207

General Purpose Bi-Directional Optical Backplane Bus

Srikanth Natarajan, Chunhe Zhao, and Ray T. Chen

Microelectronics Research Center
Department of Electrical and Computer Engineering
University of Texas, Austin, Texas 78712-1084
Tel: 512-471-4349

Abstract

We report for the first time, a bi-directional optical backplane bus for a high performance system containing nine multi-chip module (MCM) boards. The backplane bus reported herein, employs arrays of multiplexed polymer-based waveguide holograms in conjunction with a waveguiding plate within which 16 substrate guided waves for 72 (8x9) cascaded fanouts, are generated. Data transfer of 1.2 Gbit/sec at 1.3 μm wavelength is demonstrated for a single bus line with 72 cascaded fanouts. Packaging-related issues such as transceiver size and misalignment are embarked upon, to provide a reliable system with a wide bandwidth coverage. The backplane bus demonstrated, is for general-purpose and therefore compatible with such IEEE standardized buses as VMEbus, Futurebus and FASTBUS, and can function as a backplane bus in existing computing environments and can significantly reduce the bottlenecks that accompany electrical interconnects.

I INTRODUCTION

The limitations of current computer backplane buses stem from their purely electronic interconnects. These limitations include wide interconnection time bandwidths, large clock skew and large RC and RLC time constants [1]. In massively parallel architectures, the bottleneck in performance is the limited bandwidth of current interconnection networks [2]. The processing elements are densely packed on a printed wire board. This makes it impossible to provide all the interconnection channels necessary for communication among processing elements and memory units, at the maximum bandwidth.

The distributed line RLC time constant is often too large for chip-to-chip interconnects and higher level hierarchies. These factors have already created serious bottlenecks in the most advanced electronic backplane interconnect prototypes. In general, when the time bandwidths provided by electrical interconnects are too wide, they are very difficult to manage. As clock cycle time and pulse widths shrink, the bandwidth needed to preserve the rising and falling edges of the signals increases. This makes using bulky, expensive, terminated coaxial interconnections a necessity. Bus line skew is the next most important performance limitation of conventional von Neumann computers, particularly on the backplane interconnect level. It slows the signal processing and occurs particularly when signals from different parts of a circuit arrive at a gate at slightly different times. Skews of up to tens of nanoseconds may occur. This input skew may cause a gate to generate an erroneous output unless an appropriate skew delay is inserted. Also, the fan-out capabilities for electrical interconnects are restricted due to electromagnetic interference [3]. As a result, it is very difficult to exploit the performance of ultra fast logic gates in a circuit with traditional electrical interconnects.

In this paper, we report for the first time, an optical bi-directional backplane bus which is compatible with an array of IEEE standardized backplane buses such as VMEbus, Futurebus, FASTBUS etc. Fig. 1 shows the optical equivalent of a single bi-directional electronic bus line. The drive current provided by each optoelectronic transceiver powers the corresponding laser diode, whose output is bi-directionally coupled through a multiplexed waveguide hologram, into the substrate which is optically transparent and therefore functions as the medium for the backplane bus array. The surface-normal fan-out is provided by a linear hologram array located between the backplane and the processor/memory boards. Each photodiode associated

0-8194-1747-5/95/$6.00

with a transceiver detects light from either hologram array, since the light travels in both directions in the waveguiding plate shown in Fig. 1.

II BI-DIRECTIONAL OPTICAL BACKPLANE BUS ARCHITECTURE

An optical backplane bus system equivalent to an array of IEEE standardized backplane buses is described. Fig. 2 is a schematic of how the optical bus can be used as a backplane in multi-processor, high performance optoelectronic computers. There needs to be bi-directional signal flow between the backplane and the processor/memory boards, where the multichip modules (MCMs) are located. With the design that we employ, the optical bus can serve the purpose of a bi-directional backplane. Multiplexed waveguide holograms are employed to facilitate two-way communication between boards that are connected to the backplane. The enlarged portion of Fig. 2 is a section of three bus lines which are bi-directionally interconnected through arrays of multiplexed waveguide holograms where bi-directionality of the interconnect is depicted.

The desired signal is coupled into the backplane bus through the TIR hologram, which is designed to provide a total internal reflected beam within the guiding plate. An array of fan-in/fan-out holograms are recorded along each line, as shown in Fig. 2. The actual fabrication process and experimental results are dealt with in the following section.

The functionality of the bi-directional backplane bus is well understood by its application in a transceiver system . By employing the holographic bi-directional bus reported in this paper, each cardboard can send and receive information to and from every other cardboard in the system. Note that, cascaded fan-out is employed in all the communications. The control and data signals are broadcasted to all cardboards associated with the high performance, multi-processor computer. The correct response from another processor or memory board, is based on address identification.

