Experimental demonstration of the optical multi-mesh hypercube: scaleable interconnection network for multiprocessors and multicomputers

Ahmed Louri, Stephen Furlonge, and Costas Neocleous

A prototype of a novel topology for scaleable optical interconnection networks called the optical multi-mesh hypercube (OMMH) is experimentally demonstrated to as high as a 150-Mbit/s data rate ($2^7 - 1$ nonreturn-to-zero pseudo-random data pattern) at a bit error rate of 10^{-13}/link by the use of commercially available devices. OMMH is a scaleable network [Appl. Opt. **33,** 7558 (1994); J. Lightwave Technol. **12,** 704 (1994)] architecture that combines the positive features of the hypercube (small diameter, connectivity, symmetry, simple routing, and fault tolerance) and the mesh (constant node degree and size scaleability). The optical implementation method is divided into two levels: high-density local connections for the hypercube modules, and high-bit-rate, low-density, long connections for the mesh links connecting the hypercube modules. Free-space imaging systems utilizing vertical-cavity surface-emitting laser (VCSEL) arrays, lenslet arrays, space-invariant holographic techniques, and photodiode arrays are demonstrated for the local connections. Optobus fiber interconnects from Motorola are used for the long-distance connections. The OMMH was optimized to operate at the data rate of Motorola's Optobus (10-bit-wide, VCSEL-based bidirectional data interconnects at 150 Mbits/s). Difficulties encountered included the varying fan-out efficiencies of the different orders of the hologram, misalignment sensitivity of the free-space links, low power (1 mW) of the individual VCSEL's, and noise.

Key words: Hypercube, mesh, interconnection network, optical interconnects, parallel computing, scaleability, VCSEL's.

1. Introduction

Two key issues and desirable features for the construction of performant interconnection networks for parallel-processing systems are modularity (the construction of a larger network from several smaller networks) and scaleability (the increasing of the size of the network with nominal changes in the existing configuration but with a comparable increase in performance). Two popular point-to-point interconnection networks for parallel computers are the binary n-cube, or hypercube, and the mesh networks. The binary n-cube has $N = 2^n$ nodes, each of node degree n. The attractiveness of the hypercube topology is its small diameter ($\log_2 N$), high connectivity, symmetric and regular nature, simple and efficient routing algorithms, and fault tolerance. A drawback of the hypercube, however, is its lack of scaleability, which limits its use in building large-size systems out of smaller-size systems. The lack of scaleability of the hypercube stems from the fact that the node degree is not bounded and varies as $\log_2 N$. This property makes the hypercube cost prohibitive for large N.

The mesh, because of its simple regular connection and constant node degree (four), is easily implemented and highly size scaleable. For a network size of N nodes, the minimal incremental size is approximately $\sqrt{N}$. A major limitation however is its large diameter ($\sqrt{N}$).

Combining the positive features of the hypercube (small diameter, regularity, high connectivity, simple control and routing, symmetry, and fault tolerance) with those of a torus or wraparound mesh (constant node degree and size scaleability) has enabled the exploration of a novel scaleable interconnection network for a parallel-processing system called the optical multi-mesh hypercube (OMMH), and its optical design methodology has also been explored.[1,2,3,4]

The authors are with the Department of Electrical and Computer Engineering, the University of Arizona, Tucson, Arizona 85721.

Received 12 February 1996; revised manuscript received 12 June 1996.

0003-6935/96/356909-11$10.00/0

In this paper, the first version of an experimental prototype of the OMMH is demonstrated (by use of commercially available components) to verify its feasibility. By a first-version prototype, we mean a small representation of the system to characterize the conceptual design and to demonstrate that the required connectivity for the OMMH can be achieved with optical interconnects. The experimental setups are included to show how we characterized the components and overall behavior of the laboratory prototype. Driver and receiver circuits were built, along with test and interface boards for both the optics and the electronics. The results relating to the optical geometry that have a direct impact on the system scaleability and feasibility, parameters relating to the hologram that implements the space-invariant fan-out, the power requirements, interconnect media (fiber and free space), the receiver design, and driver circuits are presented and analyzed. The system parameters include the measurement of the power budget, mark ratio (number of 1's in a data pattern flowing through the network), data rate, component size, link efficiency, interconnect distances, bit error rate (BER), noise, and eye diagrams.

2. Optical Multi-Mesh Hypercube Networks

In this section we briefly define the structure of the OMMH. We also give some general features. More details can be found in Refs. 1–4.

A. Definition of Optical Multi-Mesh Hypercube Networks

An OMMH is characterized by a triplet (l, m, n), where l represents the row dimension of a torus, m represents the column dimension of the torus, and n represents the dimension of a hypercube. For two nodes (i_1, j_1, k_1) and (i_2, j_2, k_2), where $0 \le i_1 < l$, $0 \le i_2 < l$, $0 \le j_1 < m$, $0 \le j_2 < m$, $0 \le k_1 < 2^n$, and $0 \le k_2 < 2^n$, there is

1. A torus link between two nodes if (a) $k_1 = k_2$ and (b) two components, i and j, differ by 1 in one component, while the other component is identical.
2. A torus link for the wraparound connection in the row if (a) $k_1 = k_2$ and (b) $i_1 = i_2, j_1 = 0, j_2 = m - 1$, or, for the wraparound connection in the column, if (a) $k_1 = k_2$ and (b) $j_1 = j_2$, $i_1 = 0$, $i_2 = l - 1$.
3. A hypercube link between two nodes if and only if (a) $i_1 = i_2$, (b) $j_1 = j_2$, and (c) k_1 and k_2 differ by one bit position in their binary representation.

Figure 1 shows a (2, 2, 3)-OMMH interconnection, where the solid lines represent hypercube links and the dashed lines represent torus links. A (2, 2, 3)-OMMH consists of $2 \times 2 \times 2^3$ nodes = 32 nodes. The solid circles represent nodes of the OMMH network. Both ends of the torus links (dashed lines) are connected for wraparound connections if they have the same labels. In general, an (l, m, n)-OMMH contains $l \times m \times 2^n$ nodes.

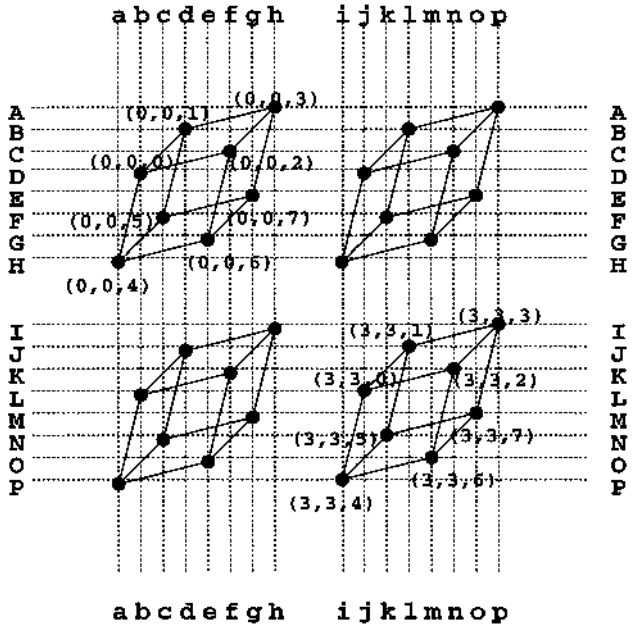

Fig. 1. Sample (2,2,3)-OMMH network with four hypercubes, eight meshes, and 32 nodes.

