

# Design and implementation of a modulator-based free-space optical backplane for multiprocessor applications

Andrew G. Kirk, David V. Plant, Ted H. Szymanski, Zvonko G. Vranesic, Frank A. P. Tooley, David R. Rolston, Michael H. Ayliffe, Frederic K. Lacroix, Brian Robertson, Eric Bernier, and Daniel F.-Brosseau

Design and implementation of a free-space optical backplane for multiprocessor applications is presented. The system is designed to interconnect four multiprocessor nodes that communicate by using multiplexed 32-bit packets. Each multiprocessor node is electrically connected to an optoelectronic VLSI chip which implements the hyperplane interconnection architecture. The chips each contain 256 optical transmitters (implemented as dual-rail multiple quantum-well modulators) and 256 optical receivers. A rigid free-space microoptical interconnection system that interconnects the transceiver chips in a 512-channel unidirectional ring is implemented. Full design, implementation, and operational details are provided. © 2003 Optical Society of America

OCIS codes: 050.1970, 200.4880, 200.4650, 200.2610, 200.0020.

## 1. Introduction

Modern multiprocessor computer architectures place stringent demands on the interprocessor connection network. In a typical system, processors and memory are distributed across several different nodes, and data must be constantly transferred between them as computing operations occur. Many researchers have suggested that optical interconnects

may provide a suitable means of implementing the interconnection network and recent reviews of this topic are given by Collet *et al.*<sup>1,2</sup> In these works it is concluded that in the case of cache-coherent multiprocessor computers, access latency to main memory is a critical limitation and that optical interconnects have the potential to reduce this. Another conclusion is that ring-based architectures are best suited to this application. Other researchers have argued that as electronic transistors approach fundamental physical limits of size and speed, the interconnection network will become one of the main determinants of performance, and that optics has the potential to provide interconnects with lower power dissipation and higher performance than is possible with electrical interconnects.<sup>3</sup> Several authors argue that optics has an inherently greater maximum bandwidth-length product than is possible with electrical interconnects, based both on fundamental physics<sup>4</sup> and the modeling of specific implementations.<sup>5,6</sup> In the most recent of these<sup>6</sup> numerical modeling suggests that optical interconnects can outperform on-chip electrical interconnects over distances in excess of 1 cm. These recent results, together with preexisting research, have motivated us to implement a scalable 256-channel free-space optical interconnection system for a four-board multiprocessor computer system. The specific architecture that was selected was a nonuniform memory access (NUMA) system with a

---

Z. G. Vranesic is with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada. When this research was done, all other authors were with the Department of Electrical and Computer Engineering, McGill University, 3480 University Street, Montreal, Quebec, Canada, H3A 2A7. T. H. Szymanski is currently with the Department of Electrical and Computer Engineering, McMaster University, Hamilton, Canada. F. A. P. Tooley is with Terahertz Photonics, Livingston, UK. D. R. Rolston is with Lighthouse Photonics Inc., Quebec, Canada. M. H. Ayliffe is with Ignis Optics Inc., San Jose, California. F. K. Lacroix is with Ciena Inc., San Jose, California. B. Robertson is with the Department of Engineering, Cambridge University, UK. E. Bernier is with Nortel Networks Inc., Ottawa, Canada. D. F. Brosseau is with the Canadian National Railway Inc., Montreal, Canada. E-mail address for A. G. Kirk is andrew.kirk@mcgill.ca.

Received 13 September 2002; revised manuscript received 21 January 2003.

0003-6935/03/142465-17\$15.00/0

© 2003 Optical Society of America

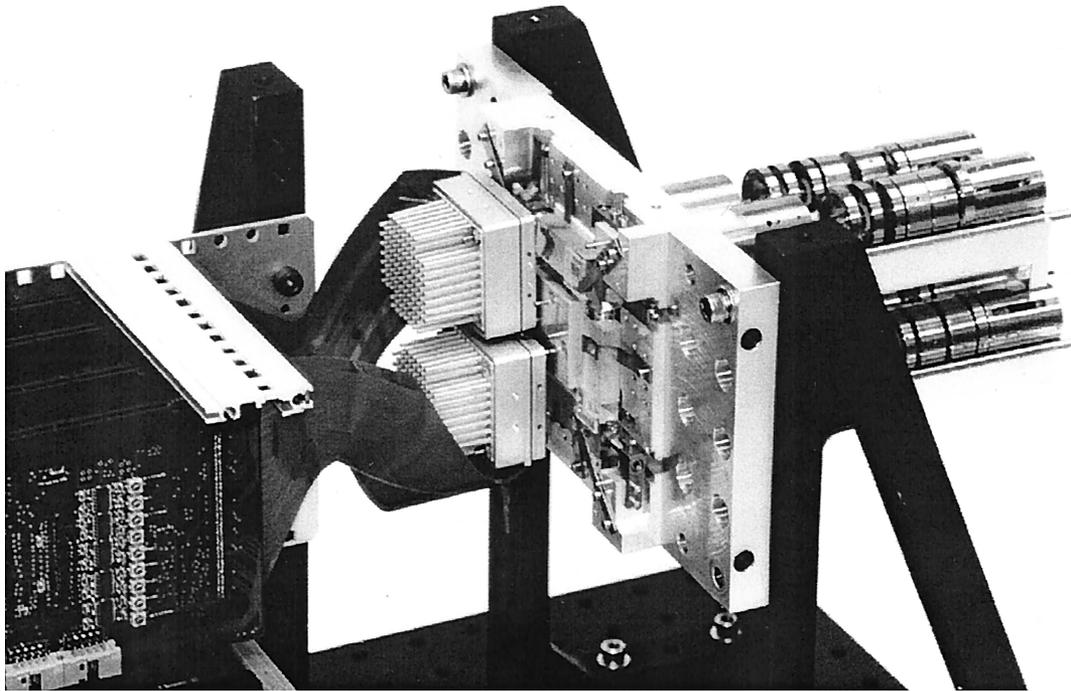


Fig. 1. Optical backplane system with two out of four chip modules inserted.

ring interconnect.<sup>7</sup> In this paper we describe the design and performance of the interconnect and then use the results obtained to draw conclusions for the design of future optical interconnection systems. In addition to creating a very highly parallel low-latency backplane interconnect that would scale to multiple nodes we also attempted to create a system that would be viable for commercial implementation. Thus one of the design goals for this system was ease of assembly, through careful optical design to obtain maximum possible tolerance to misalignment. A second desirable attribute was field serviceability. It was intended that the system was designed in such a way that active components could be removed and replaced without the need for further realignment. Although not all of these goals were fully realized, the development of this system has allowed us to answer a number of important questions about the implementation of highly parallel free-space optical interconnection systems. The techniques that were developed to meet these challenges are outlined in the remainder of this paper, which is structured as follows: In Section 2 we describe the NUMA system. In Section 3 we present the system architecture that was adopted and present the target performance requirements for the interconnect. In Section 4 we present the electrical portion of the interconnection system, and in Section 5 we describe the free-space optical system. The integration of the electrical and optical subsystems is presented in Section 6, and the system performance results are given in Section 7. A discussion of the strengths and weaknesses of this system, as revealed by experimental implementation, is given in Section 8, and conclusions are drawn in Section 9. A photograph of the entire system is pre-

sented in Fig. 1, showing the optical ring interconnect, the optoelectronic VLSI chip package, the message processor boards, and the optomechanical support structure, all of which is explained in the remainder of this paper.

## 2. Nonuniform Memory Access Multiprocessor System

In the NUMA architecture, processors have access to both local and nonlocal memories. Memory coherency is necessary to ensure that one processor is not operating on data that has been modified by another processor. Memory access latency is thus of vital importance and a significant portion of that latency is due to the backplane interconnections between processors and memories mounted on different boards. A high-bandwidth low-latency backplane is therefore required to accommodate simultaneous transfer of data between multiple processing nodes and nonlocal memories. The specific interconnect that we describe here was designed for use in the NUMAchine multiprocessor system developed at the University of Toronto.<sup>7</sup> A version of this system with purely electronic interconnects has already been implemented and thus provides a good test bed in which to evaluate the strengths and weaknesses of optical interconnects. In the all-electronic implementation of the NUMAchine a hierarchical ring interconnect is employed, as shown in Fig. 2. Local rings service groups of four stations, each of which contains four processor cards, two memory cards, and a network interface card. The local rings are interconnected by a central ring. All rings are unidirectional and are 64 bits wide with a bus frequency of 50 MHz (resulting in an aggregate data rate of 400 MB/s). In seeking to implement an optical interconnect for the

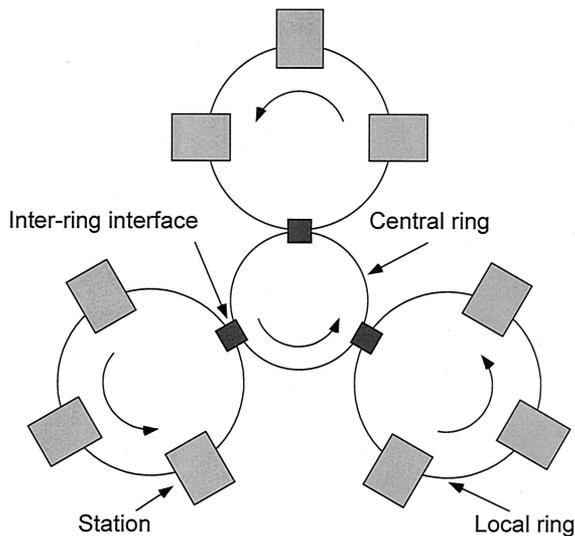


Fig. 2. NUMachine architecture showing the hierarchical ring structure.

NUMachine architecture it was decided to replace one of the local rings with an optical ring interconnect. The optical interconnect was therefore required to service at least four network interface cards, and that system was designed so that it would scale to eight nodes. This was the basis for the optical interconnect architecture. However, as explained below, the optical architecture was designed to give superior performance to the electrical backplane by permitting all nodes to access the backplane simultaneously and by offering lower latency.

