Bi-Directional Optical Backplane Bus for General Purpose Multi-Processor Board-to-Board Optoelectronic Interconnects

Srikanth Natarajan, Chunhe Zhao, and Ray. T. Chen

Abstract— **We report for the first time a bidirectional optical backplane bus for a high performance system containing nine multi-chip module (MCM) boards, operating at 632.8 and 1300 nm. The backplane bus reported here employs arrays of multiplexed polymer-based waveguide holograms in conjunction with a waveguiding plate, within which 16 substrate guided waves for 72 (8 × 9) cascaded fanouts, are generated. Data transfer of 1.2 Gbt/s 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. Theoretical treatment to minimize intensity fluctuations among the nine modules in both directions is further presented and an optimum design rule is provided. 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.**

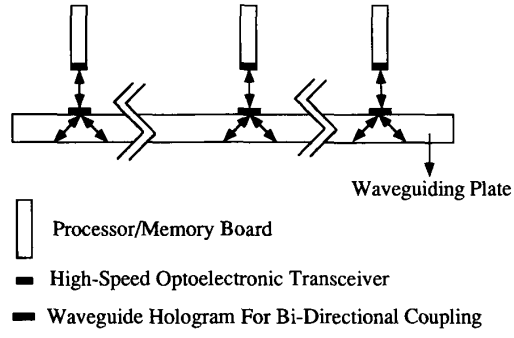

Fig. 1. Optical equivalent of a section of a single bidirectional electronic bus line.

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, such as IBM's backplane, in which the bottleneck occurs at 150 Mbt/s. For present high-speed buses, the electronic limitations are even more pronounced. For example, the VMEbus serial bus (VSB) transfers data at 3.2 Mbt/s, and its speed degrades to 363 kbts/s when the two communication points are separated by 25 m [3].

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.

Several effects preclude the use of logic in a pulsed mode. The first effect is skew, just described. The second is slow pulse rise and fall times due to RC and LC charging effects. A third effect is transmission-line and reflections due to the impossibility of exact line termination. The accepted approach is to wait for the inputs to settle before utilizing the output of a gate. Presently, for example, the RC time constant is already slower than the time it takes for a transistor to switch. Also, the fan-out capabilities for electrical interconnects are restricted due to electromagnetic interference [4]. As a result, it is very difficult to exploit the performance of ultra fast logic gates in a circuit with traditional electrical interconnects.

The difficulties associated with this RC or RLC-dominated settling time are not solved by VLSI. Indeed, as the length of a wire shrinks by a factor of S and the cross-sectional area of the wire is reduced by a factor of S^2, the capacitance of the wire decreases by a factor of S while the resistance increases by a factor of S. Therefore, the RC time constant and thus the input charging time remain the same, independent of scaling.

Manuscript received April 15, 1994; revised September 30, 1994.

The authors are with the Microelectronics Research Center, Department of Electrical and Computer Engineering, University of Texas, Austin, TX 78712-1084 USA.

IEEE Log Number 9411566.

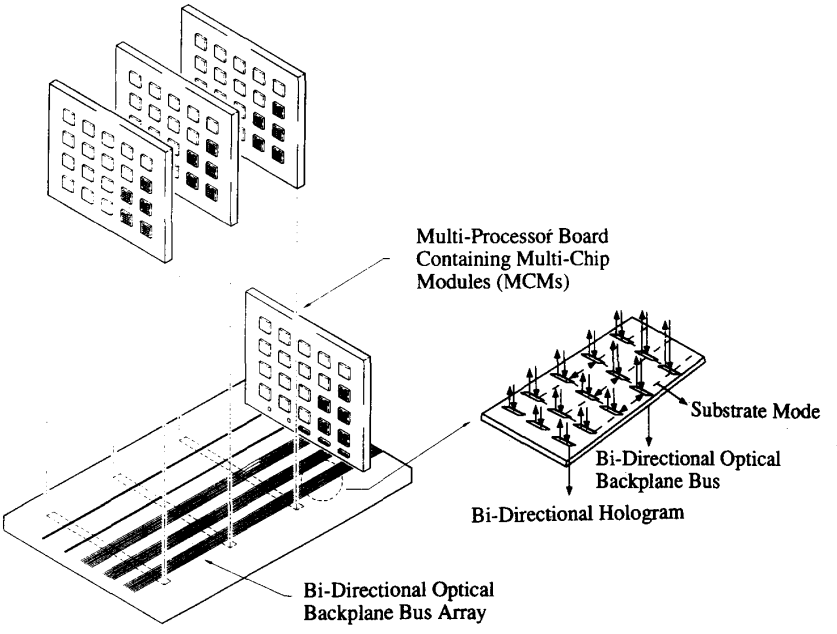

Fig. 2. Schematic of the optical bidirectional backplane bus in a multiprocessor system.