III EXPERIMENTAL RESULTS

The physical layer of the optical backplane bus is essentially a thin waveguiding plate with a set of 1-D holograms integrated on its surface. The substrate serves as the light-guiding medium. Dichromated gelatin (DCG) film is spin coated on the surface of the substrate. DCG is used as the recording material because of its large phase refractive index modulation capability, high resolution, low absorption and scattering [4]. For each bus line, two hologram arrays are recorded on the substrate afterward, to provide the required bi-directional surface-normal coupling. The recording wavelength used is 488 nm. The fabrication process [5] is optimized for all reconstruction wavelengths.

An array of multiplexed waveguide holograms is employed to provide a fully interconnected bi-directional optical bus. A total internal reflection (TIR) hologram is designed to couple the surface normal TEM_{00} laser beam into a substrate guided beam with a pre-designed bouncing angle, which is 45 degrees in our case. The second type of hologram couples an array of substrate guided beams into a 1-D array of surface-normal fan-out beams[6] with a specific coupling efficiency. The system demonstration is depicted in Fig. 3 where nine processor/memory modules are optoelectronically interconnected. Note that the input coupling hologram also functions as the output coupler when a reversely propagating guided wave is generated from other processors/memory boards. For example, the two holograms associated with the MCM board 5 (Fig. 3), couple the digital optical signal from itself to boards 4, 3, 2 and 1 through the first hologram and boards 6, 7, 8 and 9 through the second hologram. These two holograms are physically multiplexed on top of the waveguiding plate shown in Fig. 3. The same holograms will couple signals generated from boards 6, 7, 8 and 9 back to board 5 and from boards 4, 3, 2 and 1 back to board 5, when such communications are necessary.

The fabrication of the bi-directional bus is an extension of the steps mentioned above to record an 1-D recording in the opposite direction, except that the recording plate is rotated 180 degrees from the previous position, before the second recording. The mode dots in Fig. 4 correspond to the outputs of the transceivers shown in Fig. 3. Azimuthal symmetry of the mode profile is maintained [7], which significantly eases the

coupling from the laser to the backplane bus and then from the backplane bus to the detector and vice-versa. Unlike the conventional guided wave devices where a sizable portion of the surface area is needed to provide the optical signal routing, the substrate-guided-wave based board-to-board interconnect does not require any surface area for this purpose.

The main advantage of a bi-directional bus, as mentioned in the previous section, is its ability to provide a means for cascaded fan-out with signal flow in both directions. This concept is well illustrated by Fig. 4. This figure shows an array of photographs taken from the bi-directional bus, in the surface-normal direction. These photographs are taken with one MCM board generating signals and all other eight boards shown in Fig. 3 serve as receivers. Various input/output combinations are delineated. Note that the experiments are carried out with one hologram as the input coupler and all the other holograms as the output couplers. The bi-directionality is demonstrated by the fact that the same multiplexed hologram can be used as both the input and output coupler, thus aiding in bi-directional signal flow. Note that the input beam is not shown in Fig. 4. The array displayed in Fig. 4 shows four photohraphs where the second (Fig. 4a), third (Fig. 4b), fourth (Fig. 4c) and fifth (Fig. 4d) holograms are used as input couplers. Table1 indicates the input coupler and the corresponding output couplers for all the nine cases. Note that the nature of bi-directionality and cascaded fan-out is clearly indicated. Inter-MCM optoelectronic interconnects through backplane is clearly demonstrated. Unlike the optical backplane interconnects aimed at special purpose computers [8], the demonstration reported herein is suitable for general purpose multi-processor computer buses such as VMEbus, Futurebus and FASTBUS.

A square wave input can illustrate the device performance in the digital domain. One period of a square wave corresponds to two data bits. In effect a 500 MHz square wave corresponds to a 1 Gbit/sec digital signal. Fig. 5 is a photograph of the device output with the input being a 600 MHz square wave or a 1.2 GBit/sec digital signal. The input data is differential ECL and is obtained from a high speed pulse generator. A high speed sampling head is used in conjunction with an oscilloscope to display the output.

IV THEORETICAL CONSIDERATIONS OF THE SYSTEM

The change in intensity as a function of the fluctuation of the input coupling angle, from the input laser beam to the backplane bus plate, is measured. In our experiment, the angular fluctuation, i.e., $\delta\theta$ as shown in Fig. 6, is defined as the angular coverage for 1 dB off the maximum intensity. This value rather than the 3 dB points, is selected to provide a more reliable power budget for the system. A higher intensity results from a smaller deviation of the incident beam. The maximum deviation of the diffracted beam is also limited by the active area of the detector, which has a radius of 100 μm in our experiment.

Fig. 7 shows the theoretical and experimental plots quantifying the normalized intensity (efficiency) as a function of the angular misalignment. The physical parameters determining these curves are shown in the inset of Fig. 7. For the theoretical case, the overall efficiency is the product of the efficiencies of the input coupler and of the fan-out hologram. This results in a closer resemblance to the experimental data, where the incident light encounters two gratings before fan-out.