B. Size Scalability of the Optical Multi-Mesh Hypercube Network

The size of the OMMH can grow without altering the number of links per node by the expansion of the size of the torus; for example, by the insertion of 3-cubes in the row or column of the torus in Fig. 1. This feature allows the OMMH to be size scaleable and modular, as it permits the incremental expansion of small sizes compared with the doubling of the size of the hypercube. More information on the scaleability and performance analysis of the OMMH can be found in Refs. 1 and 2.

3. Optical Implementation of Optical Multi-Mesh Hypercube Networks

A. Background

The design method for implementing the OMMH network is based on two optical technologies: space-invariant free-space optics and fiber interconnects. The architecture can be viewed as a two-level system: a local-connection level representing a set of hypercube modules and a global-connection level representing the torus network that connects the hypercube modules. Because the local-connection level is geometrically compact and uses relatively short-distance connections, an extensive use of space-invariant free-space optics is used for its implementation.

Space-invariant free-space optical systems are simple to build and mass produce. This has a direct impact on the modularity aspect of the architecture. The global connections, representing the torus network and connecting the hypercube modules, are high-bit-rate, long-distance connections. Therefore, fiber-optic arrays are used for their implementation.

An additional factor in using fiber-optic arrays for torus connections is that wavelength-division multiplexing techniques can be applied to reduce the number of fibers as the network grows in size and to make better use of the transmission capacity of fiber interconnects.[5–8] More importantly, commercial fiber-optic arrays such as Motorola's Optobus (Refs. 9 and 10) are readily available.

The design methodology for the OMMH has been rigorously analyzed in Refs. 1–3, 11, and 12. Additionally, a feasibility study on the theoretical modeling of the OMMH was developed to predict the size, data rate, bit error rate, power budget, noise, efficiency, interconnect distance, pixel density, and misalignment sensitivity.[4] The numerical predictions were validated with experimental data from commercially available products to assess the effects of various thermal, systemic, and geometric parameters on the behavior of the sample model. Issues relating to the optical geometry that had a direct impact on system scaleability and feasibility, packaging and alignment, the holographic beam-steering element that implements the space-invariant fan-out, the power requirements, the interconnect medium (fiber and free space), and the receiver-design considerations were all examined. The results of the feasibility study indicated that the implementation of the OMMH was quite feasible in terms of the parameters outlined. In this paper we present the results of the experimental prototype that was developed in the laboratory to validate the theoretical modeling.

From the design methodology[1–3,11,12] the nodes belonging to a hypercube are partitioned into two sets such that any two nodes in the same set do not have a direct link. We call the two sets plane L and plane R. As an implementation example we chose to demonstrate the connectivity of a (2, 2, 3)-OMMH, as shown in Fig. 1. The conceptual view of the model for the sample OMMH is shown in Fig. 2. For clarity, the free-space and fiber links are shown separately. An imaging system is used to provide the connectivity with a space-invariant hologram between planes L and R, as shown in Fig. 2(a). A space-invariant optical interconnection module consists of an imaging system and a space-invariant hologram. The nodes of plane L are then interconnected by fibers to achieve the mesh connectivity shown in Fig. 2(b). Similarly the nodes of plane R are connected because both planes have identical connections.

B. Brief Component Description

This subsection comprises a list of components used in the experimental OMMH.

1. *Optobus*

For the fiber interconnects in the OMMH we used commercially available optical-fiber links (Optobus) from Motorola.[9,10] The Optobus is a 10-bit-wide, vertical-cavity surface-emitting laser- (VCSEL-) based bidirectional data-interconnect solution for point-to-point applications that utilizes the attributes of optical–electronic technology and offers a cost-effective system solution for transporting massive amounts of data. The initial offering of this device is 150 Mbytes/s, with the VCSEL's operating at 850 nm, a power-supply voltage of 5 V, and a BER of 10^{-13}. Later versions of the Optobus are expected to operate with a BER of less than 10^{-17} and data rates in excess of 300 Mbytes/s.

The transceiver modules consist of two completely independent transmitter and receiver subsystems. Although these two subsystems share a common substrate, they do not share a common power or ground. The transmitter subsystem consists of a transmitter waveguide unit with 10 optical outputs plus a single complementary metal-oxide semiconductor custom-laser-driver integrated circuit. The receiver subsystem contains a receiver array and two buffer arrays. The receiver is connected by the use of a simple transimpedance amplifier. The optical interface is a waveguide unit combined with an optical connector (MT) and a 10-channel fiber ribbon. For our demonstration we use four Optobuses to implement the fiber connectivity of plane L of the sample OMMH.

2. *Vertical-Cavity Surface-Emitting Laser Array*

For transmitters in the free-space optical interconnects of the OMMH, considerable attention has been focused on laser structures designed to minimize the threshold current, achieve surface emission, and maximize the external quantum efficiency and the frequency response.[13] VCSEL arrays are attractive light sources for optical interconnects because they exhibit a low threshold current, low voltage (<2 V), high speed (10 GHz), and high conversion efficiency, produce a single longitudinal mode, emit a low-divergence, circularly symmetric, aberration-free beam,[14,15,16] and are commercially available.[17] Four 2×2 VCSEL arrays (Model SN 2716, Vixel Corporation) are used for the experimental OMMH. These VCSEL's are AlGaAs semiconductor lasers, and they operate at 760 nm with a standard light output of 1 mW at a threshold current of 7 mA. The diameter of each laser diode within the array is 8 μm, and the separation between them is 125 μm.

The VCSEL's are driven by a resistor array, as shown in the upper right-hand side of the schematic diagram of the optical interconnection (Fig. 3). The light emitted from the VCSEL's is collimated by the use of a lenslet array. The diameter of each collimating lens is 125 μm, with a focal length of 1 mm.

3. *Receiver Array*

The following design considerations for optical receivers are desired: high receiver sensitivity, wide dynamic range, and transparency to the operating bit rate. The optical-receiver performance is a function of both the photodetector and the preamplifier. Together these two elements dictate many of the receiver characteristics, as well as its performance.[18,19]

The sensitivity of a receiver is defined in terms of the received optical power required to achieve a desired BER. Receiver sensitivity is essential in opti-