Given the RC parameters of VLSI, the signals will propagate at approximately 0.5–20% of the speed of light.

In this paper, we report for the first time an optical bidirectional backplane bus that 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 bidirectional electronic bus line. The drive current provided by each optoelectronic transceiver powers the corresponding laser diode whose output is bidirectionally 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 with a transceiver detects light from either hologram array, since the light travels in both directions in the waveguiding plate shown in Fig. 1.

The architecture of the demonstrated bidirectional optical backplane bus is described in Section II where the bidirectionality of the optical bus based on substrate guided waves along the backplane, is delineated. System demonstration containing nine processor/memory boards is described in Section III. Data rate at 1.2 Gbt/s for a single bus line is demonstrated at 1.3-μm wavelength. The theoretical embarkation upon the base bandwidth coverage of the backplane bus as a function of transceiver misalignment is reported in Section IV. The optimization of energy distribution among the nine surface-normal fan-outs is also presented. Other system related issues such as bus protocols and various IEEE standardized buses such as VMEbus, Futurebus, and FASTBUS are further discussed in Section V. finally, the concluding remarks are summarized in Section VI.

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 bidirectional signal flow between the backplane and the processor/memory boards, where the multichip modules (MCM's) are located. With the design that we employ, the optical bus can serve the purpose of a bidirectional 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 that are bidirectionally interconnected through arrays of multiplexed waveguide holograms where bidirectionality 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 bidirectional backplane bus is well understood by its application in a transceiver system. By employing the holographic bidirectional 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

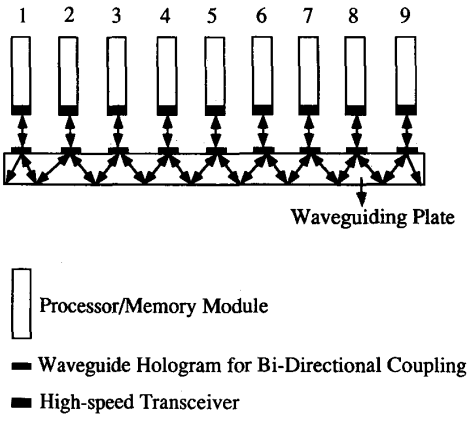

Fig. 3. Schematic depicting the functionality of the bidirectional optical backplane bus in a multi-transceiver system containing the demonstrated multiplexed waveguide hologram arrays in conjunction with a waveguiding plate.

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 [5]. For each bus line, two hologram arrays are recorded on the substrate afterward, to provide the required bidirectional surface-normal coupling. The recording wavelength used is 488 nm. The fabrication process [6] is optimized to provide two reconstruction wavelengths at 632.8 and 1300 nm.

An array of multiplexed waveguide holograms is employed to provide a fully interconnected bidirectional 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 [7] 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 4–1 through the first hologram and boards 6–9 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 6–9 back to board 5 and from boards 4–2 and 1 back to board 5, when such communications are necessary.

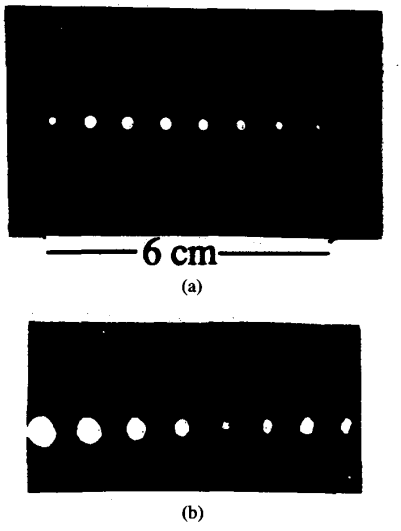

Fig. 4. (a) Photograph of the 1-to-8 bidirectional optical backplane bus operating at 632.8 nm and (b) photograph of the 1-to-8 optical backplane bus operating at 1300 nm.

