig. 21. Spots at node$_2$.

Below is a picture of the spot array at node$_2$. Figure 21 shows that an entire column of clusters is missing; of the 4×8 array of spots reflected from the modulators at node$_1$, only a 4×7 array of spots can be seen at the detector plane of node$_2$.

Figure 22, which is a close-up of the bottom right cluster, shows that the cluster focuses on the wrong targets (the left detector targets instead of the right one). The cluster could not be aligned in its proper target. This is thought to be due to a progressive elevation of the system optical axis owing to fabrication errors on the PBS as each stage is assembled.

5. Scalability Analysis

It is interesting to investigate what limits there are to the scalability of the system. Assuming that the optical design is not modified, the system is scalable in the sense that

(1) The number of nodes in the system (OE-VLSI chips) can be increased.
(2) The number of clusters can be increased, thus increasing the number of channels and the bandwidth between the nodes.

A. Increasing the Number of Nodes in the System

Increasing the number of nodes (switching the fabric size) in a system allows more data to be switched and routed. Reference 9 provides some considerations as to the scability of the hyperplane architecture. The fabric size (the number of nodes that can be implemented), the channel width (the number of bits in a channel), and the number of smart pixels present on the chip are all related. More precisely, the number of smart pixels is equal to the fabric size times the channel width. The present implementation contains 16×16 smart pixels and a channel width of 8 bits. The maximum fabric size attainable with the present implementation is thus equal to 32 nodes. The size of the OE chip would have to be increased to house more smart pixels to achieve larger fabric sizes.

The backplane is presently implemented as a four-node system arranged in a ring configuration. Figure 23 presents a flattened layout of the four-node system.

The number of nodes can theoretically be increased (to as many as 32) by turning a corner prism and repeating the pattern in a serpentine fashion as shown in Fig. 24, which illustrates an eight-node system.

Note that the eight nodes are arranged in a linear fashion. The last OE-VLSI chip in the chain is not optically linked to the first. This constitutes a problem for the present interconnect design because it is presently designed in a unidirectional ring configuration. The optical design of the interconnect would have to be modified to make it bidirectional in order to allow us to scale the system in a linear fashion. This would seriously complicate the optical design. Another potential solution would be to retain a ring configuration but to implement the system as a closed polygon. The number of nodes would then be bounded by the area that the optical layer could occupy.

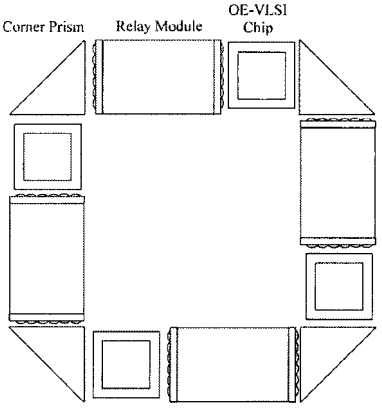

Fig. 23. Flattened four-node system in a ring configuration.

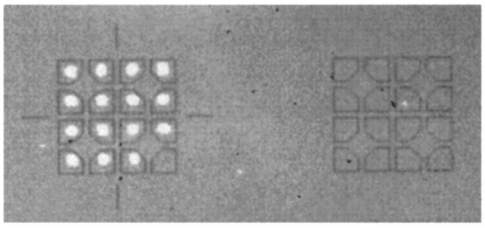

Fig. 22. Cluster at node$_2$.

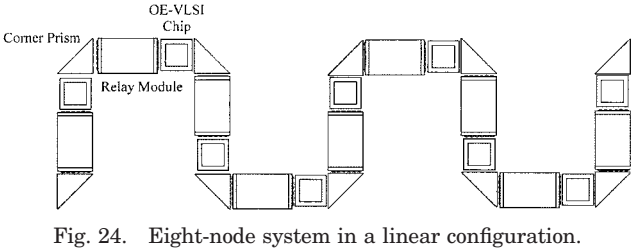

Fig. 24. Eight-node system in a linear configuration.

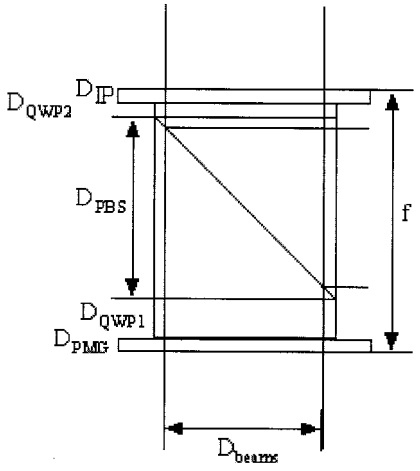

Fig. 25. Layout of a BCM.