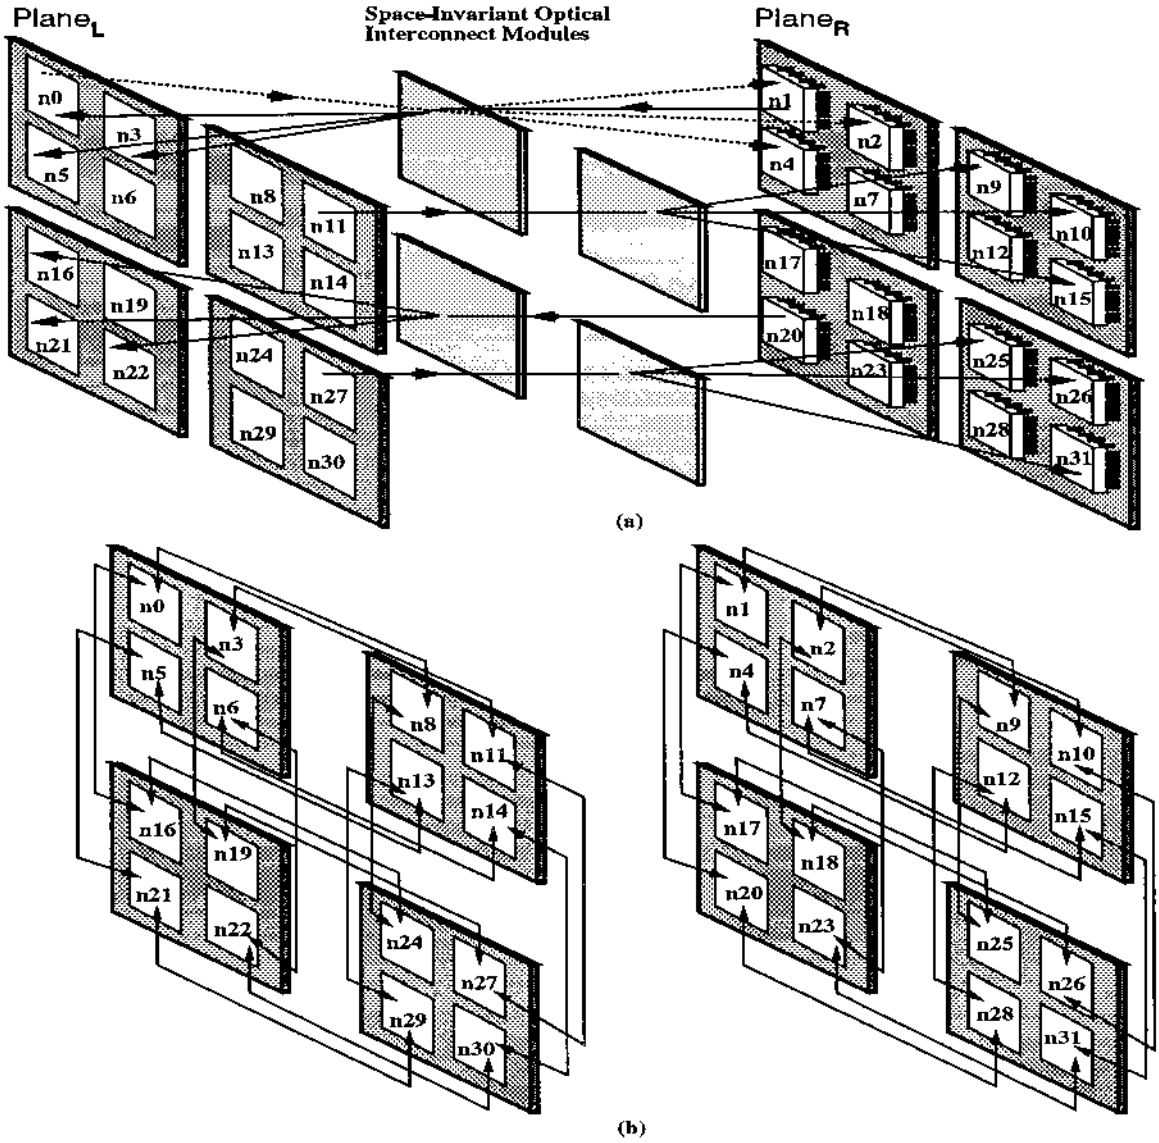

Fig. 2. Three-dimensional view of a (2,2,3)-OMMH showing plane L and plane R interconnected by four space-invariant optical interconnection modules. (a) The free-space links are shown separated from (b) the mesh links for clarity of the figure. Note that there are four hypercubes interconnected by fibers. Hypercube connectivity is achieved through free-space space-invariant interconnection modules. For example, module $n24, n27, n29, n30$ from plane L and module $n25, n26, n28, n31$ from plane R belong to the same hypercube. The connection patterns for the free-space links are all identical in both directions, as shown by the solid and the dashed arrows.

cal interconnects because it determines the power budget of the system.

The receiver dynamic range is the difference in decibels between the minimum detectable power level (receiver sensitivity) and the maximum allowable input power level. A wide dynamic range is important because it permits flexibility and convenience in the system configuration. The ability to deal with a wide range of optical power levels at the receiver means that we can accommodate any instabilities that might appear in the network that are due to environmental factors such as heating, temperature variations, aging, and laser-diode degradations. Bit-rate transparency refers to the ability of the optical receiver to operate over a range of bit rates.

Photodiodes: The important requirements on the photodiodes are the following: a high quantum efficiency, fast response time, low capacitance, and low dark current.[18] The performance of photodiodes and the trade-offs in relation to bandwidth have previously been modeled.[4] For our system we use Hamamatsu silicon photodiodes, which have a responsivity of $R = 0.5$ A/W, a dark current value of $I_d = 10^{-11}$ A, and a terminal capacitance of $C_t = 0.5$ pF at a reverse voltage of 10 V. For our (2, 2, 3)-OMMH network four such silicon quadrant photodiodes are used. Each element of the quadrant has an effective diameter of 250 μm.

Receiver Amplifier: Front amplifiers are classified into two types: high impedance (HZ) and transimpedance (TZ). The HZ amplifier offers the lowest noise level and therefore the lowest detection sensi-