The fabrication of the bidirectional bus is an extension of the steps mentioned above to record an 1-D hologram array. The recording process is done for one direction first and the same process is repeated for recording in the opposite direction, except that the recording plate is rotated 180 degrees from the previous position, before the second recording. Two photographs of the 1-to-8 bidirectional bus operating at 632.8 and 1300 nm are shown in Fig. 4(a) and (b). The hologram associated with the first MCM board, is the input coupler in this case. The mode dots in the picture are the far field patterns (without an imaging lens) taken in the surface-normal direction, where the MCM boards associated with the eight output modules are removed to facilitate the formation of the far field image. Note that the input beam from board 1 is blocked and therefore not seen in Fig. 4.

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 [8], 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 bidirectional 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. 5. This figure shows an array of photographs taken from the bidirectional 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 bidirectionality is demonstrated by the fact that the same multiplexed hologram

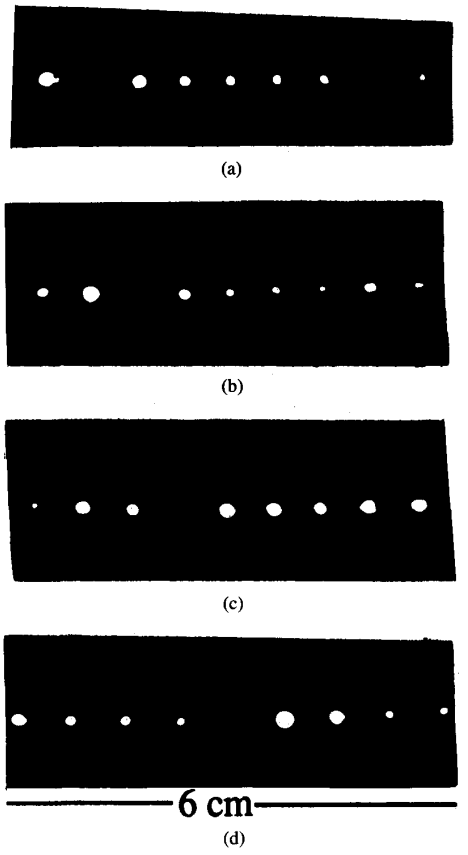

Fig. 5. An array of four photographs with the (a) second, (b) third, (c) fourth and (d) fifth holograms functioning as the input couplers in the 1-to-8 bidirectional cascaded fan-out optical backplane bus.

TABLE I
REPRESENTATION OF THE NINE POSSIBLE ARRANGEMENTS WHERE ONE MODULE ACTS AS THE INPUT MODULE WITH THE REMAINING EIGHT MODULES FUNCTIONING AS THE OUTPUT MODULES

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

can be used as both the input and output coupler, thus aiding in bidirectional signal flow. Note that the the input beam is not shown in Fig. 5. The array displayed in Fig. 5 shows four photographs where the second (Fig. 5(a)), third (Fig. 5(b)), fourth (Fig. 5(c)) and fifth (Fig. 5(d)) holograms are used as input couplers. Tables I and II, respectively, indicate the input coupler and the corresponding output couplers for all the nine cases and the measured diffraction efficiencies for all the possible input/output combinations. Note that the nature of bidirectionality 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 [9], the demonstration reported herein is suitable for general purpose multi-processor computer buses such as VMEbus, Futurebus, and FASTBUS.

TABLE II
MEASURED FAN-OUT DISTRIBUTION OF ALL THE POSSIBLE INPUT/OUTPUT COMBINATIONS OF THE 1-TO-8 BI-DIRECTIONAL BUS

Measured Output Energy (Arbitrary Unit)

Output \ Input	1	2	3	4	5	6	7	8	9
1	0	0.92	0.40	0.16	0.15	0.11	0.07	0.09	0.02
2	0.92	0	0.92	0.40	0.16	0.15	0.11	0.07	0.09
3	0.40	0.92	0	0.92	0.40	0.16	0.15	0.11	0.07
4	0.16	0.40	0.92	0	0.92	0.40	0.16	0.15	0.11
5	0.15	0.16	0.40	0.92	0	0.92	0.40	0.16	0.15
6	0.11	0.15	0.16	0.40	0.92	0	0.92	0.40	0.16
7	0.07	0.11	0.15	0.16	0.40	0.92	0	0.92	0.40
8	0.09	0.07	0.11	0.15	0.16	0.40	0.92	0	0.92
9	0.02	0.09	0.07	0.11	0.15	0.16	0.40	0.92	0