### 3. System Architecture

Figure 3 is a high-level representation of a generic single multiprocessor system that shows  $N$  processors being interconnected via a photonic backplane. Transfer of data across the photonic backplane is performed as follows: Electrical data originating from a processor/memory is communicated to an optoelectronic VLSI (OE-VLSI) chip where it is con-

verted into modulated optical signals. The OE-VLSI chip thus acts as an interface between the processor/memory node and the high-capacity optical backplane. The optical signals propagate via optical channels from one OE-VLSI chip to the next. At each stage, the optical signals are either (i) regenerated on-chip and transmitted toward the next stage or (ii) converted back into electrical signals and directed off-chip toward the local processor/memory. In this way data passes around the network in a series of hops. The transfer of data between a processor chip and an OE-VLSI chip is performed over electrical lines. As a result, the amount of data that can be transferred on and off the backplane is limited by the off-chip electrical input/output (I/O) bandwidth (off-chip data rate  $\times$  number of I/O pads). However, the amount of data that can be transferred across the backplane is determined by the OE-VLSI chip optical I/O bandwidth (on-chip data rate  $\times$  number of transceivers in the array), which is typically one to three orders of magnitude higher. This apparent mismatch between the off-chip electrical I/O bandwidth and the backplane optical I/O bandwidth is successfully managed by properly balancing the system computational and communication bandwidths and by using the aggregate electrical I/O capacity of  $N$  OE-VLSI chips to fill the high optical I/O bandwidth of the backplane. This concept is more clearly shown in Fig. 4, in which the ring interconnect is explicitly drawn. Each node transmits on a single logical channel that is then relayed to all of the other nodes. Such a system is thus an example of a fire hose architecture,<sup>8</sup> in that in principle all of the processor/memory nodes can communicate simultaneously, so long as there is no contention. The system as shown in Fig. 4, and as implemented, is an example of a sender reserve system, but other choices are also possible.

The specific system that was implemented is based on the hyperplane architecture.<sup>9,10</sup> The design requirements of this system are summarized in Table 1. Each NUMachine node requires 64 bits of I/O. In the system implementation this was achieved by

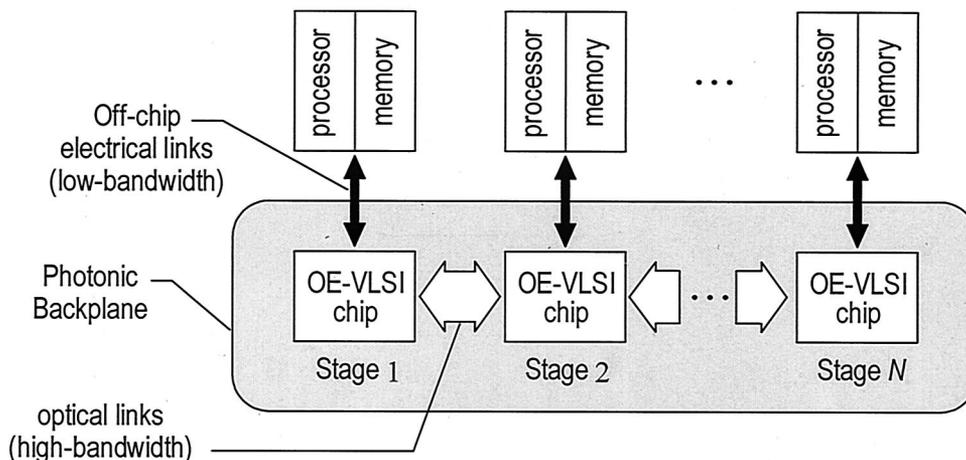


Fig. 3. Schematic representation of a multihop optical ring interconnection. Each node is either a processor or a memory.

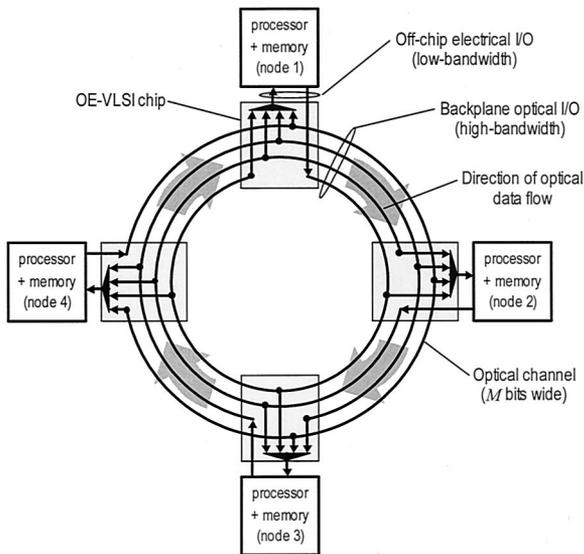


Fig. 4. Schematic representation of optical backplane ring interconnect.

Table 1. Target Values for System Parameters

Parameter	Target Value
Architecture	Ring
Number of nodes	8
I/O per node (electrical)	64 bits
I/O per node (optical data)	512 bits
I/O node (optical channels)	1024
OE-VLSI chip area	1 cm <sup>2</sup>
Data rate per I/O (electrical)	50 MHz
Data rate per I/O (optical)	200 MHz
Latency (two adjacent nodes)	20 ns
Latency (each additional node)	5 ns
Board spacing	25 mm

multiplexing two 32-bit packets. The optoelectronic chip therefore required a 32-bit wide input bus and a 32-bit wide output bus. To support up to 8 nodes the ring is therefore required to transmit 256 optical channels with 512 total I/O per chip. Dual-rail optical encoding was employed<sup>11</sup> so that each node has a total of 512 optical inputs and 512 optical outputs. Because the OE-VLSI chip area was limited to 1 cm<sup>2</sup> owing to technology and cost constraints, this requires a very high optical density, which motivated the use of free-space optical interconnects, as is explained below. In the system design, emphasis was placed on minimizing latency. For transmission from one node to an adjacent node the target for latency was 20 ns, with an additional 5 ns for each additional node to be traversed. The target clock rate for electrical I/O was 50 MHz, with an optical clock rate of 200 MHz. In the system that was actually constructed each OE-VLSI chip was connected to a message processor board, which was designed to communicate with the NUMachine station. This was done because the NUMachine system had already been developed with an electrical interconnect. To avoid any confusion of terms, the following definitions will be used for the various data paths in the system. The term logical channel refers to the 32-bit wide logical bus that is available to each NUMachine node. The term optical data channel refers to one of the 256 optical data paths that link each OE-VLSI chip. The term physical optical channel refers to one of the 512 dual-rail physical optical links between one OE-VLSI chip. An additional physical constraint was also imposed. To have a system that would be compatible with standard rack-mounted electrical equipment it was required that the board-to-board spacing should be 25 mm.

A more detailed view of the data path of the system that was implemented is shown in Fig. 5. In this

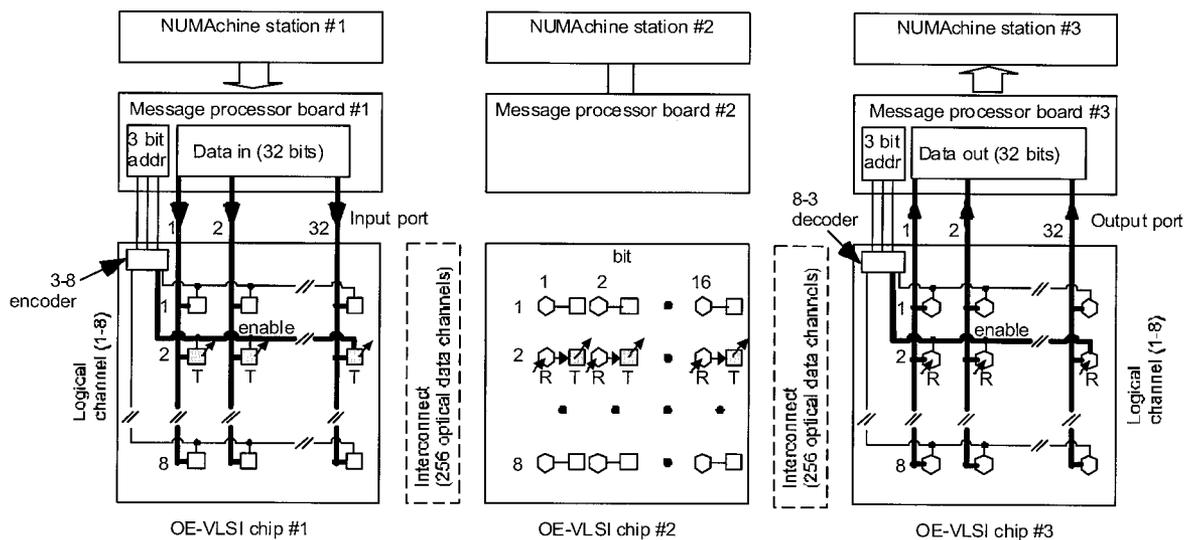


Fig. 5. Representation of the electrical and optical data paths showing the transmission of data from station #1 to station #3 on logical channel #2, with bypass through OE-VLSI chip #2. Note that all OE-VLSI chips contain both transmitters (T) and receivers (R), but not all are shown here.

example NUMAchine station #1 transmits a data packet to station #3 by use of logical channel #2. The 32-bit wide data stream is transferred electrically from station #1 to the corresponding message processor board and then onto OE-VLSI chip #1. The data are prefixed by a 3-bit address allowing the destination node to be selected. Once this reaches the OE-VLSI chip this address code is then used to enable logical channel #2 for transmission. The first 32 bits that are transmitted on this channel contain the destination address (it should be noted that this scheme could be extended to permit broadcast to all nodes, and that this address space is large enough to enable a 32 board system).<sup>9</sup> The address and subsequent data are transmitted across the optical interconnect by using 32 of the available 256 optical data channels. In the diagram OE-VLSI chip #1 is depicted as having only transmitters but in reality all chips have both transmitters and receivers arranged in transceiver units (sometimes also referred to as smart pixels). The signals are received at OE-VLSI chip #2 and the initial address bits are evaluated. It is determined that the data packet is not destined to be extracted at this node but should be rebroadcast (i.e., the transceivers here are in transparent mode for this channel). The address and data bits therefore traverse the short distance from the receivers on chip #2 to the transmitters and are retransmitted to OE-VLSI chip #3. Here the address recognition circuitry determines that the packet is destined for station 3 and the data is extracted. An address decoder is used to inform the message processor board as to the channel that the data arrived on (in a sender-reserve system this also provides information as to the identity of the transmitting node). While this transmission process is underway the other logical channels can be used simultaneously to route data between different pairs of nodes. However, a contention scheme is required when two transmitters both seek to send data to the same destination node.

As stated previously, the on-chip clock rate (200 MHz) is specified as being 4 times greater than the on-board data rate (50 MHz). This optical speedup therefore implies that once data has been injected onto an OE-VLSI chip it can traverse several nodes (ideally four) within one electrical clock cycle. In this way, although data traveling around the ring crosses several OE-VLSI chips, the latency is not significantly increased. This topic is discussed in more detail in the next sections that deal with the physical implementation of the system.