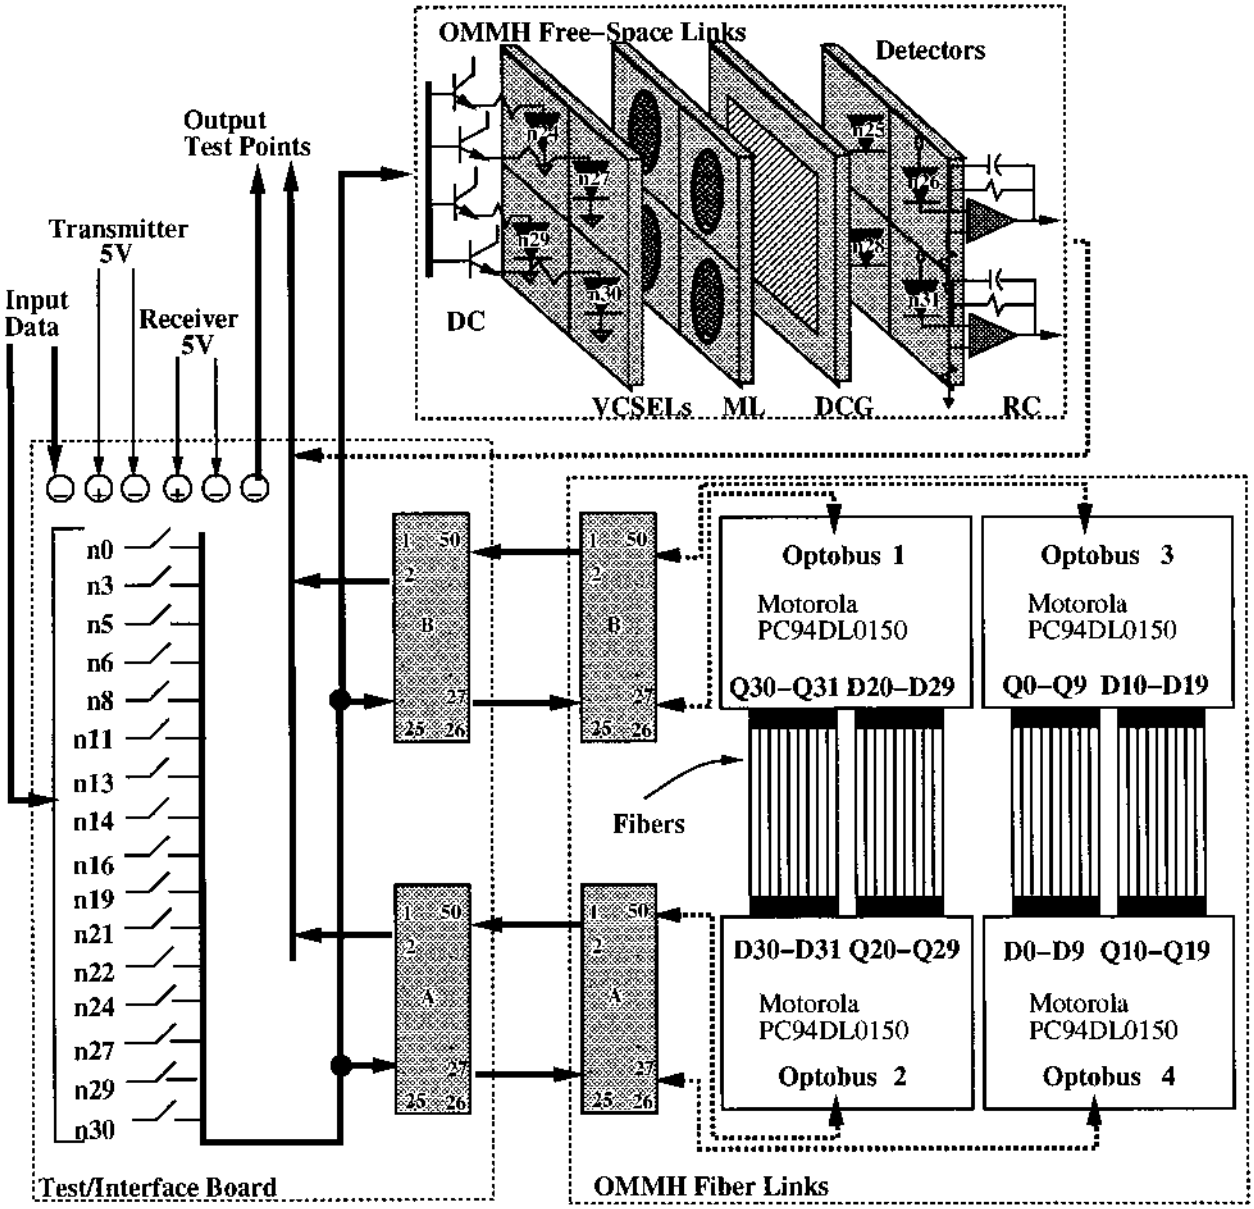

Fig. 3. Setup for the experimental OMMH. The test-interface board interfaces the free-space and fiber links of the OMMH, providing the power-ground, input data, and output test points. The fiber links of plane L are implemented by four Optobus modules. The free-space links for a single hypercube are also shown. Definitions: VCSEL, vertical-cavity surface-emitting laser; D, transmitter; Q, receiver; A and B, 50-pin latch–eject connectors; DCG hologram, dichromated-gelatin hologram; DC, VCSEL drive circuit; RC, receiver circuit; ML, microlens array. Symbols: , switch; , flat ribbon cable; , wire wrap; , 10-channel fiber ribbon.

tivity. Nevertheless, because of the high load impedance at the front end, the frequency response is limited by the RC time constant at the input. Another major drawback of the HZ amplifier is its limited dynamic range.[18]

The TZ amplifier is a popular approach to avoiding the dynamic-range problem. In addition, with the TZ configuration we can achieve a much higher bandwidth. Because of the above advantages of the TZ amplifier we chose to use this configuration for our receiver circuitry in the free-space optical interconnects in the OMMH. A wide-band operational amplifier (Model CLC 425) from Comlinear Corporation was used. This operational amplifier features a gain-to-bandwidth product of 1.9 GHz, an input-voltage noise level of 1.05 $nV/\sqrt{Hz}$, and a skew rate of 350 $V/\mu S$. The receiver amplifiers were fabricated in the laboratory on special boards[20] designed to operate at high frequencies. For compensating for noise from the power supplies a pair of 6.8-μF tantalum capacitors were used across both power supplies (−5 V and +5 V).

4. *Space-Invariant Hologram*

The beam-steering element that was used in the implementation of the OMMH network is a Dammann

grating[21] that was fabricated on a hologram. Dichromated-gelatin (DCG) holograms were developed for the OMMH. DCG holograms are phase holograms and can provide very high diffraction efficiencies.[22,23] They require a very stable environment during exposure time, and they can be affected by humidity. First the DCG emulsion was fabricated in the laboratory, and then the hologram was recorded by use of the procedure followed in Ref. 22. Diffraction efficiencies of 12% were achieved for the ± 1 orders (discussed below in Subsection 5.C). When the hologram was designed emphasis was placed on equalizing the optical power of each output order. Nevertheless, as a result of fabrication errors, more power was coupled into the zero order. Approximately 30% of the input power was coupled into the zero order and 48% into the ± 1 orders (both vertical and horizontal). The rest of the power was coupled into higher diffraction orders.

C. Optical-System Setup

As can be seen from Fig. 2, the proposed optical implementation is symmetric and modular. The mesh links for plane L are identical to those of plane R. Additionally, the free-space interconnect patterns of plane L are also identical to those of plane R. As a result, we chose to physically implement one plane, say plane L of Fig. 2(b), to demonstrate the fiber links, and one hypercube defined by nodes ($n24$, $n27$, $n29$, $n30$) on plane L and ($n25$, $n26$, $n28$, $n31$) on plane R to demonstrate the free-space links. This setup is sufficient to demonstrate the OMMH connectivity, as is shown in the example in Subsection 5.C. Figure 3 shows the experimental OMMH for plane L of Fig. 2(b) and the single hypercube connecting planes L and R of Fig. 2(a). A photograph of the laboratory prototype is shown in Fig. 4.

Data to the network are input to a test-interface board (Fig. 3), which subsequently routes it to both the fiber and free-space interconnects to achieve the desired connectivity of the OMMH. The test-interface board contains leads for the input data signal from the pseudo-random data transmitter, the power and ground connections of the laser drivers and receiver circuitry of the OMMH links, switches corresponding to individual nodes labeled in Fig. 2, connections from the fiber and free-space links to the output test points, and leads from the output test points to the Anritsu BER receiver and Hewlett Packard high-bandwidth oscilloscope. The purpose of the test points is to allow us to probe the network with the test equipment at individual nodes to examine the data signal that is being received by that node and subsequently being sent to other nodes. We can then characterize the data transfer over the network in terms of the data rate, noise, BER, mark ratios, and eye diagrams.