The wavelength variation ($\delta\lambda$) for the surface-normal coupling is related to the angular perturbation ($\delta\theta$) through the following relation [9] :

$$\delta\theta/\delta\lambda = K/[4\pi n \sin(\phi-\theta)] \qquad (1)$$

where λ is the optical wavelength, θ is the incident angle and ϕ is the grating slant angle [5], and n the emulsion index. And K is the magnitude of the grating vector. In our experiment n= 1.5, $\phi = 45^0$ and $\theta = 0^0$.

The schematic showing these physical parameters is illustrated in Fig. 6. Since all the quantities in Eq. 1 are known, the wavelength shift $\delta\lambda$ can be easily derived. The frequency shift $\delta\upsilon$ corresponding to this wavelength shift can be calculated according to the following first order approximation:

$$\delta\upsilon = (c/\lambda^2)\delta\lambda \qquad (2)$$

The wavelength coverage determined by transforming the angular fluctuation of 0.5 degree, to wavelength fluctuation, is 42 nm, which is equivalent to 7.5 THz (10^{12}/sec) base bandwidth.

In high-speed digital optical systems, the active area of a photodetector is inversely proportional to the modulation speed of the signal beam. In our experiment, the radius of the active area is 100 microns. The diffracted beam may traverse upto a radius of the active area of the detector. The beam strikes the center of the detector and then moves along the radius as the input angle θ changes. To determine the relationship between the diffracted beam deviation and the radius of the active area of the detector, as indicated in Fig. 6, we further derive the following equation:

$$\delta r = 2DN(\tan\phi_1 - \tan\phi_2) \qquad (3)$$

where δr represents the radius of the active area of the detector, which is 100 μm in our case. D= 3000 μm (substrate thickness), N= 10 (10th fan-out) and ϕ_1 is the input diffraction angle, which is 45 degrees in our case. From Eq. 3, ϕ_2 can be calculated and $\phi_1 - \phi_2$, shown in Fig. 6 is the maximum deviation for the detector with an active area radius of 100 microns ($\phi_1 - \phi_2 = \delta\phi$). Fig. 6 is a schematic where the limitation of the detector area size on the maximum deviation of the diffracted beam is clearly indicated. The input angular misalignment will cause a shift in the diffraction angle. Eq. 4 relates the input angle deviation to the diffraction beam deviation. The shift in diffraction angle and thus the wavelength deviation can be calculated based on Eq. 1 and Eq. 4, which is

$$\delta\phi = -[\cos(\theta)/n\cos(\phi)]\delta\theta \qquad (4)$$

The bandwidth coverage due to the angular shift of TIR beams is determined to be 5 nm. This clearly indicates that the detector size limits the bandwidth of the system that incorporates the bus, though the bus itself has a bandwidth, two orders of magnitude higher. The frequency shift associated with this bandwidth is 0.89 THz.

As far as the power budget is concerned, the performance of a multiplexed system is always limited by the output channel with the minimum power. For the bi-directional optical bus we report in this paper, the power of the output channels is determined by the diffraction efficiencies of the two sets of holograms. Due to the bi-directionality of the optical bus, it is impossible to get uniform intensity fan-outs for all the cases where the modulated optical signals are incident from different channels. So, the best way to overcome this least-power-limitation problem is to try to balance the output power from different channels. This can be achieved by optimizing the diffraction efficiencies of the holographic gratings. The details are presented in another publication.

V BUS PROTOCOLS

If an optical bus has to serve as a backplane in existing computing environments, it must be able to provide the devices that communicate via the bus, a mechanism to base their timing [10]. This kind of timing requirements are satisfied by standard bus data transfer protocols. In a multi-processor system, board-to-board data transfer

between a sender and a receiver must be co-ordinated. The data transfer protocols provide this timing co-ordination between devices that are sending and receiving data.

Synchronous, uncompelled asynchronous and compelled asynchronous[10] are the three standard type of data transfer protocols used in electronic backplane buses. In synchronous data transfer, a synchronous bus provides a bus clock line whose rising and/or falling edges are signals for data transfers to take place. In uncompelled asynchronous data transfer, the bus timing is controlled by the data sender rather than by the bus itself. In case of the compelled asynchronous data transfer, compelled protocols specify an exchange of signals (handshake) between sender and receiver upon each transmission. Our demonstration shown in the previous section provides the hardware needed to optoelectronically interconnect these backplane systems.

The main advantage of the bi-directional optical bus that we have designed and demonstrated, is that equivalent data transfer protocols can be designed for the optical backplane bus. Bi-directional holograms can be designed along the bus lines to provide for clock distribution and data transfer. A uni-directional high-speed optical backplane bus with a modulation speed of 0.89 THz has been demonstrated [13]. The bi-directional bus can also be designed for such high speeds. This implies that compelled and uncompelled asynchronous high-speed data transfer can be efficiently done through the bi-directional optical backplane bus.