#### 4. Electrical System Implementation

##### A. Overview

The electrical portion of this interconnection system consists of the message processor boards that act as the source and destination of data, OE-VLSI chips that contain the active optoelectronic devices and thus act as an interface to the optical interconnect, and flexible electrical printed circuit boards that are used to connect them. Each OE-VLSI chip is con-

nected directly to the flexible electrical printed circuit board, which provides mechanical decoupling between the message processor boards and the chip. This approach was followed to maintain a high degree of mechanical alignment between the optoelectronic chips while also allowing the message processor boards to be inserted and removed from a rack with standard alignment tolerances.<sup>12</sup>

##### B. Optoelectronic Devices

Two different optical transmitter technologies have been widely used for parallel optical interconnects. These are vertical cavity surface emitting lasers (VCSELs) and multiple quantum well (MQW) modulators. In designing the system it was decided to use reflection mode p-i-n diode modulators.<sup>13</sup> Although arrays of up to 1080 VCSELs and p-i-n photodiode detectors attached to complementary metal-oxide semiconductors (CMOS) have been reported<sup>14</sup> these were not available at the time that the project was initiated, whereas large arrays of MQW devices have been previously demonstrated.<sup>15</sup> The modulators were patterned on a GaAs substrate and then flip-chip solder-bump bonded to the silicon with substrate removal. The same MQW devices were used for both modulators and the detectors. Each OE-VLSI chip has 256 optical data channel inputs and 256 optical data channel outputs, both of which use dual-rail encoding so that 512 physical optical inputs and outputs are required, resulting in a total of 1024 MQW diodes for each chip. As is explained in Section 5, the optical system design required the MQW diodes to be arranged in  $8 \times 8$  clusters of  $4 \times 4$  devices, as shown in Fig. 6. The clusters are pitched at  $800 \mu\text{m}$  in both directions, and the diodes within each cluster are pitched  $90 \mu\text{m}$  in both directions. The modulator diodes are  $50 \mu\text{m}$  in diameter and the detector diodes are  $80 \mu\text{m} \times 80 \mu\text{m}$ . The chip is also organized such that columns of modulator clusters alternate with columns of detector clusters. The device dimensions were selected as a trade-off between misalignment tolerance in the optical system (that motivates large devices) and low device capacitance (and hence high speed, which motivates small devices). A simplified diagram of the layer structure is shown in Fig. 7. The  $500\text{\AA}$  thick GaAs layer (cap layer) was *p*-doped ( $5 \times 10^{18} \text{ Be}$ ) to provide an ohmic *p*-contact for the later deposition of a metal contact that was also used as a reflecting mirror ( $400 \text{\AA}$  Ti/ $400\text{\AA}$  Pt/ $1000 \text{\AA}$  Au). The thickness of the MQW region (60 layers) was selected as a balance between responsivity, insertion loss, and voltage swing.<sup>16</sup> Etch stop layers were grown between the *n* region and the GaAs substrate to facilitate substrate removal after flip-chip bonding. The layer structure was grown via molecular beam epitaxy and devices were later etched in a wet etching process. The devices were designed to provide a maximum change in reflectivity at a temperature of  $40^\circ\text{C}$  and a wavelength of  $852 \text{ nm}$ . A photograph and schematic diagram of a modulator prior to flip-chip bonding is shown in Fig. 8.

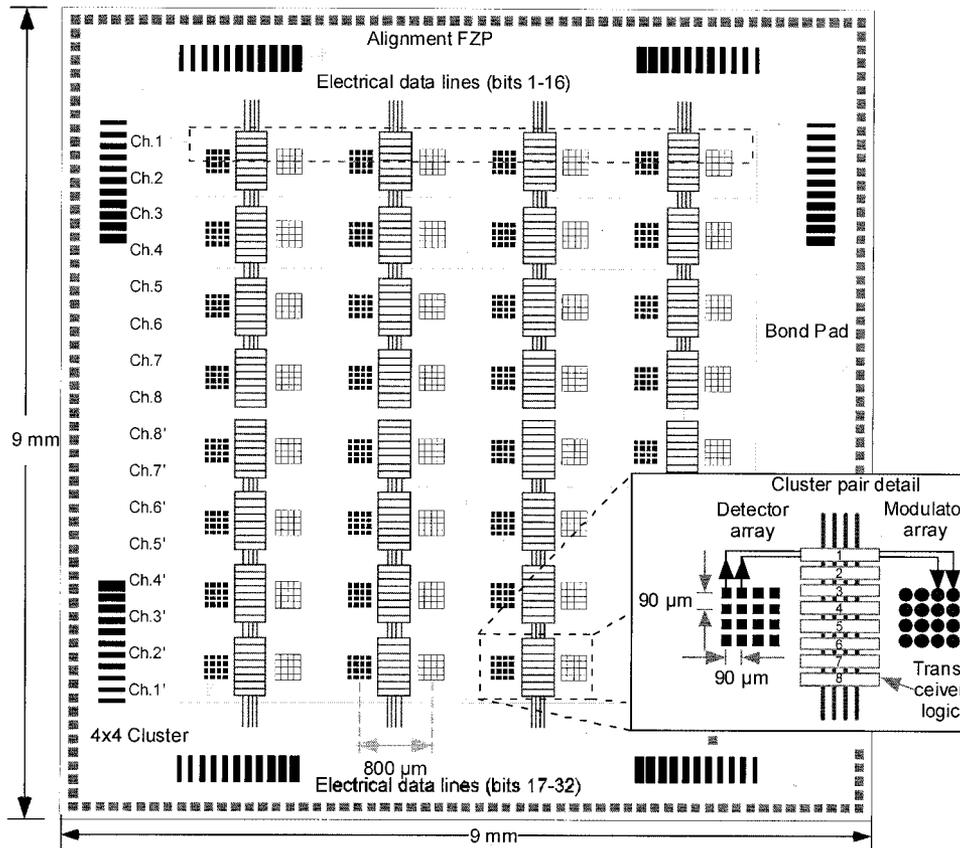


Fig. 6. The layout of the optoelectronic chip (FZP: Fresnel zone plate).

### C. Optoelectronic Very Large Scale Integration Chip

Initially the OE-VLSI chip was designed and fabricated as a  $9 \times 9$  mm device in a  $p$ -substrate ( $n$ -well) 5-Volt, 3-metal layer,  $0.8 \mu\text{m}$  gate-width BiCMOS process. However, owing to design errors these

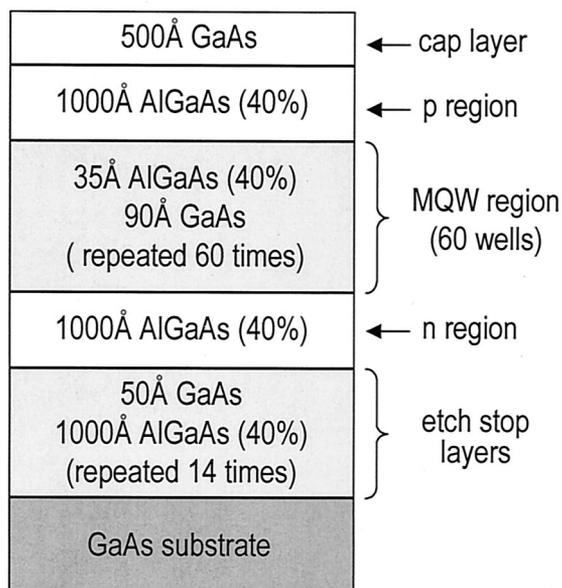


Fig. 7. MQW modulator layer structure.

chips were not fully functional, and so a second version was fabricated in an  $n$ -substrate ( $p$ -well), 5-Volt, 2-metal layer,  $1.5 \mu\text{m}$  gate-width, CMOS process. As a result of the increase in transistor-gate size, the maximum on-chip clock speed was therefore limited to 50 MHz. As described previously the chip contains 256 transceivers, each of which has 2 detectors (dual-rail optical input) and 2 MQW modulators (dual-rail optical output), together with electrical input, output, and control lines. A schematic diagram of one transceiver is shown in Fig. 9. The transceiver consists of a transimpedance amplifier used as an optical receiver, a bit-inversion multiplexer (to permit inversion of the dual-rail optical data during testing), a D-flip flop, a by-pass multiplexer, an output concentrator multiplexer, a transmit multiplexer, and a transmitter driver. There are 93 transistors per transceiver and almost 30,000 transistors on the chip (including bond-pad circuitry). To simplify the testing of the chip, each transceiver is provided with a mechanism allowing it to operate either synchronously or asynchronously. During asynchronous operation, an external clock is not required. However, during normal operation the optical and electrical data is latched in and out of the transceiver. The transceiver has three states: the add state, which converts electrical input data to optical data on one logical channel; the retransmit state, which allows optical data presented at the receivers to pass di-

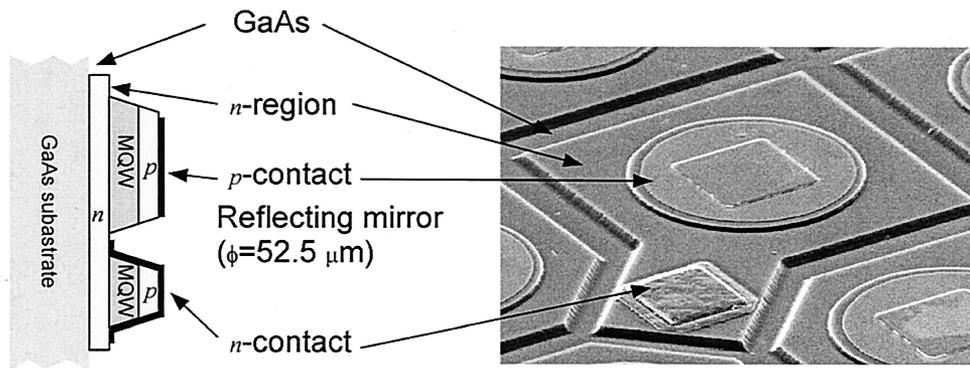


Fig. 8. MQW modulator prior to flip-chip attachment to the CMOS chip.

rectly to optical output data at the modulators, and the drop state, which allows optical input data from one logical channel to be converted into electrical output data which can then be transferred off the chip. As described in Section 3, this design allows data to be pipelined across a series of optically linked transceivers at a data rate determined by the on-chip clock speed. At each intermediate transceiver the data is transferred from the detectors to the modulators each time the D-flip flop is toggled. Once the data reaches a transceiver where the drop line has been enabled the data is transferred to the output concentrator (which is described below). When the chips are operated in an asynchronous mode, data can be transmitted around the ring with a delay that is determined only by the speed of the combinational logic circuits and the optical time of flight.

As shown in Fig. 6 and described in Section 4.B, the optoelectronic devices were arranged in clusters of  $4 \times 4$  devices, with alternate columns of modulator and detector clusters. The transceivers were therefore laid out and connected as shown in the inset to

Fig. 6. Each transceiver logic block is linked to a pair of detectors and a pair of modulators, so that there are 8 transceivers per modulator/detector cluster. The system design requires 32 electrical inputs and 32 electrical outputs per chip, with 8 logical channels. To accommodate this, and to ensure uniform loading on bond pads on all sides of the chip, the layout was partitioned into two symmetrical halves (top and bottom). The top half of the chip was used for bits 1–16 of each channel, while the bottom half was used for bits 17–32. Thus channel 1 was transmitted from the top two rows of the highest cluster and the bottom two rows of the lowest cluster (represented as 1' in Fig. 6). As a result the chip contained two electrical input concentrators and two electrical output concentrators, each 16 bits wide. These are fanned out to each channel. The address decoder is used to determine which channel the data is transmitted and received on.