A node in the theoretical network can be either a processing element, memory module, or switch. For implementation a switch in conjunction with dedicated VCSEL transmitters and detectors represent a node. The input data signal is first routed to

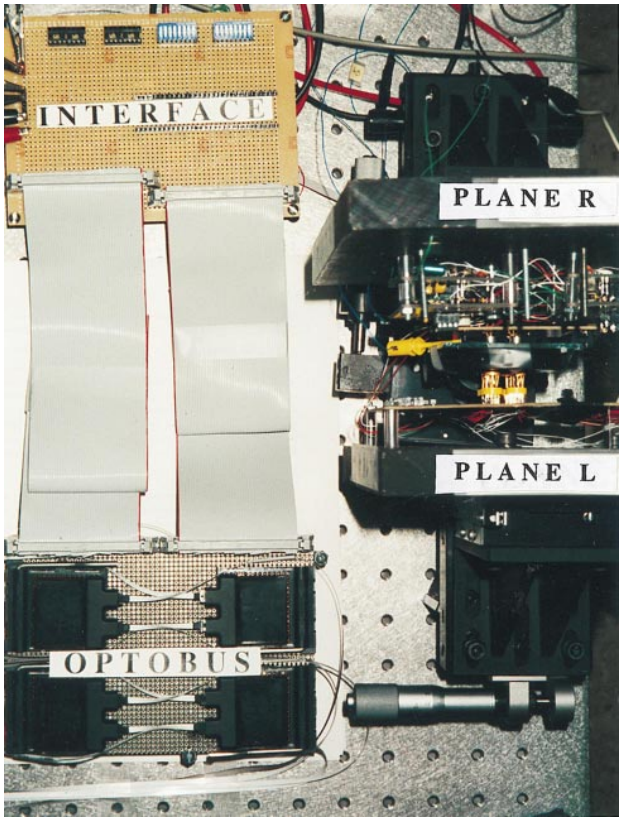

Fig. 4. Photograph of the laboratory prototype of the OMMH.

switches on the test-interface board, which allows data to flow through that node and subsequently onto the fiber and free-space links dictated by the connectivity of the OMMH. For example, if the switch representing node $n30$ is turned on, then data signals sent from the VCSEL's for that node should be received by the detectors for nodes $n14$ and $n22$ via the fiber links. At the same time, nodes $n26$, $n28$, and $n31$ should receive data via the free-space links of the OMMH. In addition, the switches facilitate individual testing of each link in characterizing the overall behavior of the network. Individual testing would determine which links are functional and allow us to measure the parameters outlined above.

Thus, when a switch is turned on, it indicates that the node associated with that switch desires to communicate with other nodes, as determined by the connectivity of the OMMH. The signal is then routed to the Optobus from the test–interface board by connectors A and B (50-pin latch–eject connectors from Digi-Key Corporation). For clarity in Fig. 3, only a few pin numbers are shown on A and B. The connectors latch onto flat ribbon cables that connect the test–interface board to the board for the OMMH fiber links. Simultaneously, the signal is also sent from the test–interface board to the VCSEL associated with that node and needed to establish the free-space connectivity on the board for the OMMH free-space links.

Each fiber link shown in the plane L of Fig. 2 is represented by an Optobus link (D–Q pair), where D

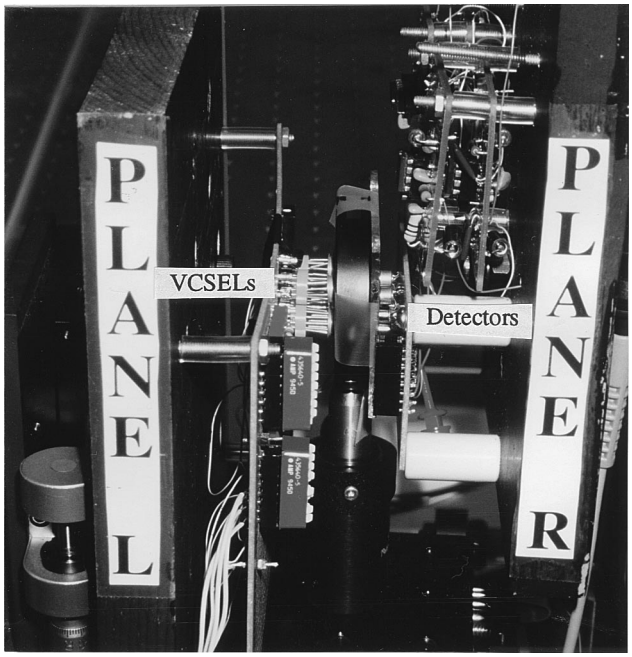

Fig. 5. Photograph of the laboratory prototype of the free-space links in the OMMH.

represents the transmitter input and Q the corresponding receiver output. The modulated data from the transmitter (D) module of the Optobus is then sent along the corresponding Optobus fiber link to the detector (Q) on the receiver module of the other Optobus. A power-supply voltage of 5 V with separate power and ground for the transmitters and receivers is used. The input-voltage swing of the Optobus is between 0.25 and 0.8 V.

A detailed representation of the links for the single free-space unidirectional hypercube is shown in Fig. 3. The VCSEL is modulated by the pseudo-random input data by means of the laser-driver circuitry. The modulated light from the VCSEL is then collimated by the microlens (ML) onto the DCG hologram, where the space-invariant fan-out operation is performed on the detectors of the receiving nodes. Two power supplies of ± 5 V were used to power the receiver circuits, and another power supply at $+5$ V was used to power the drive circuits. The input-voltage swing for the free-space interconnects is between 2 and 4 V. A photograph taken in the laboratory of the free-space links is shown in Fig. 5. Decision thresholding, amplification, and conversion take place from optical to electrical at the detectors for both the fiber and the free-space links, and the signals are subsequently sent to the output-test–probe points of the test–interface board.

4. System Measurements

In this section we briefly define the critical system measurements. We also define the experimental setup and the equipment used to generate and measure them.

A. Measurement Parameters

For the experimental demonstration we measure BER's, eye diagrams, and data rates for the OMMH. The BER is defined as the probability of an error per bit. A BER of 10^{-17} means an average of one error every 10^{17} bits. To determine the BER in our system, two parameters need to be measured: the error rate and the error count. Both are measured with the BER equipment discussed below in Subsection 4.B.

Because of the noise inherent in any receiver, there is always a finite probability that a bit will be incorrectly identified by the decision circuit. This probability can be determined from the eye diagram formed by the superposition of the electrical pulses corresponding to different bits on top of each other. The eye diagram provides a visual way to monitor the link performance. Closing of the eye is an indication that the receiver is not performing properly because of noise and timing jitter. A pseudo-random data generator and a high-bandwidth oscilloscope (discussed Subsection 4.B) is used to determine the eye diagrams for the OMMH connectivity.

The pseudo-random pattern (PRBS) from the data generator has characteristics that strongly resemble an actual line signal. It is used to generate the eye diagram of each link. The Anritsu Model ME522A transmitter is used to generate a PRBS signal ($2^7 - 1 = 127$-bit period) at a data rate of 150 Mbytes/s with non-return-to-zero (NRZ) coding. In addition, the PRBS pattern is set such that the mark ratio (ratio of 1's in the bit pattern) can be set to 1/2 and to 7/8. It is critical to be able to send data patterns of varying mark ratios because this will determine the reliability of the link to send varying data accurately.

B. Experimental Setup for Measurements

For generating the eye diagram for each link, first the ME522A transmitter is set to transmit data at 150 Mbytes/s, preprogrammed to the NRZ for different PRBS data patterns and mark ratios. The transmitter clock is connected to channel 1 of the oscilloscope, and the data from the transmitter are sent through the input of the test–interface board and subsequently through the fiber and free-space links that establish the required connectivity of the OMMH. The data are then routed back to the test–interface board and then to channel 2 of the oscilloscope.