The bi-directional optical backplane bus can be made transparent to all existing bus systems, if equivalent data transfer protocols (as already discussed) are employed. The distinct advantage of the optical equivalent of the VMEbus will be a drastic reduction in round-trip delay time through the bus, due to the absence of capactive loading effects. Another favorable aspect of the optical bus would be the elimination of any settling time for the bus lines, when they change from one steady state to a new steady state [10].

VI FACTORS INFLUENCING THE COMMERCIAL VIABILITY OF THE BI-DIRECTIONAL OPTICAL BACKPLANE BUS

Fig. 7 shows that the angular full-width at half-maximum diffraction efficiency is $\pm 2^\circ$, which translates to a large angular tolerance of $\pm 1^\circ$ before the diffraction efficiency drops significantly. The angular misalignment of the input beam θ, will cause a spatial shift of the fan-out beam due to the shift of the diffraction angle ϕ. Based on phase matching conditions this spatial shift can be expressed as [13]:

$$\delta L = 2D[\tan(\phi)-\tan(\phi \pm \delta\phi) \qquad (5)$$

where D is the substrate thickness. For our device D=3000μm, $\theta = 0^\circ$, and $\phi = 45^\circ$. If we consider angular misalignments of 1° and 0.1°, they correspond to ~ 200 μm and ~ 2 μm spatial shifts respectively. The above equation transforms a 3-D spatial and angular alignment problem into a 2-D planar one. Existing 2-D planar alignment techniques employed in fabricating silicon VLSI circuits, can be used to integrate the bi-directional optical backplane bus to multichip modules (MCMs) [13]. This implies that it is indeed feasible to realize a commercially exploitable device.

Crosstalk analysis is critical in devices where channels are designed by multiple exposures on the same emulsion area. The recording angle and the diffraction angle are different for each exposure. These parameters

are changed to fabricate holograms with different grating periods. A fine example for such a device is a wavelength division demultiplexer (WDDM). Here the idea is to disperse signals at selected wavelengths [14]. In this case there can exist a state where there is optical signal in say, channel A while channel B is supposed to have no optical signal. But due to crosstalk, from channel A, some optical signal maybe detected in channel B.

In our device, even though the emulsion area is exposed twice, the recording parameters remain unchanged. So the periodicity of the gratings is the same. Crosstalk is not critical in our case since it is one optical signal that is guided in the substrate after it is coupled in by the input coupler and it is the same optical signal that is coupled out by the other diffraction holograms. In fact cross-coupling is an essential part of device operation. The system however can have a degradation of the signal-to-noise ratio due to crosstalk between the optoelectronic components involved, but this is a limitation associated with any communication system that incorporates optoelectronic devices at the source and detector end of each transceiver.

There are other factors that contribute to losses in the system that incorporates the optical backplane bus. There is a significant coupling loss associated with the optical bus. Only 20% of the input signal is coupled into the substrate. Reflection from the surface and transmission through the substrate contribute to this reduced coupling efficiency. However the coupled signal is sufficient for the efficient operation of any transceiver system, where the optical bi-directional backplane bus acts as the interconnect medium. There is also a 0.1 dB/cm propagation loss in the device, which translates to 0.6 dB in our device which is 6 cm long. The propagation loss is well within limits that are acceptable for existing commercial applications. An electronic bus line is in effect a transmission line. As a result of this, there exists a finite settling time due to the presence of mismatched impedance. In case of an optical bus line, reflected power can be made to be zero by use of index-matched anti-reflection coatings at the ends. With this type of optical isolation, settling-time effects can be eliminated.

Packaging is a very critical part of fabricating any commercial transceiver system. Since a transceiver system includes electronic and optical components, the primary criterion that has to be addressed is the alignment between the various components. The source, detector and the holographic optical elements must be perfectly aligned to one another [15]. In effect, if alignment between components is done with care, packaging limitations can be considerably reduced.

VII CONCLUSIONS

We have discussed the advantages and applications of a bi-directional optical backplane bus using polymer-based multiplexed waveguide holograms in conjunction with a thin waveguiding plate. The architecture of this type of backplane bus has also been explained in considerable detail. A system with nine MCM boards is demonstrated with cascaded fan-out capability and fully equivalent to an electronic backplane bus. Data transfer rate of 1.2 Gbit/sec at 1.3 μm wavelength is further demonstrated with a single bus line. Packaging-

related issues, such as misalignment of transceiver arrays and detector size limitation are further addressed, in order to provide a reliable system.

The polymer-based bi-directional optical bus reported herein, is transparent to all the higher layers of electronic bus systems and can be applied to all buses of interest (VME, Futurebus, CAMAC, FASTBUS, and other high performance backplane buses) including the most advanced futurebus (IEEE896.1, up to 250 Mbit/sec data highways) as long as the appropriate protocols governing the rules of data transaction are provided.