There are a total of 232 bond pads on the chip. Of these, 24 are ground pads, and 20 power pads. There are 16 bond pads for MQW diode biasing and 2

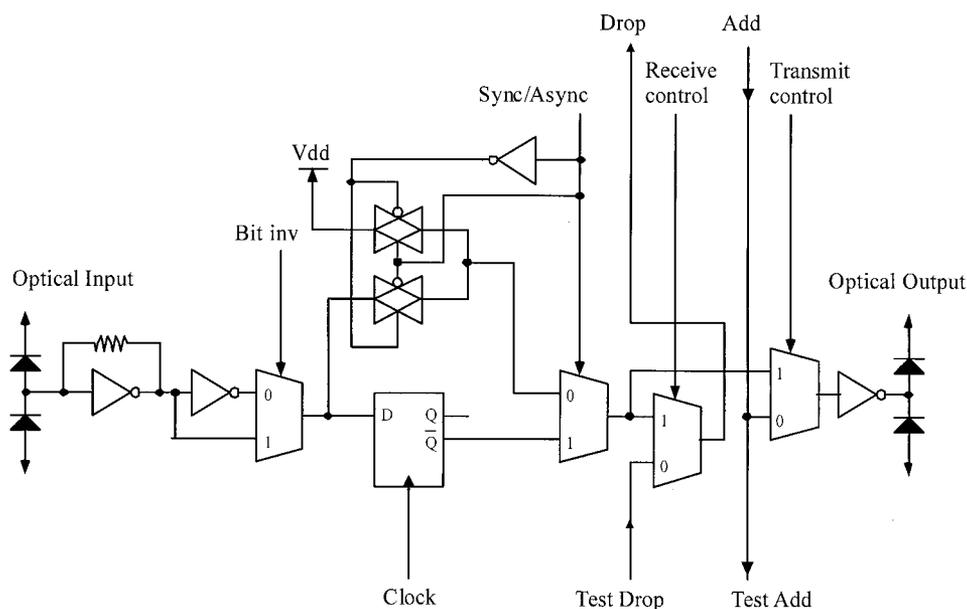


Fig. 9. Schematic diagram of a transceiver circuit.

bond pads for transimpedance amplifier optical receiver biasing. There are 64 I/O pads for electrical data, and 50 pads used for control and clocking of the channels. A total of 32 bond pads are used for array test structures and a total of 24 are used for active alignment techniques (described in Section 6). To assist with the alignment of the chip to the optical system, passive structures are also designed into the chip and are described in Section 6.

The power dissipation of the chip was estimated through the use of HSPICE simulations. The time-average current through a single transceiver during a receive-transmit phase (when both the detectors and the modulators are activated) was found to be 3.13 mA, with a total power dissipation of 15.6 mW. Extending this to all 256 transceivers results in an estimated power dissipation of 4.0 W. The power dissipation of the bond drivers was estimated to be 3.5 W, resulting in a total worst-case dissipation for the chip of 7.5 W. Further details relating to the design of the optoelectronic chip can be found in Ref. 17.

#### D. Electronic Packaging

To evaluate the optical interconnection system the message processor board was developed as a stand-alone board that allows for more extensive testing of the optoelectronic chip and of the system. This is a 10-layer PCB with 50  $\Omega$  impedance matched traces and was designed to both generate and process test vectors through the use of an on-board field programmable gate array (FPGA) chip. To test system performance, the message processor board routes signals on and off the FPGA via 189 I/O lines while receiving its controlling signals from either an on-board switch array or a PC-based LabVIEW 24-bit user interface. The FPGA can then output six distinct 16-bit test vectors (for which each bit may be activated or masked). The data stream can then be directed to two or more of the 32 available channels.

A flexible printed circuit board is used to connect the message processor board to the optoelectronic chip. It is 18.5 cm long, has 4 layers, and is designed with impedance-controlled I/O lines. Time domain reflectometry (TDR) measurements show a line impedance of  $53 \pm 3 \Omega$ . The flexible printed circuit board is connected to the message processor board via a 200 pin connector. Gold-plated fingers are provided for wirebonding to the optoelectronic VLSI chip. Decoupling capacitors (0.15 mF) are also mounted next to the optoelectronic chip to minimize switching noise.

### 5. Optical System

#### A. Optical System Design

A variety of different options are available for the implementation of parallel optical interconnects. These include guided wave solutions, such as parallel fiber arrays,<sup>18,19</sup> optical fiber image guides and image conduits,<sup>20,21</sup> and waveguides embedded into printed circuit boards,<sup>22,23</sup> together with free-space optical

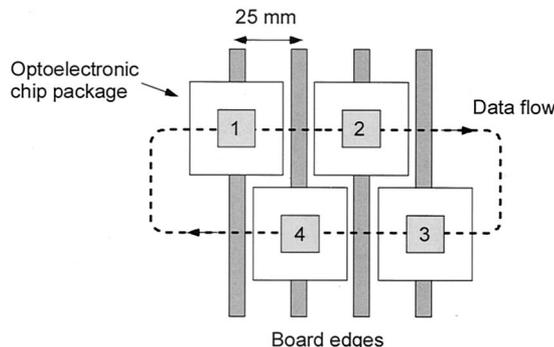


Fig. 10. Spatial relationship between the optoelectronic chips and the message processor board.

relays. Owing to the large number of physical optical channels (1024) and the requirement for a read-out spot array for the MQW modulators, it was concluded that guided wave solutions would not be appropriate. Within the field of free-space optical systems a range of options also exists, including microchannel relays,<sup>24</sup> macrolens imaging systems,<sup>25</sup> reflective systems,<sup>26</sup> planar optics designs,<sup>27</sup> and hybrid macro-micro clustered optical systems.<sup>28,29</sup> The physical dimensions of the optical assembly are determined by a number of factors. As stated previously, the board-to-board spacing was required to be 25 mm. Unfortunately, the size of the OE-VLSI chips and associated electrical packaging results in a minimum chip-to-chip spacing that is larger than this value. However, the ring architecture permits the optoelectronic chips to be placed on a 50-mm center-to-center spacing, while still respecting the 25-mm board-to-board spacing, because alternate boards can be connected to the upper and lower sections of the ring, respectively. This concept is illustrated in Fig. 10. The combination of a relatively large optical throw (50 mm) and the high spatial density of the array would make microlens relays impractical.<sup>30</sup> A planar optics solution could be attractive as a means of providing an optical relay, but the polarization control necessary to provide a beam combination for the MQW modulator readout would be difficult to implement in that technology. Reflective systems have performed well as interconnects for multichip modules, but in this case the relatively long interconnect distance militated against this approach. The remaining choices are macrolens systems (with a single lens or lens pair to image the entire chip) and a clustered optical system. However, the aberration-free field of macrolenses is limited unless multielement lenses are used.<sup>31</sup> Therefore to deliver a system that would be as compact and as tolerant to misalignment as possible, it was decided to implement the optical system as a clustered interconnect, based on a previously published design.<sup>32</sup>

The optical system is a ring, each arm of which is a telecentric  $8-f$  optical relay, where the focal length in air is 8.5 mm. A schematic diagram of the optical path that interconnects two nodes is shown in Figure 11. It should be noted that this system has a  $90^\circ$  fold

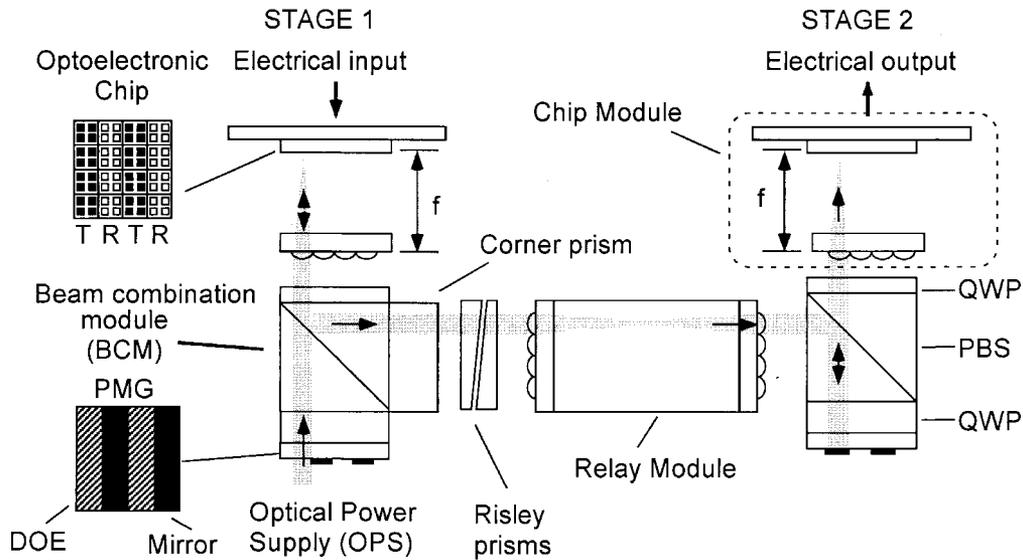


Fig. 11. Flattened layout of optical relay between 2 stages (PMG, patterned mirror-grating; T, transmitter (modulator); R, receiver; QWP, quarter-wave plate; PBS, polarizing beam splitter). Note that the relay module contains an  $8 \times 8$  minilens array, rather than  $4 \times 4$  as shown here.