The experimental setup for the BER measurement is the same as that for generating the eye diagram, except that the Hewlett Packard Model 54522A oscilloscope is replaced with an Anritsu ME522A BER receiver. The basic idea is to compare the PRBS pattern generated by the transmitter, and, after data are cycled through the OMMH, to check the number of errors detected by the receiver. The ME522A receiver has to be preprogrammed to the PRBS pattern sent by the ME522A transmitter to be able to compare the patterns and generate the error count and the error rate. In addition, we photograph patterns of light beams from the VCSEL's after they have

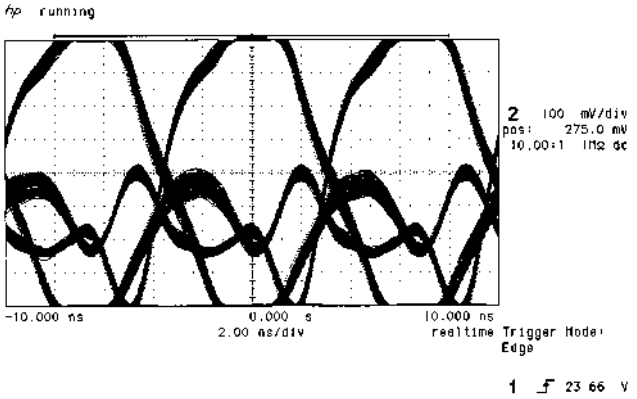

Fig. 6. Noise received at node 11 (FIBER), with the eye diagram superimposed to show the signal and the noise.

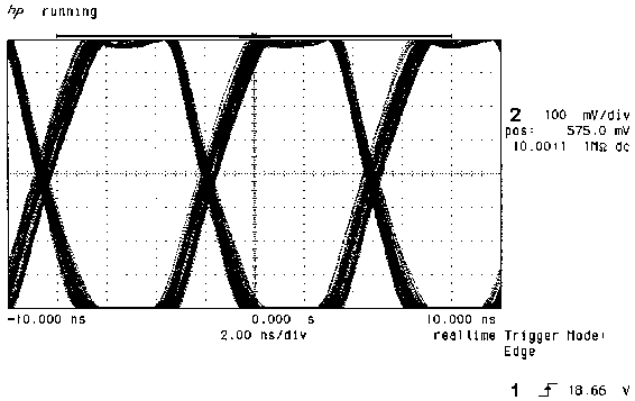

Fig. 7. Measured eye pattern for node $n27$ transmitting to node $n26$ (FREE-SPACE) at a 150-MHz data rate ($2^7 - 1$ NRZ pseudorandom data pattern). This node receives from the zero order of the hologram, and therefore the eye opening is greater than the nodes receiving from the ± 1 orders. The eye opening is 0.8 V.

fanned out from the hologram plane and entered the detector plane to show the required free-space connectivity of the OMMH.

5. Experimental Results

In this section we first present the individual results of both the fiber-mesh connectivity and the free-space hypercube connectivity of the OMMH. We then document an experimental demonstration of the OMMH connectivity, using a combination of fiber and free-space links by way of an example.

A. Fiber-Mesh Links in the Optical Multi-Mesh Hypercube

For the fiber links in the OMMH, the typical output-voltage swing is 0.25 V, with a channel-to-channel skew of $\approx$0.5 ns. A 150-MHz clock signal, which can be applied simultaneously to all 32 channels, is used to determine the mesh connectivity and functionality. After the link is determined to be functional and the mesh connectivity is provided by the Optobus modules, data can be applied. Data at 150 Mbytes/s have been transmitted over each of the 32 data channels with error rates of less than 10^{-13}. With the channel off, the measured noise is superimposed with the eye diagram to show the difference between the signal and the noise (Fig. 6). There is no contention for node $n3$ as there are separate links for communication. Additionally, we have measured the eye opening for node $n30$ transmitting to node $n22$ via the fiber links at 150 Mbytes/s, with a mark ratio of 1/2 at 0.4 V. A mark ratio of 7/8 yielded an opening of 0.8 V. We observe that, for higher mark ratios, the eye opening is larger, which tells us that data patterns with varying mark ratios can be transmitted.

For the BER measurement, we used a single channel with the links operating over a period of 18.5 h. The error rate was measured at 0.5, with an error count of 4.3×10^{12} yielding a BER of $\approx 10^{-13}$.

B. Free-Space Hypercube Links in the Optical Multi-Mesh Hypercube

For the free-space links in the OMMH, the total interconnection length was 4 mm. The total interconnection length basically depends on the focal length of the lenslet array. The above length value does not include the thickness of the holographic plate and the microlens substrate. If these thicknesses are introduced, then the total length of the hypercube interconnect is 10 mm.

Data at a rate of 150 Mbytes/s have been transmitted over the hypercube links. One unidirectional hypercube module was aligned, and every node of plane L was able to transmit to three nodes in plane R to establish the hypercube connectivity of the OMMH.

The eye diagrams of node $n27$ transmitting to nodes $n26$, $n25$, and $n31$ are shown in Figs. 7–9. Node $n26$ corresponds to the zero order coming out of the hologram. Because more power was coupled into the zero order the eye opening of node $n26$ is approximately 0.8 V, whereas for the other two orders the eye opening is 0.4 V for node $n31$ and 0.16 V for node $n25$. Figure 10 represents the noise measured at the output of node $n26$ superimposed with the signal to show the difference between signal and

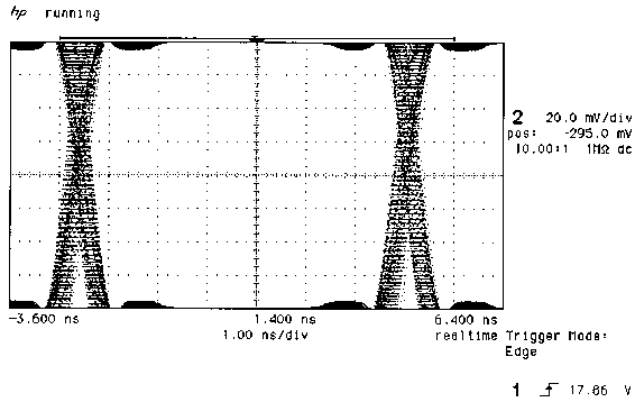

Fig. 8. Measured eye pattern for node $n27$ transmitting to node $n25$ (FREE-SPACE) at a 150-MHz data rate ($2^7 - 1$ NRZ pseudorandom data pattern). The eye opening is 0.16 V.

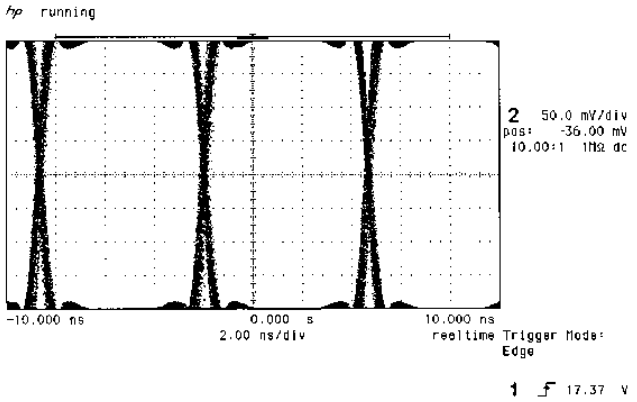

Fig. 9. Measured eye pattern for node $n27$ transmitting to node $n31$ (FREE-SPACE) at a 150-MHz data rate ($2^7 - 1$ NRZ pseudo-random data pattern). The eye opening for this node is 0.4 V.

noise. All the hypercube eye diagrams were taken at a mark ratio of 1/2. To measure the BER, the method used for the mesh links was applied. Node $n26$ yielded a BER less than 10^{-13}, whereas the two other nodes yielded a BER of the order of 10^{-13}.

To align the system a couple of x–y micrometer positioners were used. One was used to move the VCSEL's and the other the detectors, whereas the lenslet array and the hologram were stable. First, the VCSEL's were aligned with the lenslet array and the holograms and second with the detectors.