REFERENCES

[1] P. Sweazey, "Limits of performance of backplane buses," in Digital Bus Handbook, edited by J. De Giacomo, McGraw-Hill, New York, 1990.

[2] A. Husain, T. Lane, C. Sullivan, J. Bristow, and A. Guha, "Optical backplanes for massively parallel processors and demonstration in the connection machine," presented at GOMAC, Las Vegas, NV, Nov. 1990.

[3] Ray. T. Chen, Suning Tang, Maggie M. Li, David Gerald, and Srikanth Natarajan, " 1-to-12 surface normal three-dimensional optical interconnects," *Applied Physics Letters.*, vol. 63, n0. 20, pp. 1883-1885, 1993.

[4] Hans Peter Herzig and Rene Dandliker, "Diffractive components: holographic optical elements," in Perspectives for Parallel Optical Interconnects, edited by Ph. Lalanne and P. Chavel, Springer-Verlag, New York 1993.

[5] Ray. T. Chen, Maggie Lee, Srikanth Natarajan, Chuan Lin, Z. Z. Ho, and Dan Robinson, "Single-mode Nd^{3+}-doped graded-index polymer waveguide amplifier," *IEEE Photonics Technology Letters.*, vol. 5, no. 11, 1993.

[6] Suning Tang, and Ray. T. Chen, "Surface normal 1-to-27 optical fan-out using substrate guided waves in conjunction with two dimensional waveguide hologram array," *Integrated Photonics Research.*, vol. 3, pp. 277-279, 1994.

[7] Ray. T. Chen, Hey Lu, Daniel Robinson, Michael Wang, Gajendra Savant, and Tomasz Jannson, "Guided-wave planar optical interconnects using highly multiplexed polymer waveguide holograms," *Journal of Lightwave Technology.*, vol. 10, no. 7, pp. 888-897, 1992.

[8] J. P. G. Bristow, "Recent progress in optical backplane," presented at OE/LASE '94, Los Angeles, California 22-29 January 1994.

[9] Herwig Kogelnik, "Coupled wave theory for thick hologram gratings," *The Bell System Technical Journal.*, vol. 58, no. 9, pp. 2909-2947, 1969.

[10] Ray. T. Chen, "Architecture and building blocks for VME optical backplane bus," *Proc.* SPIE., vol. 1849, pp. 196-208, 1993.

[11] IEEE Futurebus^{+} P869.1: Logical Layer Specifications, Published by IEEE, New York, 1990.

[12] Srikanth Natarajan, Ray T. Chen, Suning Tang, and Robert A. Mayer, "High speed optical backplane bus with modulation and demodulation capabilities," *Proc.* SPIE, vol. 2153, pp. 344-351, 1994.

[13] Suning Tang, R. T. Chen, Dave Gerold, M. M. Li, Srikanth Natarajan, Jielun Lin, N. Chellappan and M. Peskin, "Design considerations for high packaging density optical bus array," *Proc. of* SPIE, vol. 2153, pp. 227-235, 1994.

[14] Ray T. Chen and H. Lu, "Polymer-based 12-channel single-mode wavelength division multiplexer on GaAs substrates," *Proc.* SPIE, vol. 179, pp 44, 1992.

[15] Susant K. Patra, Jian Ma, Volkan H. Ozguz and Sing H. Lee, "Alignment issues in packaging for free-space optical interconnects," *Optical Engineering.*, vol. 33, no. 5, pp. 1561-1570, May 1994.

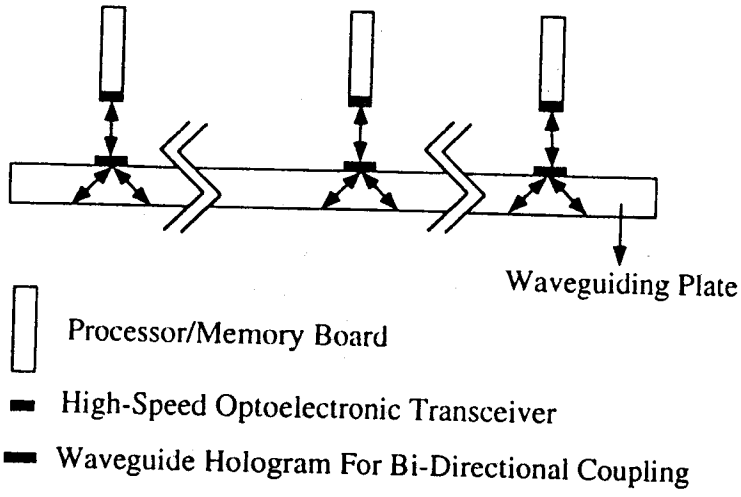

Figure. 1. Optical equivalent of a section of a single electronic bus line.