(out of the plane of the diagram) at the corner prism. The optical array is divided into  $8 \times 8$  clusters, each of which contains a  $4 \times 4$  array of optical data channels. The design philosophy behind the optical system was to partition the system into modules. Elements within the modules have either a low tolerance to lateral misalignment (typically  $1\text{--}5 \mu\text{m}$ ) or low angular misalignment tolerances (typically frac-

tions of a degree) and are assembled by using precision alignment techniques. However, after assembly the modules have a high tolerance to misalignment with other elements in the system (typically  $10\text{'s of } \mu\text{m}$ ). The modules were designed to be inserted into a precision machined baseplate by using passive alignment, as shown in Fig. 12.<sup>33</sup> The baseplate was designed such that it could be computer

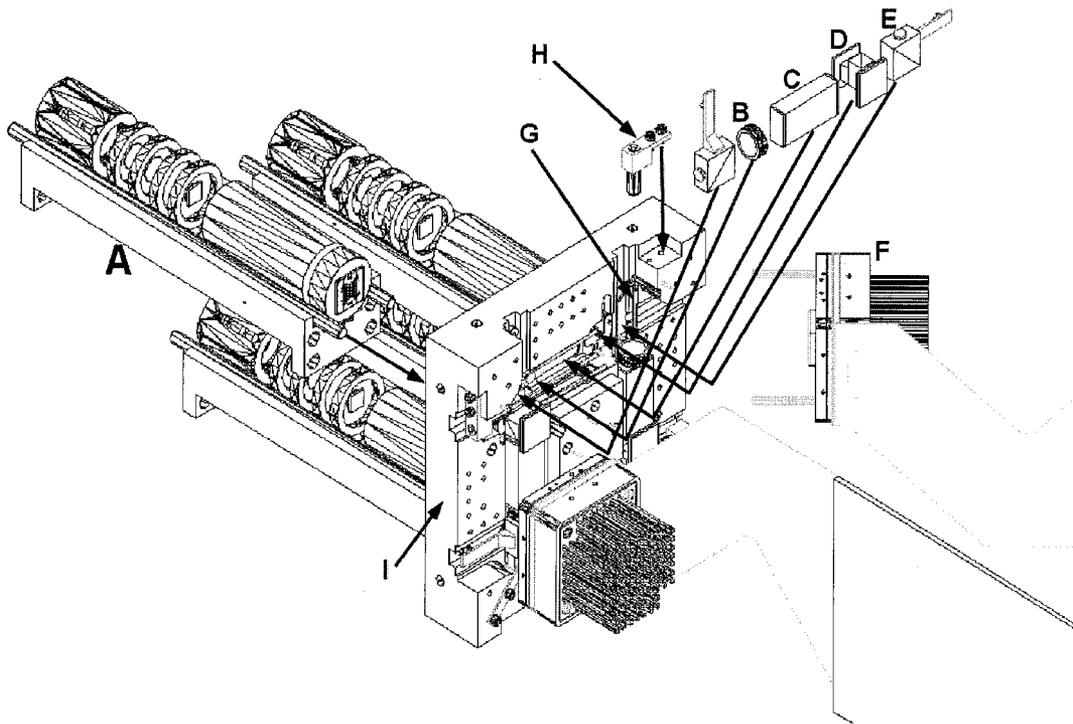


Fig. 12. Drawing showing exploded three-dimensional view of backplane. (A, optical power supply; B, Risley prisms; C, relay module; D, beam combination module; E, corner prism; F, OE-VLSI chip module; G, hardened steel rod; H, adjustment screw; I, baseplate).

numeric control machined in one pass from a single side in order to maximize precision for module insertion. A system of clamps and registration features was devised for each module to maintain this precision. Several adjustment mechanisms were also incorporated and are detailed in the next sections.

The system is partitioned into four modules: the optical power supply module, the beam combination module, the chip module, and the relay module. Complete details of the implementation of the optical system are given in Refs. 33–35, and so only a short description of each module is provided here.

#### B. Optical Power Supply Module

The optical power supply module (OPS) generates the continuous-wave (CW) beam that is required to read out the state of the MQW modulators. The OPS generates a 512-beam right-hand circular polarised CW spot array at a wavelength of 852 nm through the use of two stages of diffractive fan out.<sup>35</sup> The spots are focused onto the OE-VLSI via an  $8 \times 8$  array of mini-lenses. These lenses are 8-level diffractive elements, have a focal length of 8.5 mm, and a square aperture of 800  $\mu\text{m}$  (this results in a 1% clipping of the outer beams in the relays). The spots within each cluster are on a  $90 \times 90 \mu\text{m}$  pitch and the cluster-to-cluster spacing is 800  $\mu\text{m}$ . The OPS also contained a pair of Risley prisms to enable alignment of the beams onto the modulators on the OE-VLSI chip.

#### C. Beam Combination Module

The beam combination module (BCM) is a polarization-based unit designed to route three arrays of beams: the CW spot array of beams incoming from the OPS and directed to the modulators, the intensity-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 receivers. It is comprised of five components: two quarter-wave plates, a polarizing beam splitter (PBS), and a patterned-mirror grating. The patterned-mirror grating present on the BCM is composed of alternating stripes of diffractive fan-out gratings and gold mirrors. The diffractive portion forms part of the optical power supply system. It is used to split each CW beam in the  $4 \times 8$  array output from the OPS into a  $4 \times 4$  cluster. The reflective part is used to reflect the spot array incoming from the relay module onto the detectors.

#### D. Chip Module

The combination of the minilens array and the OE-VLSI chip (in addition to the electrical and thermal packaging for the chip) is described as the chip module and is described in more detail in Section 6. The minilenses perform the task of focusing the beams onto the OE-VLSI chip. The beam waist of the spots at the surface of the OE-VLSI chip is required to be 13.1  $\mu\text{m}$  to ensure that they are not clipped at the modulators or detectors.

#### E. Relay Module

The relay module is a 4-*f* telecentric relay that contains the same diffractive minilens arrays as in the chip module. Its purpose is to relay the optical signals between successive stages. It consists of a block of SF56A glass (with a refractive index of 1.76) and a minilens array at either end. The total length of the relay module is 29.56 mm. The choice of high-index glass was motivated by a desire to maximize the physical length of the module for a given optical-path length. A combination of visual-alignment features and interferometric alignment<sup>36</sup> was used to align the minilenses in the relay prior to bonding. Inspection of the relay modules after assembly showed the lens-to-lens misalignment was better than  $\pm 2.5 \mu\text{m}$ . The relay path also contains Risley prisms that were designed to permit the final alignment of the beams onto the detectors of the OE-VLSI chip in the next stage. Gross alignment of the system was designed to be achieved by adjusting the position of the corner prism, which was mounted on a micrometer screw.

### 6. Electrical-to-Optical Packaging

The interface between the electronic/optoelectronic portions of this system and the optical interconnect is of critical importance. This function is performed at the chip module. The chip module consists of the optoelectronic VLSI chip, and its associated electrical and thermal packaging together with an  $8 \times 8$  minilens array, as shown in Fig. 11.<sup>37</sup> The decision to package the minilens array with the chip module arose from a study of the optical-alignment tolerances. When the minilens is packaged as part of the chip module the lateral misalignment tolerance between the chip module and the beam combination module is  $\pm 26 \mu\text{m}$ . In contrast, if the minilens array were to be packaged with the beam combination module, the lateral misalignment tolerance for the chip module would be reduced to  $\pm 8 \mu\text{m}$ .<sup>37</sup> One of the design goals for this system was to ensure that the chip module could be removed and kinematically reinserted into the system. The larger lateral alignment tolerance was therefore an attractive feature of packaging the minilens array with the chip module. The penalty paid for the larger lateral alignment tolerance was a reduced tilt tolerance of  $\pm 0.03^\circ$  between the minilens and the beam combination module. This contrasts with a tilt-misalignment tolerance (between the OE-VLSI chip and the minilens array) of  $\pm 0.12^\circ$  for the individually packaged solution. However, the low tilt tolerance can be achieved by making use of the flat optical surfaces of the minilens array and the flat glass spacer that is attached to the beam combination module as reference surfaces. The required lateral-alignment tolerance was achieved through the use of precision machined dowel pins in the chip module and corresponding holes in the optical interconnect baseplate.

The alignment between the minilens array and the optoelectronic chip is critical. The minilens array

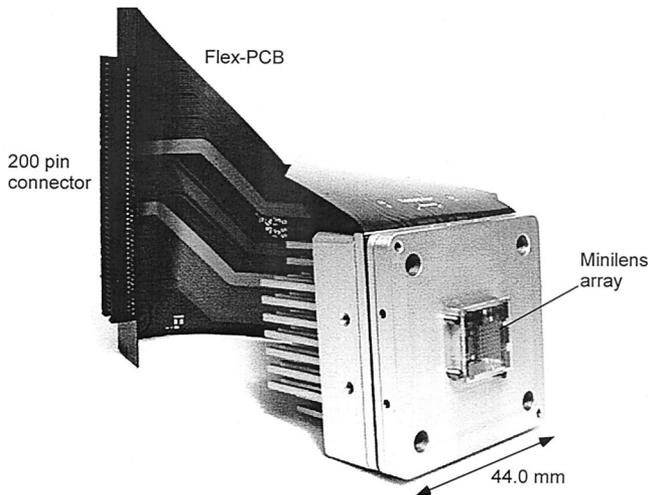


Fig. 13. Optoelectronic chip package module showing minilens array, TE cooler fins, and flexible printed circuit board.

must be aligned to a high degree of precision in all six degrees of freedom. However, it is positioned 8.5 mm from the chip, which makes the use of passive mechanical features problematic. Off-axis Fresnel lenses on the chip (as shown in Fig. 6) were used to align it to the minilens array. Further details are given in Ref. 37.

The chip module was also required to maintain the chip within  $\pm 5^\circ\text{C}$  of the specified operational temperature of  $40^\circ\text{C}$  to ensure that the modulator reflectivity remained within 90% of the peak value. To achieve this a thermoelectric cooler (TEC) was incorporated into the package in addition to a thermistor probe. The chip was mounted on a copper heat spreader that was attached to the TEC. The hot side of the TEC was then attached to an omnidirectional heat spreader. Experimental studies showed that this was able to withstand the worst-case thermal load generated by the chip.<sup>37</sup> Figure 13 shows the completed module.

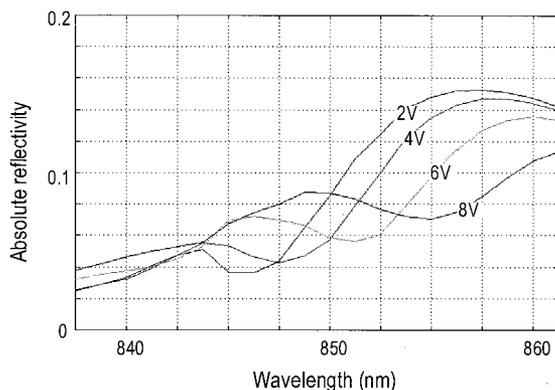


Fig. 14. MQW reflectivity curve as a function of the applied voltage.

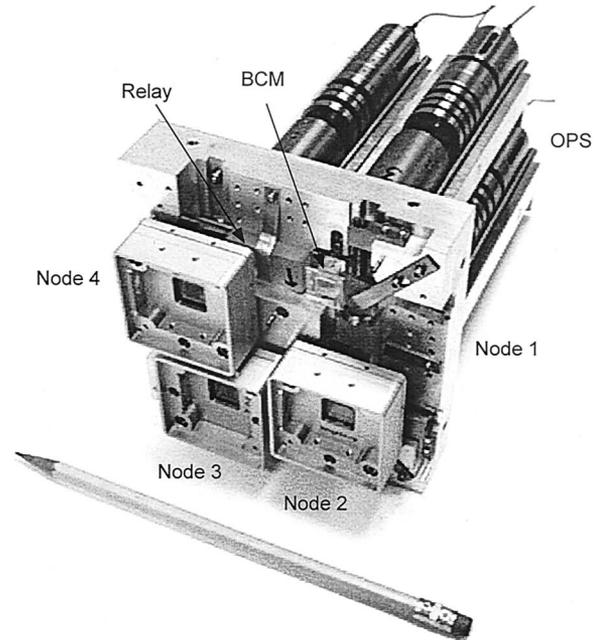


Fig. 15. Assembled optical system with a pencil for scale.