B. Increasing the Number of Optical Channels

The number of optical channels linking each node can be increased by increasing the number of clusters. This increases the bandwidth between each node and the processing capacity of the system. However, assuming that the optical design of the interconnect remains the same, the number of clusters is constrained by a number of factors. Although there is no minimum, a maximum exists because if the focal length remains constant then the space to fit more optical beams is simply not available. Figure 25 illustrates the geometry of a BCM.

A total distance of one focal length (f) separates the PMG from the IP plane. This a design constraint as one focal length must separate the Fourier plane array generator (the PMG) from the minilens attached to the IP for the system to remain telecentric. The focal length ($f = 8.5$ mm) is assumed to be fixed. Note that increasing the focal length will not be beneficial as an increased focal length will also increase the minilens dimensions and thus decrease the interconnection density. It can be seen from Fig. 14 that increasing the diameter of the input matrix of optical beams (D_{beams}) also increases the size of the PBS (D_{PBS}) by the same proportion. However, the size of the PBS has to remain smaller than one focal length. Ideally, having $D_{\text{beams}} = D_{\text{PBS}}$ would ensure that the area of the optical components is fully used and no space is lost. In equation form,

$$f = \frac{D_{\text{PMG}}}{n_{\text{PMG}}} + \frac{D_{\text{QWP1}}}{n_{\text{QWP}}} + \frac{D_{\text{PBS}}}{n_{\text{PBS}}} + \frac{D_{\text{QWP2}}}{n_{\text{QWP}}} + \frac{D_{\text{JP}}}{n_{\text{JP}}}, \quad (3)$$

where f is the focal length, D_x represents the length of the components, and n_x is the index of refraction of each component. Note that in practice, $D_{\text{beams}} < D_{\text{PBS}}$ guarantees that some margin is left for possible misalignments.

Equation (3) can be used to calculate the maximum D_{beams} that can be used, assuming that the focal length is fixed and equal to 8.5 mm.

We assume that the components are at least 1-mm thick to ensure mechanical robustness and that the index of refraction of the glass used can vary from 1.48 (BK7) to 1.76 (SF56A). Solving gives $D_{\text{PBS}} = D_{\text{beams}} = 10.96$ mm.

Calculations indicate that a 13×13 array of 800 μm $\times$ 800 μm clusters occupies a 10.4 mm $\times$ 10.4 mm area. A 13×13 array of 4×4 clusters will provide 2704 optical beams to the OE chips, more than doubling the number of channels compared with the present implementation. This represents the absolute maximum achievable by use of this optical interconnect design.

Note that scaling to a 13×13 array might be challenging as the minimum feature size (the width of the minimum feature that can be etched) of the diffractive fan-out grating will decrease from 2 to 0.5 μm. Half a micrometer is approximately the minimum feature size that can be commercially fabricated in binary diffractive optics technology.

6. Discussion

The implementation of a four-stage, scalable optical interconnect for photonic backplane applications was described. The complete backplane consists of four 12f interconnects linked in a ring configuration. This is one of the most complex optical backplane demonstrators ever assembled.

The main goals of this project were to demonstrate

(1) Passive alignment of the optical modules within the supporting optomechanical structure.
(2) Repeatable insertion and extraction of the modules.
(3) Stability.

The first goal was not achieved. Fabrication errors in the angle of the PBS reflecting plane induced large misalignments in the beams propagating in the interconnect. Use of a ring configuration meant that the tolerance stackup continued throughout the whole 48f system. The presence of chamfers on the edges of the PBS means that part of the beams were blocked or deviated upon propagation. The modules had to be actively aligned and compensation pads used. This greatly complicated system assembly. The use of pads also meant that the optical axis of the system deviated from the mechanical axis of the baseplate. This gap between the baseplate mechanical axis and the interconnect optical axis was perceptible when closing the ring was attempted. There was a difference of approximately 800 μm between the actual position of the beams and where they should have been. This made it impossible to close the ring.

However, the second and third goals of the demonstrator were achieved. Module insertion and extraction were found to be highly repeatable. Once achieved, the system alignment was stable. No drift in alignment could be observed in the course of several days. This is a significant achievement as it means that once aligned, the system stays aligned.

It is believed that the use of highly accurate polarizing beam-splitter cubes possessing well-controlled reflecting angles could largely solve the alignment

problem at least for the assembly of short interconnects. Although the optical signals are terminated at each 12*f* relay, the use of a ring configuration where the first chip is optically linked to the last one means that we are really aligning a 48*f* optical system. The passive alignment of a 48*f* system requires extremely severe component-fabrication tolerances owing to tolerance stackup effects. Such a system might be too costly to be practical. It might be appropriate to investigate the use of active alignment techniques to decrease costs when long systems such as this one are assembled. For example, multiple sets of Risley prisms properly positioned along the relay (for example, Risley prisms positioned both before and after the relay module) might be used to bring the system into alignment. The use of an active alignment step might be less costly than the use of severe tolerances that guarantee passive alignment. The strength of this design is that once aligned, the system stays aligned.