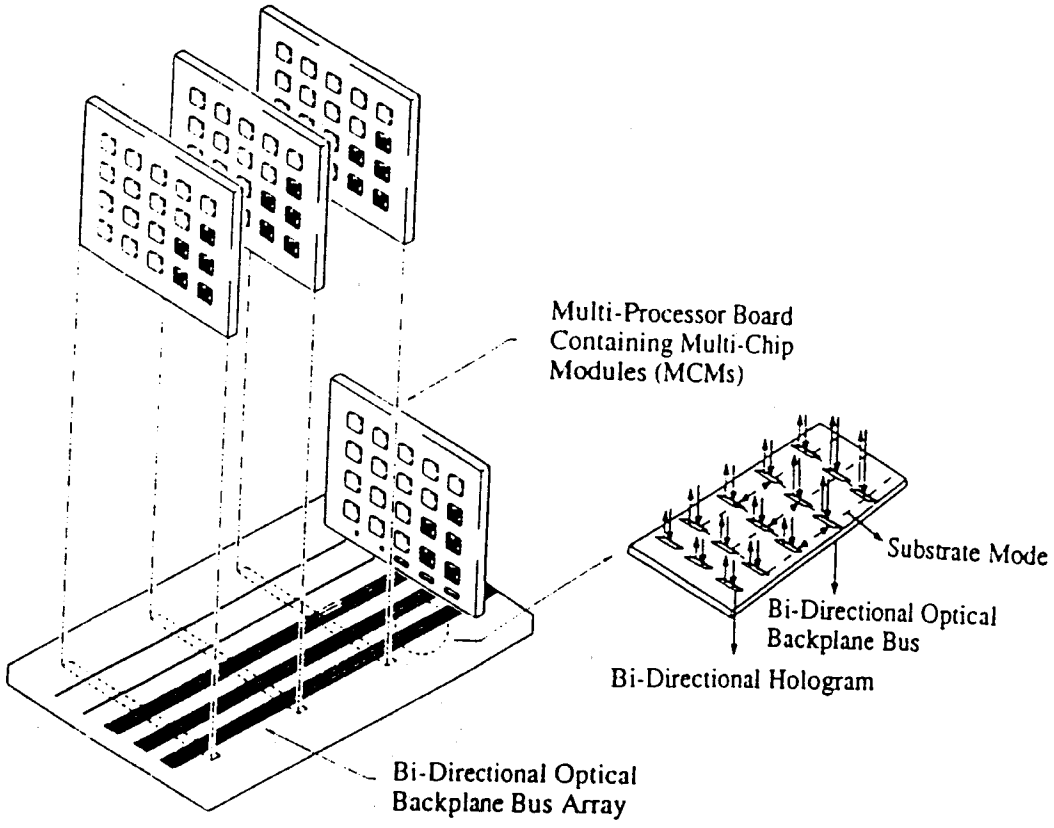

Figure. 2. Schematic of the bi-directional optical backplane bus in a multiprocessor system.

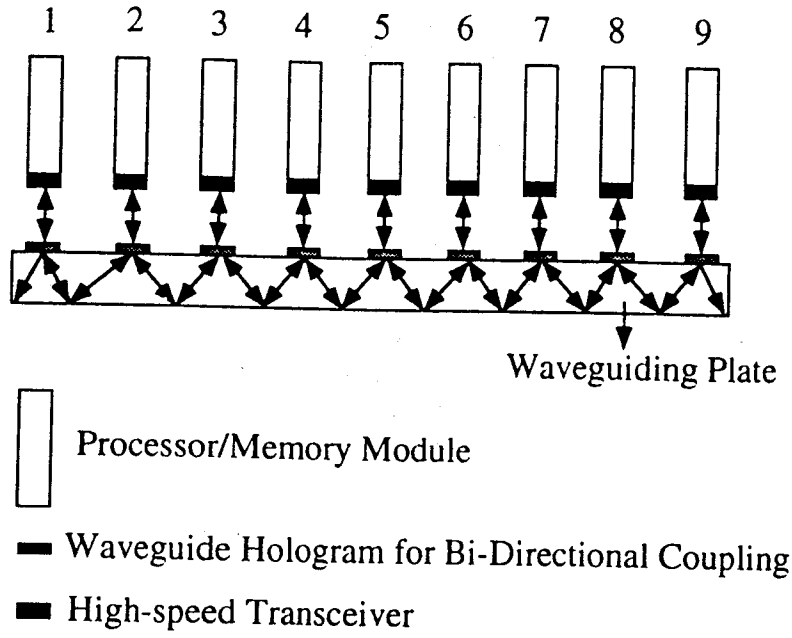

Figure. 3. Schematic depicting the functionality of bi-directional optical backplane bus in a multi transceiver system, with the multiplexed holograms in conjunction with a waveguiding plate.