## 7. System Integration and Characterization

### A. MQW Modulators and Receivers

The MQW modulators and receivers were tested after flip-chip bonding. The reflectivity curve of the modulators as a function of the applied voltage is shown in Figure 14. Here it can be seen that the best high-voltage reflectivity  $R_{\text{high}}$  was approximately 15% while the low-voltage reflectivity  $R_{\text{low}}$  was approximately 7%. This represents a much lower reflectivity than the predicted values of  $R_{\text{high}} = 80\%$  and  $R_{\text{low}} = 30\%$ . Investigation of the modulators revealed that the reflectivity of the gold backside mirror was only 21% instead of the  $>90\%$  value that had been predicted. By measuring the performance of devices before and after flip-chip bonding, it was found that the degradation occurred during fabrication, most probably during the annealing of the ohmic contacts. The performance of the MQW devices as detectors was also measured. In this case they performed as expected, with a peak responsivity of 0.5 A/W at 852 nm.

### B. Optical System Assembly and Performance

The completed optical system after assembly is shown in Fig. 15 (note that the chip modules in the photograph contain chrome alignment targets rather than the OE-VLSI chips). The modular, kinematic design strategy that had been selected was very successful for the optical power supply module and the chip module. It was shown that the power supply modules could be inserted and removed from the optical system many times while maintaining lateral and angular tolerances. Over a cycle of 30 insertions and removals, the lateral alignment remained within  $\pm 5.8 \mu\text{m}$  and the angular alignment remained

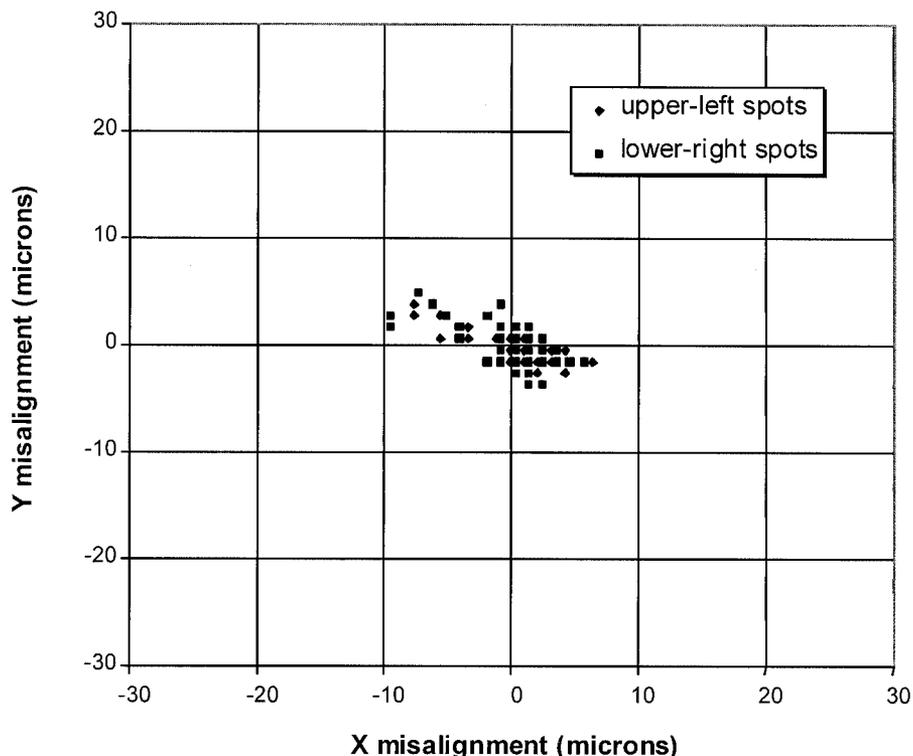


Fig. 16. Repeatability data for 30 insertion/removal cycles of the kinematic chip module.

with  $\pm 20$  arc minutes.<sup>35</sup> Similarly insertion and removal tests were carried out on the chip module. In this case the standard deviation of the spots on the detectors was found to be  $2.2 \mu\text{m}$  over 30 insertions and removals,<sup>38</sup> which is well within the required range. Figure 16 shows the results of this test.

However, the same passive module insertion strategy was less successful for the optical relays and was not able to deliver the required degree of precision. The focused spots at the detector plane in each module were typically misaligned by  $100 \mu\text{m}$ . This degree of misalignment was not correctable through use of the Risley prisms. Examination of the optical components showed that the beam combination modules did not meet the required specifications. The  $45^\circ$  reflecting plane was misaligned (relative to the reference plane of the BCM) by an average of  $0.3^\circ$  and  $87 \mu\text{m}$ . This was due to a combination of several factors, including a failure to properly specify the angular tolerance of the BCM PBS to the manufacturer and a larger than expected thickness of epoxy during the module assembly. To compensate for this it was necessary to modify the assembly scheme. The thickness error in each BCM was determined and pads were placed at the appropriate points on the base of the BCM to correctly route the beams through the system. Once this has been done the optical system was stable and performed well.

The optical system was tested by connecting the input to one of the optical power supply modules to an 850-nm diode laser. The optical power supply insertion efficiency (from input fiber to modulator plane) was found to be 33% (not including fan-out loss).<sup>35</sup>

Figure 17(a) shows the 512-spot array at the receiver plane after transmission between two stages. A detailed image of one cluster before transmission (at a modulator plane) is shown in Figure 17(b), and the spot cluster after transmission is shown in Fig. 17(c). The stage-to-stage throughput efficiency (from modulator to detector) was found to be 23%. This compares to a predicted throughput efficiency of 29%. The dominant source of loss in the optical system was the diffraction efficiency of the 8-level diffractive relay lenses (90% including reflection losses). Additional loss was thought to be due to clipping at the apertures of the system due to misalignment. The waist at the chip was slightly larger than the required beam waist ( $14.1 \mu\text{m}$  in comparison with  $13.1 \mu\text{m}$ ). This indicates a possible slight departure from telecentricity in the optical system owing to axial misalignments of some components, but would not significantly impact system performance. The array uniformity at the modulator plane and the detector plane was measured. At the modulator plane the standard deviation in beam powers was 5.1%. The nonuniformity at the detector plane was significantly worse than this (with a standard deviation of 31%). The power distribution for one output cluster is shown in Fig. 18. In general spots at the edge of the array had less power, indicating that clipping had occurred in the system and also that the diffractive minilenses were less efficient off axis as the feature size decreased.<sup>39</sup> Further details of the optical system assembly and performance are given in Refs. 33 and 34.

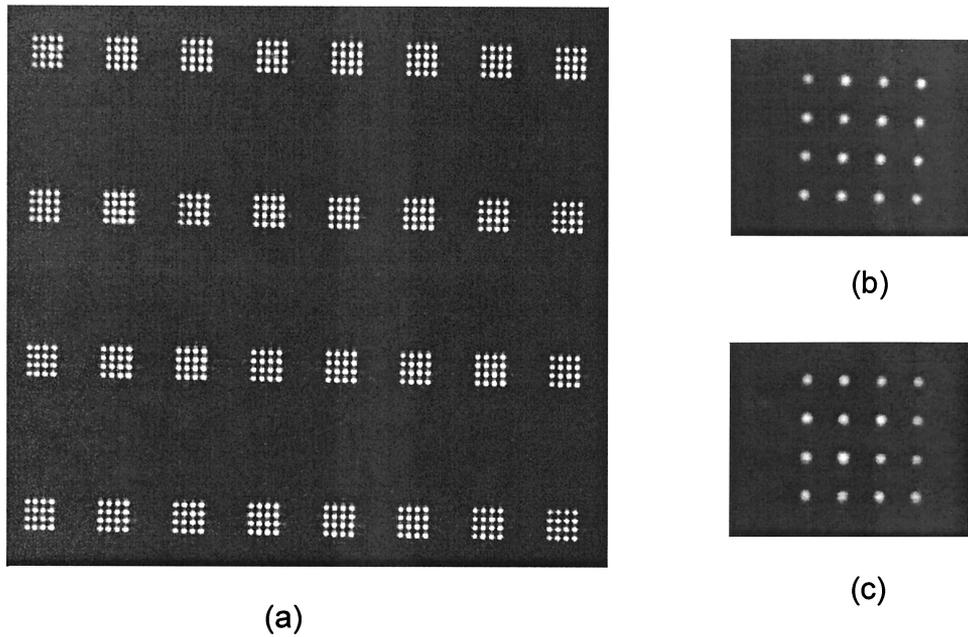


Fig. 17. Spot array at the receiver plane after transmission through the optical system (a), detail of spot cluster at the modulator plane before transmission (b), detail of spot cluster after transmission (c).

### C. Electronic System

Performance of the digital portion of the electronic system was tested in several different ways. Prior

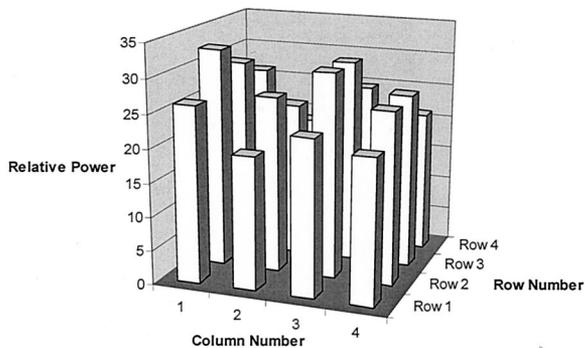


Fig. 18. Total relative power for a cluster of spots at the receiver array.

to flip-chip bonding of the optoelectronic devices, a series of probe tests were carried out on the chip while it was attached to a test board. The chip was also tested after packaging to the flexible PCB by running a series of test vectors from the FPGA on the message processor board. All aspects of the digital system performed as expected. The performance of the flexible PCB was also evaluated. It was capable of transmitting signals up to 1 Gb/s.