Alternatively, using a linear configuration instead of a ring configuration to implement the backplane would mean that the last OE chip would not have to be aligned relative to the first one. This would limit tolerance stackup effects and greatly decrease the requirements on component tolerances and might render passive alignment achievable. However, a more complex bidirectional optical system would be required.

This research was supported by a grant from the Canadian Institute for Telecommunications Research under the National Centres of Excellence Program of Canada. F. Lacroix gratefully acknowledges funding from Fonds pour les chercheurs et l'aide à la recherche (FCAR). Thanks to Rhys Adams for help in assembling the system.

References

1. *The National Technology Roadmap for Semiconductors* (Semiconductor Industry Association, San Jose, Calif., 1997), p. B1. See www.sematech.org.
2. *International Technology Roadmap for Semiconductor 1998 Update* (Semiconductor Industry Association, San Jose, Calif., 1998), p. 4. See www.sematech.org.
3. D. A. B. Miller, "Physical reasons for optical interconnections," Int. J. Optoelectron. **11,** 155–168 (1997).
4. F. A. P. Tooley, "Challenges in optically interconnecting electronics," J. Sel. Top. Quantum Electron. **2,** p. 3–13 (1996).
5. F. A. P. Tooley, A. Z. Shang, and B. Robertson, "Alignment tolerant smart pixels," in *Advanced Applications of Lasers in Materials Processing/Broadband Optical Networks/Enabling Technologies and Applications/Smart Pixels/Optical MEM's and Their Applications: IEEE/LEOS 1996 Summer Topical Meetings* (Institute of Electrical and Electronics Engineers, New York, 1996), pp. 55–56.
6. K. Hirabayashi, T. Yamamoto, and S. Hino, "Optical backplane with free-space optical interconnections using tunable beams deflectors and a mirror for bookshelf-assembled terabit per second class asynchronous transfer mode switch," Opt. Eng. **37,** 1332–1342 (1998).
7. M. A. Neifeld and R. K. Kostuk, "Error correction for free-space optical interconnects: space–time resource optimization," Appl. Opt. **37,** 296–307 (1998).
8. F. Lacroix, M. Chateauneuf, X. Xue, and A. G. Kirk, "Experimental and numerical analyses of misalignment tolerances in free-space optical interconnects," Appl. Opt. **39,** 704–713 (2000).
9. T. H. Szymanski and H. S. Hinton, "A reconfigurable intelligent optical backplane for parallel computing and communications," Appl. Opt. **35,** 1253–1268 (1996).
10. B. Robertson, "Design of an optical interconnect for photonics backplane applications," Appl. Opt. **37,** 2974–2984 (1998).
11. 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. **39,** 733–745 (2000).
12. D. R. Rolston, D. V. Plant, T. H. Szymanski, H. S. Hinton, W. S. Hsiao, M. H. Ayliffe, D. Kabal, M. B. Venditti, P. Desai, A. V. Krishnamoorthy, K. W. Goossen, J. A. Walker, B. Tseng, S. P. Hui, J. E. Cunningham, and W. Y. Jan, "A hybrid-SEED smart pixel array for a four-stage intelligent optical backplane demonstrator," IEEE J. Sel. Top. Quantum Electron. **2,** 97–105 (1996).
13. F. Lacroix, "Analysis and implementation of a clustered, scalable and misalignment tolerant optical interconnect," Master's thesis (McGill University, Montréal, Québec, Canada, 1999).
14. M. H. Ayliffe, D. Kabal, F. Lacroix, E. Bernier, P. Khurana, A. G. Kirk, F. A. P. Tolley, and D. V. Plant, "Electrical, thermal and optomechanical packaging of large 2D optoelectronic device arrays for free-space optical interconnects," J. Opt. A. **1,** 267–271 (1996).
15. B. Robertson, Y. Liu, G. C. Boisset, M. R. Tagizadeh, and D. V. Plant, "In-situ interferometric alignment systems for the assembly of microchannel relay systems," Appl. Opt. **37,** 9253–9260 (1998).
16. J. Jahns, "Diffractive optical elements for optical computers," in *Optical Computing Hardware*, J. Jahns and S. H. Lee, eds. (Academic, Boston, 1994), Chap. 6, pp. 137–167.
17. T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, L. A. D'Asaro, L. M. F. Chirovsky, S. P. Hui, B. Tseng, D. Kossives, D. Dahringer, and R. E. Leibenguth, "1-Gb/s two-beam transimpedance smart-pixel optical receivers made from hybrid GaAs MQW modulators bonded to 0.8μm silicon CMOS," IEEE Photon. Technol. Lett. **8,** 422–424 (1996).
18. W. J. Smith, *Modern Optical Engineering*, 2nd ed. (McGraw-Hill, New York, 1990).