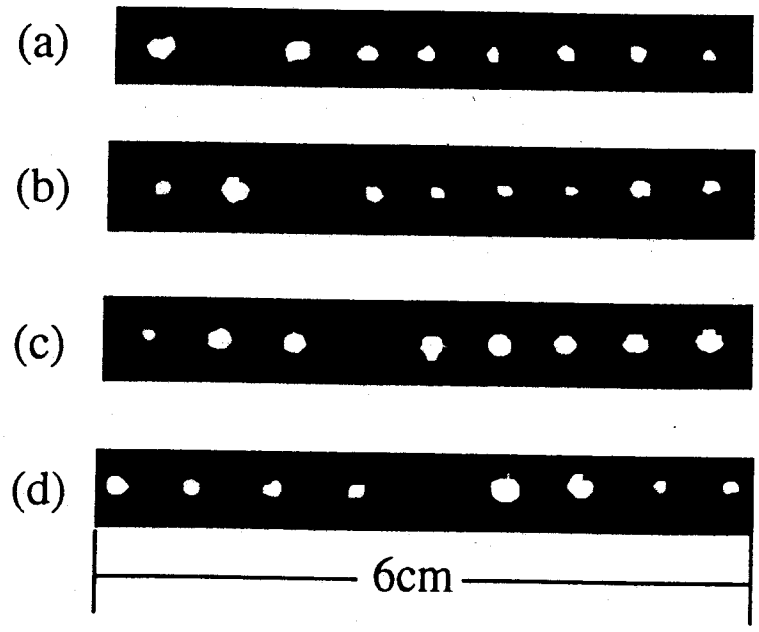

Figure. 4. An array of 4 photographs with the 2nd (4a), 3rd (4b), 4th (4c) and 5th (4d) holograms functioning as the input couplers, in the 1-to-8 bi-directional cascaded fan-out optical backplane bus.

Input MCM Board	Output MCM Board
1	2, 3, 4, 5, 6, 7, 8, 9
2	1, 3, 4, 5, 6, 7, 8, 9
3	1, 2, 4, 5, 6, 7, 8, 9
4	1, 2, 3, 5, 6, 7, 8, 9
5	1, 2, 3, 4, 6, 7, 8, 9
6	1, 2, 3, 4, 5, 7, 8, 9
7	1, 2, 3, 4, 5, 6, 8, 9
8	1, 2, 3, 4, 5, 6, 7, 9
9	1, 2, 3, 4, 5, 6, 7, 8

Table. 1. Representation of the 9 possible arrangements where one module acts as the input module, with remaining 8 modules functioning as the output modules.

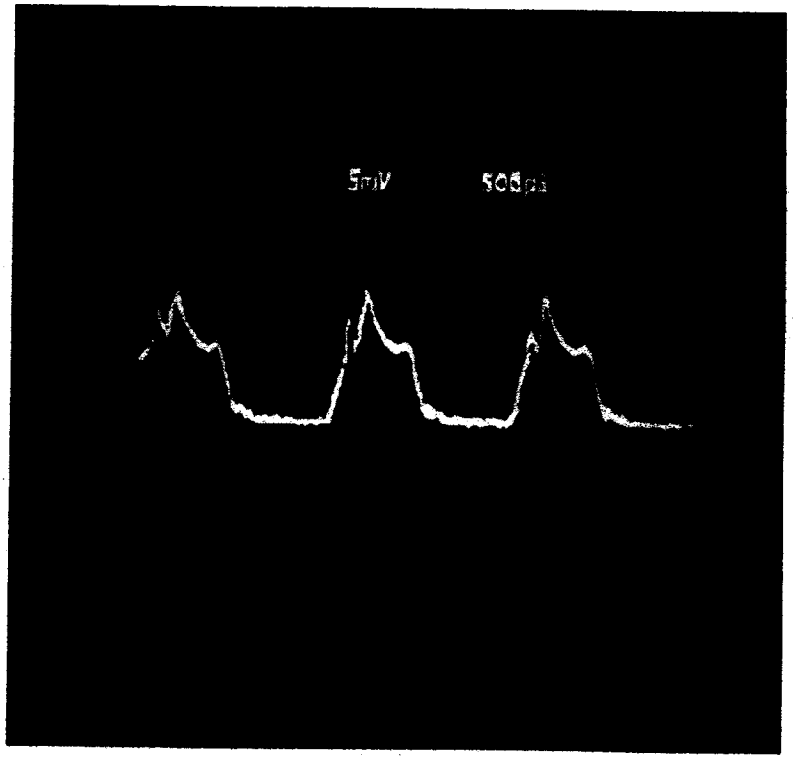

Figure. 5. Receiver output from the transceiver system incorporating the bi-directional optical backplane bus for a 1.2 Gbit/sec input signal.