C. Experimental Demonstration of the Optical Multi-Mesh Hypercube Connectivity

As an example of the OMMH connectivity we show the detailed communication of node $n3$ with node $n31$ in the OMMH. Figures 2(a) and 2(b) show the intermediate nodes and labels for the links that establish such a connection. To establish such a link, it is first necessary that node $n3$ connect with node $n11$ and send the data pattern ($2^7 - 1$ NRZ) via the fiber-mesh link. The modulated signal at node $n11$ is then sent from the VCSEL for node $n11$ to node $n27$, where it is received via another fiber-mesh link by the photodiode and receiver circuitry for node $n27$. The

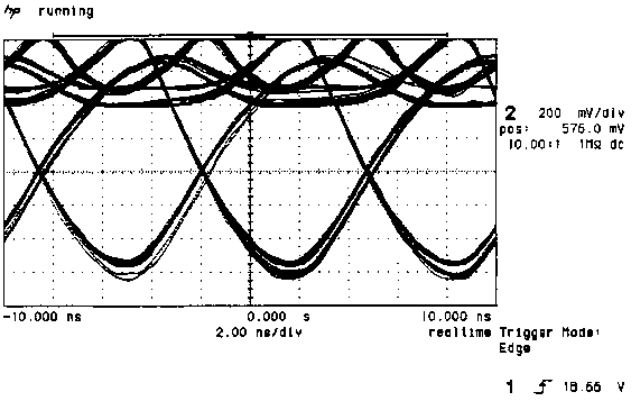

Fig. 10. Measured noise at node $n26$ superimposed with the eye diagram to show the difference between the signal and the noise.

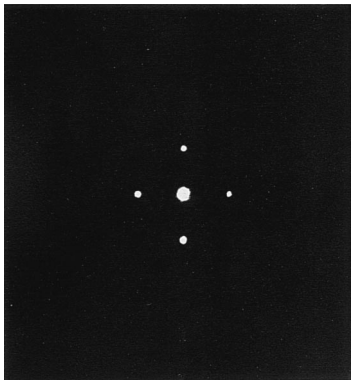

Fig. 11. Photograph of the space-invariant fan-out pattern implementing free-space connectivity in the OMMH.

eye diagrams for the respective fiber links are similar to the eye diagram shown in Fig. 6. The opening of the eye is approximately 0.8 V, which is consistent with the output-voltage swing of the Optobus.

Node $n27$ then sends the data pattern as modulated light from the VCSEL on node $n27$ to a collimating lenslet array (ML) and subsequently to the hologram, where the space-invariant fan-out operation is performed. The fan-out pattern leaving the hologram is shown in the photograph in Fig. 11. Higher orders also emanated, as shown, because the hologram acts as a diffraction grating. Light from these orders does not fall on the detectors for nodes $n25$, $n26$, $n28$, or $n31$. It is therefore lost power and needs to be removed as heat if compact systems need to be built.

The beams exiting the hologram are incident on the photodiodes and receiver circuitry for nodes $n25$, $n26$, and $n31$. The corresponding eye diagrams for these nodes and the free-space links were shown in Figs. 7–9 in Subsection 5.B.

The power budget for the free-space link for nodes $n27$ to $n25$ is shown in Table 1 for the ± 1 orders. A measured laser power of 1 mW exiting the VCSEL and incident on the hologram after some minor losses at the lenslet array is split into 0.12 mW for the ± 1 orders and approximately 0.3 mW for the zero order. Power detected at the detector plane fluctuated and was averaged at 45 μW for the ± 1 orders. To balance the power-budget table we therefore had to add in some engineering margin at -3 dB to account for quantum efficiency and other inaccuracies in measurements. No power budget was done for the fiber links because Motorola's Optobus is a packaged product in which all the optical-signal degradation is done internally.

6. Discussion

In this section we discuss the problems and difficulties of the implementation of the OMMH network and possible solutions for improving the performance of a future OMMH. We also describe the limitations as to the size scaleability of the OMMH. The OMMH was optimized to operate at 150 Mbytes/s as a result of the limitation imposed by the Optobus.

Table 1. Free-Space Link Power Budget

Parameter	Value	Power Ratio (dBm)
VCSEL power	1 mW	0
Coupling-efficiency lens (ML)	0.999	-4.34×10^{-3}
Hologram coupling efficiency	0.12	−9.208
Coupling efficiency into the detector	0.76	−1.367
Engineering margin		−3
Power incident on the detector	45.55 μW	−13.415
Detector quantum efficiency	0.5	
Photocurrent at the receiver	33.35 μA	

Difficulties included noise, misalignment sensitivity of the free-space links, low power (1 mW) of the individual VCSEL's used for the free-space links, and the varying fan-out efficiencies of the different orders of the hologram.

Reduction of the noise in the prototype OMMH is possible if printed circuit boards with multiple signal layers for the drive circuits and the test–interface board are packaged and used. The current noise level is acceptable in relation to the signal-to-noise ratio determined by the BER value of 10^{-13} and the eye diagrams. By reducing the noise, a future prototype OMMH can achieve lower BER's, down to as low as 10^{-17}.

For the fiber links in the OMMH, misalignment is not an issue because we are using a packaged product—Optobus. However, the MT connectors connecting the fiber ribbon and the VCSEL's are a major concern as to the reliability of the mating connector. Poor mating leads to power loss and reduced data rates in our system. A couple of channels in the Optobus links were burned after 18 h of continuous testing, which suggests that burn-in times of the Optobus should be more thoroughly investigated. The fibers in the Optobus are directly coupled against the VCSEL's. This means that a lot of power is not efficiently coupled in the fiber, as could be achieved with microlenses. Heat removal and cooling effects, therefore, were a major issue with the use of the Optobus. Perhaps with the use of microlenses, which can be lithographically etched onto VCSEL's, more power could be coupled into the fibers and would result in higher bit rates for the links. Four Optobuses were used for the OMMH. This meant 40 links were available, of which we used 32. We chose the 32 best links by individual testing of each Optobus link. Links that were not used were found either to operate at less than 150 MHz, to be not functional, or to burn out after prolonged testing. Future versions of the Optobus are expected to deal with the issues of better mating between the fiber and VCSEL, higher bit rates (300 Mbytes/s), and better uniformity across the different channels.

For the free-space links in the OMMH, misalignment was a major issue. The very small diameter of the VCSEL's (8 μm) and the very short interconnect distances (10 mm) made alignment extremely difficult. Longitudinal misalignment was controlled by the observation of the spot sizes coming out of the hologram and the movement of the detector plane closer to the hologram so that as much light as possible was coupled into the detectors. Angular and lateral misalignment were a much more difficult proposition. We had to deal with lateral misalignment of the VCSEL's and the collimating lenslet array and at the same time of the lenslet array and the detectors. The hologram was space invariant, so lateral misalignment was not an issue. However, angular misalignment of the hologram with respect to the detector plane had to be precisely adjusted by the movement of the x–y positioners on the detector plane. Thus the availability of extremely accurate equipment is critical to aligning the system. Submillimeter accuracy was available, and it was sufficient to align the system. However, submicrometer accuracy would make alignment easier. Misalignment lead to other problems, such as cross talk at the receiver side, lower signal-to-noise ratios, and lower data rates.

The power budget was important, as is outlined in the power budget table (Table 1). The current power of the VCSEL's (1 mW) allowed us to operate at a data rate of 150 Mbytes/s. As future versions of the Optobus scale to data rates greater than 300 Mbytes/s, an increase in the power of the VCSEL's for the free-space links would be needed. An increase in power of the VCSEL's would yield a higher power level incident on the hologram, which would result in more power coupled to the detectors. The fabrication of VCSEL's is a rapidly emerging technology that promises improved power output in the near future. Such an improvement would definitely have a direct impact on the performance of the OMMH in terms of higher signal-to-noise ratios, higher data rates, and lower BER's.