### D. Optoelectronic Chip

The performance of the OE-VLSI chip transceivers was investigated. The add state was tested by illuminating one modulator within a pair with a cw 850-nm readout beam. Digital electrical data from the message processor board was then transferred to the appropriate transceiver, and the modulated optical signal was detected with a high-speed avalanche photodiode detector. The detected signal is shown

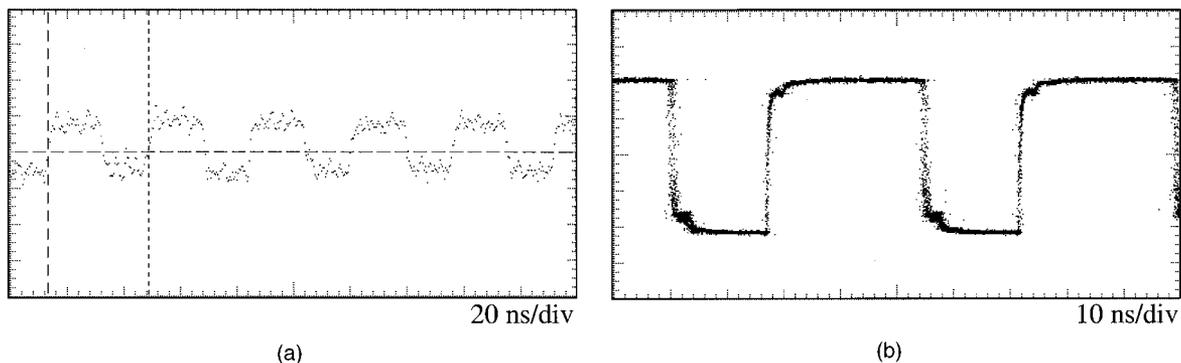


Fig. 19. Optoelectronic chip test results: (a) add mode (optical output from electrical input), (b) drop mode (electrical output from optical input).

in Fig. 19(a). The rise time of the modulator was 1.6 ns and the fall time was 1.8 ns. These data rates represent the limitation of the pad speed of the optoelectronic chip. The maximum data rate of the transmitters was 56 Mb/s, which is consistent with the predicted speed of the 1.5  $\mu\text{m}$  CMOS.

The extract state was also tested. A dual-rail binary modulated optical signal generated by a complementary pair of 850-nm wavelength VCSE lasers was directed onto two detectors that made up a receiver. The electrical signal generated on the chip was then relayed along the flexible PCB onto the message processor board via the on-board FPGA. In open-loop mode (i.e., with the feedback resistor open circuited) the receiver demonstrated an excellent sensitivity (approximately 2 mW and 4 mW of differential power). However, the rise time of the receiver amplifier in this mode was extremely slow (approximately 1.4 ms), limiting the maximum data rate to 1 MHz. In closed-loop mode the sensitivity was not as good (requiring 24 mW and 44 mW of differential power) but a data rate of 75 MHz was achieved. Figure 19(b) shows the results of these tests. The rise time of the optical-to-electrical signal was 1.1 ns and the fall time was 1.9 ns.

#### E. System Performance

By transferring data from the message processor board to and from the transceivers on the optoelectronic chip it was possible to verify that the digital portions of the system were performing as expected. However, the verification of the operation of the system as a whole was prevented by the low reflectivity of the modulators. To operate at reasonable speed, the receivers required 44 mW. The losses in the optical system consisted of the OPS transmission efficiency (33%), the optical system transmission efficiency (23%), the 1/512 optical fan-out loss and the 15% high-state reflectivity of the modulators. Therefore to deliver the required power of 44 mW to a modulator, an optical power of approximately 2W would be required at the input to the OPS. This calculation does not take into account the nonuniformity of the optical system, or the insertion loss that would be experienced when trying to couple a light source to the input fiber of an OPS module. This power level was not available at the time that the system was being tested and so the board to board transmission of data could not be demonstrated.

#### 8. Discussion

In developing this system, we have attempted to demonstrate that free-space optics has the potential to provide high-density interconnections at the board-to-board level. This system has been designed and constructed to obtain maximum benefit from the parallelism provided by free-space optics. In particular, although the interconnect was configured as a series of point-to-point links between boards, the use of the hyperplane architecture permits data to hop over multiple nodes with low latency. Despite the fact that the system did not operate in its entirety, a

number of important and valuable lessons were learned during the design, construction, and testing phases of this project.

The optical components were implemented as a set of modules that were designed to be integrated together without requiring critical alignment after assembly. In practice this goal was not achieved completely for the optical relay, because of an incomplete specification of the tolerances of some of the parameters of the optical components and also because of insufficient control of epoxy thickness. In both cases these problems could be eliminated in a future system by more careful tolerancing of the components. The optical system for the chip module was implemented successfully as a module that could be passively inserted and removed. In the future it would be desirable to implement some of these modules as single components molded in plastic. The relay module is a strong candidate for this process because it consists simply of two diffractive surfaces separated by a transparent block. An example of the integration of free-space optical systems via plastic molding is given in Refs. 40 and 41. However, most widely used optical polymers have a thermal expansion coefficient that is an order of magnitude larger than that of glass, and this may limit practical applications in systems that must operate over a large temperature range. Another example of the way in which the optical system could be improved would be to reduce the size of the optical power supply module. At present each optical power supply module is 170 mm long. An equivalent folded microoptical system that has a length of only 85 mm and that requires many fewer components has been designed and successfully tested.<sup>42</sup> However, the optical system would be further simplified if the modulators were replaced with emitting elements such as VCSELs. This would remove the need for the optical power supply altogether. An example of a VCSEL-based optical interconnect with a similar optical throw and that contains 256 optical channels is given in Refs. 43 and 44. However, the use of VCSELs would require a modification of the optical design to accommodate the high divergence of a VCSEL beam. A final potential modification would be to employ a flexible optical interconnection (such as fiber arrays or fiber image guides) in place of the rigid link used here. This would be greatly simplified if VCSELs were used. Fiber arrays of 64 channels have been reported<sup>19</sup> but at present the largest commercial fiber arrays have 48 channels. The use of a flexible optical interconnect would allow the flexible printed circuit board to be eliminated so that the optoelectronic chip could be mounted directly on the message processor board. For these short interconnect distances it remains an open question as to whether the cost of terminating an optical fiber array at each end can be competitive with molded plastic microoptics. It is also interesting to consider the scalability of the optical system (assuming the same cluster size and beam spacing). Calculations show that the maximum optical field is limited by the di-

mensions of the PBS.<sup>33</sup> The maximum field size is  $10.4 \times 10.4$  mm, which results in a  $13 \times 13$  array of  $4 \times 4$  clusters. This would allow an I/O of 1352, thus doubling the size of the system and permitting it to scale to 16 boards (each with 64 bits of optical I/O). The challenge in this case would be to deliver sufficient optical power to drive all the modulators, and also to arrange the electrical wiring of the OE-VLSI chip. In the current system, even if the modulator reflectivity had been 75%, which was the lower end of the predicted range, the required input power at the OPS would have been 400 mW. This represents an achievable value but one 1 W power supply laser would drive two stages at most. Doubling the size of the system would require almost 1 W per stage. It might also be of interest to investigate the possibility of using optical fan out to interconnect each transmitter to more than one receiver, in contrast to the simple point-to-point links implemented here. Examples of systems based on polarization-controlled fan out are given in Refs. 45 and 46. However, this approach places further stress on the alignment requirements of the optics and reduces the power available at the receivers.

Improvements could also be made to the OE-VLSI chip. One example of this would be the inclusion of flow control and error correction circuitry. References 47–49 contain a detailed description of two different optoelectronic chips for parallel interconnect applications that contain such circuitry. Other researchers have designed OE-VLSI chips that implement more complex protocols and that contain carrier-sense-multiple-access/collision-detection.<sup>50</sup> Improvements in the receiver sensitivity could also be made by increasing the complexity of the receiver circuitry.

## 9. Conclusions

We have presented the design and implementation of an optical interconnection architecture for multiprocessor computer systems. The system is designed to interconnect four boards (scalable to at least eight) in a unidirectional ring configuration. Each board has access to a 32-bit bus that is interconnected to all the other boards in a nonblocking fashion. Each board is connected to an OE-VLSI chip via a flexible electrical ribbon cable. The chip contains 8 discrete 32-bit logical channels. Within each channel data can be added or dropped from the optical ring data, or can be transparently retransmitted to the adjacent chip in the ring. Data is encoded optically by use of dual-rail logic so that each chip requires a total of 512 MQW modulators and 512 detectors to provide 256 bits of input and output. A rigid microoptical system is used to transmit the optical signals between the optoelectronic chips. This results in a very high density of 1,250 channels/cm<sup>2</sup>. The microoptical components are assembled into modules that are then integrated into the optical system. The modules were designed to allow for passive insertion. However, owing to some uncontrolled component dimensions it was necessary to use active alignment to

complete the assembly of the optical system. The optical transmission efficiency from modulator to detector was measured to be 23%. Owing to degradation of the gold mirrors the reflectivity of the MQW modulators was too low to permit complete system operation. However, the transceivers were shown to operate correctly in adding and dropping data to and from the optical ring.

This work was supported by a grant from the Canadian Institute for Telecommunications Research under the NCE program of Canada, the Nortel/NSERC Industrial Research Chair, NSERC, and FCAR research grants. We gratefully acknowledge the support of the Canadian Microelectronics Corporation in fabricating the CMOS VLSI chip used in this system. David Rolston, Eric Bernier, and Frederic Lacroix were supported by scholarships from NSERC and FCAR. We thank the following individuals for their help. Assistance with MQW wafer design, growth, fabrication, patterning, and flip-chip bonding was provided by John Trezza (formerly of Sanders/Lockheed Martin), Anthony Springthorpe (Nortel Networks), David Neilson (formerly with the NEC Research Institute, Princeton), Edwis Richard and John Curry (formerly of the Université de Montreal). PCB design and testing at McGill was provided by Feras Michael, Alan Chuah, and David Kabal. Assistance with some of the diffractive optical elements was provided by Mohammad Taghizadeh of Heriot-Watt University. We thank Dominic Goodwill (of Nortel Networks) for valuable support and advice throughout this project.

## References

1. J. H. Collet, D. Litaize, J. Van Campenhout, C. Jesshope, M. Desmulliez, H. Thienpont, J. Goodman, and A. Louri, "Architectural approach of the role of optics in monoprocessor and multiprocessor machines," *Appl. Opt.* **39**, 671–682 (2000).
2. J. H. Collet, W. Hlayhel, and D. Litaize, "Parallel optical interconnects may reduce the communication bottleneck in symmetric multiprocessor computers," *Appl. Opt.* **40**, 3371–3378 (2001).
3. A. F. J. Levi, "Optical interconnects in systems," *Proc. IEEE* **88**, 750–757 (2000).
4. D. A. B. Miller, "Rationale and challenges for optical interconnects to electrical chips," *Proc. IEEE* **88**, 728–749 (2000).
5. M. R. Feldman, S. C. Esener, C. G. Guest, and S. H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," *Appl. Opt.* **27**, 1742–1751 (1989).
6. E. D. Kyriakis-Bitaros, N. Haralabidis, M. Lagadas, A. Georgakilas, Y. Moisiadis, and G. Halkias, "Realistic end-to-end simulation of the optoelectronic links and comparison with the electrical interconnections for system-on-chip applications," *J. Lightwave Technol.* **19**, 1532–1542 (2001).
7. R. Grindley, T. Abdelrahman, S. Brown, S. Caranci, D. DeVries, B. Gamsa, A. Grbic, M. Gusat, R. Ho, O. Krieger, G. Lemieux, K. Loveless, N. Manjikian, P. McHardy, S. Srblic, M. Stumm, Z. Vranesic, and Z. Zilic, "The NUMachine multiprocessor," in *Proceedings of the International Conference on Parallel Processing* (IEEE Computer Society), 487–496 (2000).
8. A. V. Krishnamoorthy and D. A. B. Miller, "Firehose architectures for free-space optically interconnected VLSI circuits," *J. Parallel and Distributed Computing* **41**, 109–114 (1997).