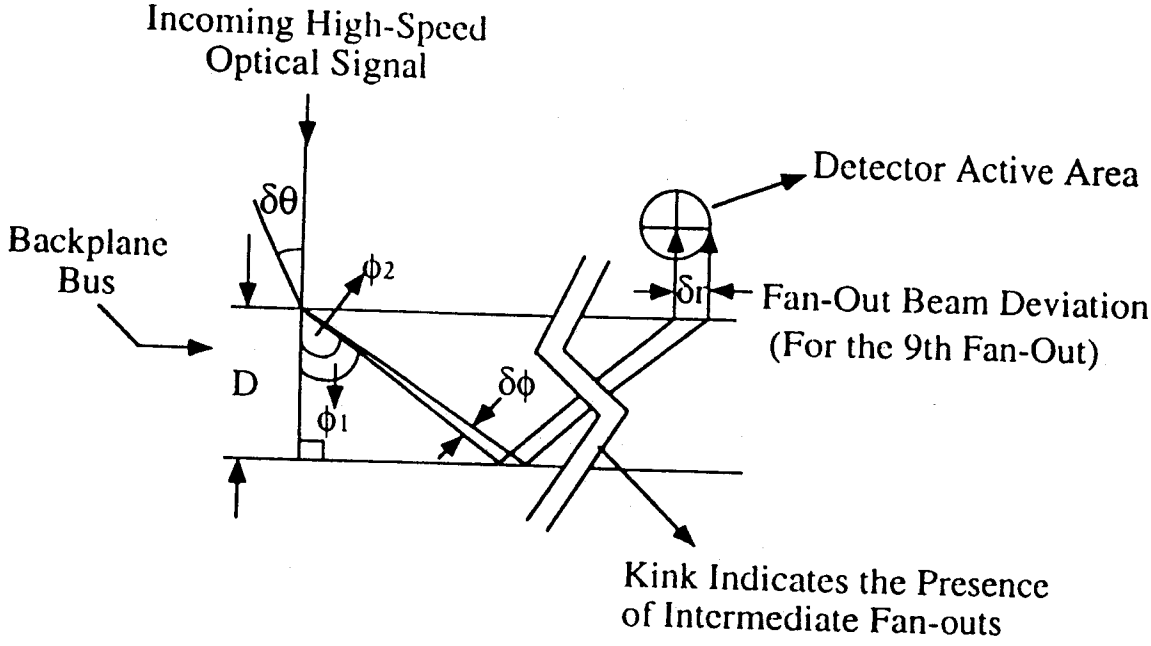

$$\delta r = 2DN(\tan\phi_1 - \tan\phi_2)$$

N= Number of Fan-Outs

Figure. 6. Schematic of the related parameters for surface-normal and intra-guiding-plate optical interconnects.

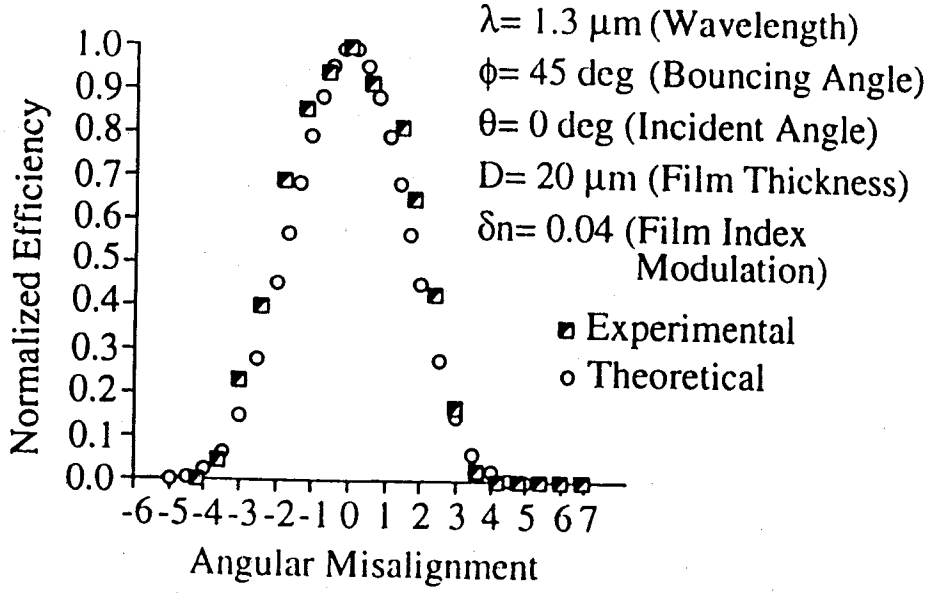

$$\delta\theta/\delta\lambda = K/[4\pi n\sin(\phi-\theta)]$$

K= Grating Vector Magnitude
n= 1.5 (Emulsion Index)

Figure. 7. Experimental and theoretical plots showing the relationship between angular misalignment and normalized intensity (efficiency) for the 1-to-8 optical backplane bus operating at 1300 nm.