The use of computer-generated holograms (CGH's) as the space-invariant diffractive element would provide higher diffraction efficiencies, as well as a balanced power distribution in the diffraction orders. For comparison with DCG holograms, a CGH was designed but not fabricated. The design values of the diffraction efficiencies of the CGH were 15% in the ± 1 orders and 20% for the zero order. This result suggests that perhaps CGH's are better for constructing fan-out elements, as, to an extent, one can achieve higher diffraction efficiencies, equalize the optical output power coupled into the different orders, and minimize the power coupled into unwanted orders.

However, the efficiency of the DCG hologram was sufficient to obtain the desired data rate of 150 Mbytes/s (limitation imposed by the Optobus) and as such the CGH was not necessary. As we scale to higher data rates however, the CGH will be used.

Finally, the size scaleability of the sample OMMH can be achieved by the expansion of the size of the mesh. For example, by the addition of 3-cubes in the row or column of the mesh shown above in Fig. 1. The mesh links are dictated by the Optobus. With the addition of a single Optobus module, 10 bidirectional mesh links can be implemented. A bidirectional OMMH would therefore scale up in size by the addition of VCSEL and receiver arrays along the rows and columns of planes L and R (Fig. 2) and by the addition of more Optobus modules to provide mesh connectivity without altering the existing configuration. In addition, up to a $(l, m, 9)$-OMMH can be implemented, as 8×8 VCSEL arrays are commercially available. Planes L and R would then contain 8×8 arrays of VCSEL's and detectors. A new hologram would then have to be fabricated to achieve the space-invariant connectivity for the 9-cube.

The maximum size of the hypercube modules is determined by the power limitations of the free-space-invariant optics. For large size hypercubes a significantly large number of fan-out beams would be required from the hologram. Thus power limitations should be considered. A way to improve the power limitations is to use beam-steering elements with higher diffraction efficiencies (i.e., CGH's).

7. Conclusions

We have experimentally demonstrated the first version of a scaleable optical interconnection network called OMMH using commercially available components. The OMMH network lends itself well to optical implementations for which high-bandwidth, point-to-point links are needed. Our initial results indicate that the required connectivity of the OMMH can be achieved by the use of optical interconnects transmitting at a data rate of 150 Mbytes/s with BER's of less than 10^{-13}. The OMMH was optimized to operate at 150 Mbytes/s because of the limitations imposed by Motorola's Optobus. A careful balance of the benefits of optics and electronics was critical to the success of the OMMH network. In addition, the difficulties in demonstrating our system included the varying fan-out efficiencies of the different orders of the hologram, misalignment sensitivity of the free-space links, low power (1 mW) in the individual VCSEL's, and noise. Future work in packaging driver circuits, higher diffraction efficiencies of the holograms (CGH's), increased laser power, and better alignment techniques would lead to an increase in the data rates, a reduction in the noise, and a lower BER for the OMMH network.

References

1. A. Louri and H. Sung, "A scaleable optical hypercube-based interconnection network for massively parallel computing," Appl. Opt. **33,** 7588–7598 (1994).
2. A. Louri and H. Sung, "An optical multi-mesh hypercube: a scaleable optical interconnection network for massively parallel computing," J. Lightwave Technol. **12,** 704–716 (1994).
3. A. Louri and H. Sung, "3D optical interconnects for high-speed interchip and interboard communications," Computer **27,** 27–37 (1994).
4. A. Louri and S. Furlonge, "Feasibility study of a scaleable optical interconnection network for massively parallel processing systems," Appl. Opt. **35,** 1296–1308 (1996).
5. T. S. Wailes and D. G. Meyer, "Multiple channel architecture: a new optical interconnection strategy for massively parallel computers," J. Lightwave Technol. **9,** 1702–1716 (1991).
6. G. R. Hill, "Wavelength domain optical network techniques," Proc. IEEE **77,** 121–132 (1989).
7. M. G. Hluchyj and M. J. Karol, "ShuffleNet: an application of generalized perfect shuffles to multihop lightwave networks," J. Lightwave Technol. **9,** 1386–1397 (1991).
8. R. J. Vetter and D. H. C. Du, "Distributed computing with high-speed optical networks," IEEE Trans. Comput. 8–18 (1993).
9. D. Bursky, "Parallel optical links move data at 3 Gbits/s," Electron. Design **42,** 79–82 (1994).
10. Motorola Corporation, "Optobus™," Tech. rep. BR1459/D, (Logic Integrated Circuits Division, Motorola, Tempe, Ariz., 1995).
11. A. Louri and H. Sung, "Efficient implementation methodology for three-dimensional space-invariant hypercube-based free-space optical interconnection networks," Appl. Opt. **32,** 7200–7209 (1993).
12. A. Louri and H. Sung, "Design methodology for three-dimensional space-invariant hypercube networks using graph bipartitioning," Opt. Lett. **18,** 2050–2052 (1993).
13. C. Fan, B. Mansoorian, D. A. VanBlerkom, M. W. Hansen, V. H. Ozguz, S. C. Esener, and G. C. Marsden, "Digital free-space optical interconnections: a comparison of transmitter technologies," Appl. Opt. **34,** 3103–3115 (1995).
14. S. Tang, R. T. Chen, L. Garrett, D. Gerold, and M. M. Li, "Design limitations of highly parallel free-space optical interconnects based on arrays of vertical cavity surface-emitting laser diodes, microlenses, and photodetectors," J. Lightwave Technol. **12,** 1971–1975 (1994).
15. J. Neff, "Optical interconnects based on two-dimensional VCSEL arrays," in *IEEE Proceedings of the First International Workshop on Massively Parallel Processing Using Optical Interconnections*, Cancun, Mexico, 26–27 April 1994, pp. 202–212.
16. N. C. Craft and A. Y. Feldblum, "Optical interconnects based on arrays of surface-emitting lasers and lenslets," Appl. Opt. **31,** 1735–1739 (1992).
17. G. Olbright, "VCSEL's could revolutionize optical communications," Photon. Spectra **29,** 98–101 (1995).
18. T. V. Muoi, "Receiver design for high-speed optical-fiber systems," J. Lightwave Technol. **2,** 243–267 (1984).
19. M. E. Landgraf, C. A. Eldering, S. T. Kowel, and P. F. Brinkley, "Optical interconnection techniques for hypercube," in *Optical Information Processing Systems and Architectures II*, B. Javidi, ed., Proc. SPIE **1347,** 580–591 (1990).
20. 8-Pin Monolithic Amplifier Evaluation Boards (Comlinear Corporation, Fort Collins, Colo., January 1993).
21. H. P. Herzig, P. Ehbets, D. Prongue, and R. Dandliker, "Fan-out elements recorded as volume holograms: optimized recording conditions," Appl. Opt. **31,** 5716–5723 (1992).
22. B. Robertson, M. R. Taghizadeh, J. Tutunen, and A. Vasara, "Fabrication of space-invariant fan-out components in dichromated gelatin," Appl. Opt. **29,** 1134–1141 (1990).
23. B. Robertson, M. R. Taghizadeh, E. Restall, and A. C. Walker, "Space-invariant holographic optical elements in dichromated gelatin," Appl. Opt. **30,** 2368–2375 (1991).