9. T. H. Szymanski and H. S. Hinton, "Reconfigurable intelligent optical backplane for parallel computing and communications," *Appl. Opt.* **35**, 1253–1268 (1996).
10. T. H. Szymanski and H. S. Hinton, "Optoelectronic smart pixel array for a reconfigurable intelligent optical interconnect," U.S. Patent 6,016,211 (18 January 2000).
11. H. S. Hinton, T. J. Cloonan, F. B. McCormick, A. L. Lentine, and F. A. P. Tooley, "Free-space digital optical systems," *Proc. IEEE* **82**, 1632–1648 (1994).
12. Y. S. Liu, G. C. Boisset, M. H. Ayliffe, R. Iyer, and D. V. Plant, "Design, implementation and characterisation of a four stage hybrid optical system for a free-space optical backplane demonstrator," *Appl. Opt.* **37**, 2895–2914 (1998).
13. D. A. B. Miller, "Novel analog self-electro-optic-effect devices," *IEEE J. Quantum Electron.* **29**, 678–698 (1993).
14. M. B. Venditti, E. Laprise, J. Faucher, P.-O. Laprise, J. S. Ahearn, and D. V. Plant, "Design and verification of an OE-VLSI chip with 1080 VCSELs and PDs heterogeneously integrated with CMOS," *Proc. IEEE/LEOS 2001 Annual Meeting, PD-1.4* (2001).
15. A. L. Lentine, K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. J. Tseng, R. E. Leibenguth, J. E. Cunningham, W. Y. Jan, J.-M. Kuo, D. W. Dahringer, D. P. Kossives, D. D. Bacon, G. Livescu, R. L. Morrison, R. A. Novotny, and D. B. Buchholz, "High-speed optoelectronic VLSI switching chip with >4000 optical I/O based on flip-chip bonding of MQW modulators and detectors to silicon CMOS," *IEEE J. Sel. Top. Quantum Electron.* **2**, 77–84 (1996).
16. D. T. Neilson, "Optimization and tolerance analysis of QCSE modulators and detectors," *IEEE J. Quantum Electron.* **33**, 1094–1103 (1997).
17. D. Rolston, "The design, layout and characterization of VLSI optoelectronic chips for free-space optical interconnects," Ph.D. dissertation (McGill University, Montreal, Canada, 2000).
18. J. M. Sasian, R. A. Novotny, M. G. Beckman, S. L. Walker, M. J. Wojcik, and S. J. Hinterlong, "Fabrication of fiber bundle arrays for free-space photonic switching systems," *Opt. Eng.* **33**, 2979–2985 (1994).
19. C. V. Cryan, "Two-dimensional multimode fiber array for optical interconnects," *Electron. Lett.* **34**, 586–587 (1998).
20. T. Maj, A. G. Kirk, D. V. Plant, J. F. Ahadian, C. G. Fonstad, K. L. Lear, K. Tatah, M. S. Robinson, and J. A. Trezza, "Interconnection of a two-dimensional array of vertical-cavity surface-emitting lasers to a receiver array by means of a fiber image guide," *Appl. Opt.* **39**, 683–689 (2000).
21. D. M. Chiarulli, S. P. Levitan, P. Derr, R. Hofmann, B. Greiner, and M. Robinson, "Demonstration of a multichannel optical interconnection by use of imaging fiber bundles butt coupled to optoelectronic circuits," *Appl. Opt.* **39**, 698–703 (2000).
22. Y. Li, T. Wang, H. Kosaka, S. Kawai, and K. Kasahara, "Fiber-image-guide-based bit-parallel optical interconnects," *Appl. Opt.* **35**, 6920–6933 (1996).
23. R. T. Chen, L. Lin, C. Choi, Y. J. Liu, B. Bihari, L. Wu, S. Tang, R. Wickman, B. Picor, M. K. Hibb-Brenner, J. Bristow, and Y. S. Liu, "Fully embedded board-level guided-wave optoelectronic interconnects," *Proc. IEEE* **88**, 780–793 (2000).
24. F. B. McCormick, F. A. P. Tooley, T. J. Cloonan, J. M. Sasian, and H. S. Hinton, "Optical interconnects using microlens arrays," *Opt. Quantum Electron.* **34**, 6471–6480 (1992).
25. C. Berger, J. Ekman, X. Wang, P. Marchand, H. Spaanenburg, F. Kiamilev, and S. Esener, "Parallel distributed free-space optoelectronic computer engine using flat 'plug-on-top' optics package," *Optics in Computing 2000*, R. A. Lessard and T. V. Galstian, eds., *Proc. SPIE* **4089**, 1037–1045 (2000).
26. M. W. Haney, M. P. Christianson, F. Milojkovic, G. J. Fokken, M. Vickberg, B. K. Gilbert, J. Rieve, J. Erkman, P. Chandramani, and F. Kiamilev, "Description and evaluation of the fast-net smart pixel-based optical interconnection prototype," *Proc. IEEE* **88**, 819–828 (2000).
27. D. Fey, W. Erhard, M. Gruber, J. Jahns, H. Bartelt, G. Grimm, L. Hoppe, and S. Sinzinger, "Optical interconnects for neural and reconfigurable VLSI architectures," *Proc. IEEE* **88**, 838–848 (2000).
28. A. W. Lohmann, "Image formation of dilute arrays for optical information processing," *Opt. Commun.* **86**, 365–370 (1991).
29. 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.* **35**, 1220–1233 (1996).
30. H. Thienpont, C. Debaes, V. Baukens, H. Ottevaere, P. Vynck, P. Tuteleers, G. Verschaffelt, B. Volckaerts, A. Hermanne, and M. Hanney, "Plastic microoptical interconnection modules for parallel free-space inter- and intra-MCM data communication," *Proc. IEEE* **88**, 769–779 (2000).
31. D. T. Neilson and C. P. Barrett, "Performance trade-offs for conventional lenses in free-space digital optics," *Appl. Opt.* **35**, 1240–1248 (1996).
32. B. Robertson, "Design of an optical interconnect for photonic backplane applications," *Appl. Opt.* **37**, 2974–2984 (1998).
33. F. Lacroix, E. Bernier, M. H. Ayliffe, F. A. P. Tooley, D. V. Plant, and A. G. Kirk, "Implementation of a compact, four-stage, scalable optical interconnect for photonic backplane applications," *Appl. Opt.* **41**, 1541–1555 (2002).
34. F. K. Lacroix, "Design, analysis and implementation of free-space optical interconnects," Ph.D. dissertation (McGill University, Montreal, Canada, 2001).
35. 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).
36. B. Robertson, Y. Liu, G. C. Boisset, M. R. Taghizadeh, and D. V. Plant, "*In situ* interferometric alignment systems for the assembly of microchannel relay systems," *Appl. Opt.* **35**, 9253–9260 (1997).
37. M. H. Ayliffe, M. Chateaufneuf, D. R. Rolston, A. G. Kirk, and D. V. Plant, "Design and testing of a kinematic package supporting a  $32 \times 32$  array of GaAs modulators flip-chip bonded to a CMOS chip," *IEEE J. Lightwave Technol.* **19**, 1543–1599 (2001).
38. M. H. Ayliffe, "Alignment and packaging techniques for two-dimensional free-space optical interconnects," Ph.D. dissertation (McGill University, Montreal, Canada, 2001).
39. C. Alleyne and A. G. Kirk, "Transmission uniformity of diffractive parallel optical interconnect relays: A numerical analysis based on rigorous coupled wave theory," *Proc. 15th IEEE LEOS Annual Meeting, Vol 2*, 901–902 (2002).
40. D. T. Neilson and E. Schenfeld, "Free-space optical relay for the interconnection of multimode fibers," *Appl. Opt.* **38**, 2297–2300 (1999).
41. C. Debaes, M. Vervacke, H. Onevaere, W. Meeus, P. Tuteleers, M. Brunfaut, V. Baukens, J. Van Campenhout, H. Thienpont, "Demonstration of manufacturable free-space modules for multichannel intra-chip optical interconnects," *Proc. 15th IEEE LEOS Annual Meeting, Vol 1*, 63–64 (2002).
42. F. Thomas-Dupuis, M. Châteaufneuf, and A. G. Kirk, "Assembly and characterization of a folded spot array generator for a modulator-based free-space optical interconnect," *Proc. Intl. Topical Meeting on Optics in Computing, Vol. 1*, Taipei, Taiwan, 328–330 (2002).
43. M. Châteaufneuf, A. G. Kirk, D. V. Plant, T. Yamamoto, J. D. Ahearn, and W. Luo, "512-channel vertical-cavity surface-emitting laser based free-space optical link," *Appl. Opt.* **41**, 5552–5561 (2002).
44. D. V. Plant, M. B. Venditti, E. Laprise, J. Faucher, K. Razavi,

- M. Chateauneuf, A. G. Kirk, and J. D. Ahearn, "A 256 Channel Bi-Directional Optical Interconnect Using VCSELs and Photodiodes on CMOS," *J. Lightwave Technol.* **19**, 1093–1103 (2001).
45. A. G. Kirk, F. Lacroix, F. Mathieu, and F. Tooley, "Demonstration of a free-space optical broadcast network," in *Optics in Computing*, Vol. 8, 1997 OSA Technical Digest (Optical Society of America, Washington, D.C., 1997), pp. 165–167.
46. J. A. B. Dines, D. T. Nelson, J. F. Snowdon, and B. S. Wherrett, "A comparison of massively parallel interconnect topologies employing optical highways," in *Optics in Computing*, Vol. 8, 1997 OSA Technical Digest (Optical Society of America, Washington, D.C., 1997), pp. 186–188.
47. T. H. Szymanski and V. Tyan, "Error and flow control for a terabit free-space optical backplane," *IEEE J. Sel. Top. Quantum Electron.* 838–846 (1999).
48. T. H. Szymanski, "Bandwidth optimization of optical datalinks using error control codes," *Appl. Opt.* 1761–1775 (2000).
49. J. Faucher, M. B. Venditti, E. Laprise, and D. V. Plant, "Application of parallel forward error correction in two-dimensional optical data links," *IEEE J. Lightwave Technol.* (to be published) (2003).
50. C.-H. Chen, B. Hoanca, C. B. Kuznia, A. A. Sawchuk, and J.-M. Wu, "TRANslucent smart pixel array (TRANSPAR) chips for high throughput networks and SIMD signal processing," *IEEE J. Sel. Top. Quantum Electron.* **5**, 316–329 (1999).