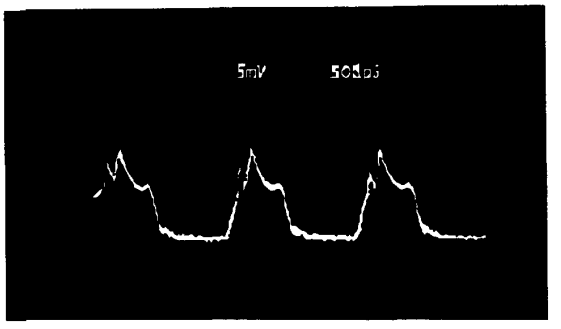

Fig. 6. Receiver output from the transceiver system incorporating the bidirectional optical backplane bus for a 1.2-Gbt/s input signal.

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-Gbt/s digital signal. Fig. 6 is a photograph of the device output with the input being a 600-MHz square wave or a 1.2-Gbt/s 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 LIMIT OF THE SYSTEM

The change in intensity as a function of the variation 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. 7, is defined as the angular coverage for 1 dB off the maximum intensity. This

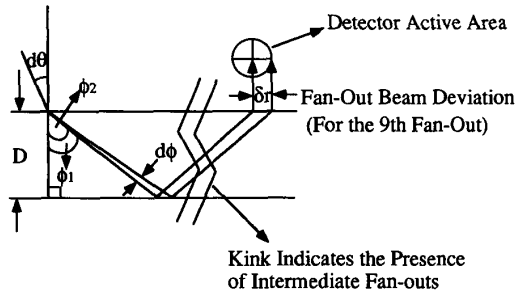

Fig. 7. Schematic of the related parameters for surface-normal and intra-guiding-plate optical interconnects.

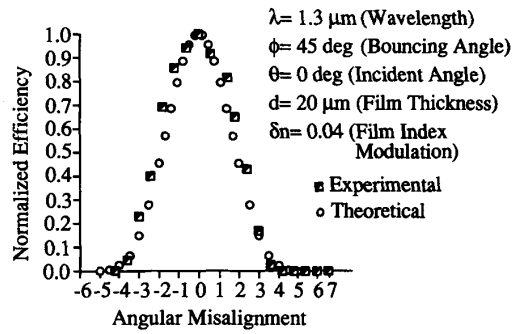

Fig. 8. Experimental and theoretical plots showing the relationship between angular misalignment and normalized intensity (efficiency) for the 1-to-8 optical backplane bus operating at 1300 nm.

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. 8 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. 8. 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 [10]:

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

where λ is the optical wavelength, θ is the incident angle, ϕ is the grating slant angle [6], and n is the emulsion index. K is the magnitude of the grating vector. In our experiment $n = 1.5$, $\phi = 45°$, and $\theta = 0°$. The schematic showing these physical parameters is illustrated in Fig. 7. Since all the quantities in (1) are known, the wavelength shift $\delta\lambda$ can be easily derived. The frequency shift $\delta\nu$ corresponding to this wavelength shift can be calculated according to the following first order approximation:

$$\delta\nu = (c/\lambda^2)\delta\lambda. \tag{2}$$

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}/s) 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 μm. The diffracted beam may traverse up to 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. 7, we further derive the following equation:

$$\delta r = 2DN(\tan\phi_1 - \tan\phi_2) \tag{3}$$

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 (3), ϕ_2 can be calculated and $\phi_1 - \phi_2$, shown in Fig. 7 is the maximum deviation for the detector with an active area radius of 100 μm ($\phi_1 - \phi_2 = \delta\phi$). Fig. 7 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. Equation 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 (1) and (4), which is

$$\delta\phi = -[\cos(\theta)/n\cos(\phi)]\delta\theta]. \tag{4}$$

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 bidirectional 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 bidirectionality 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.

A schematic demonstrating the diffraction mechanisms of the bidirectional optical bus is given in Fig. 9. After having recorded the two sets of hologram arrays, two grating vectors,

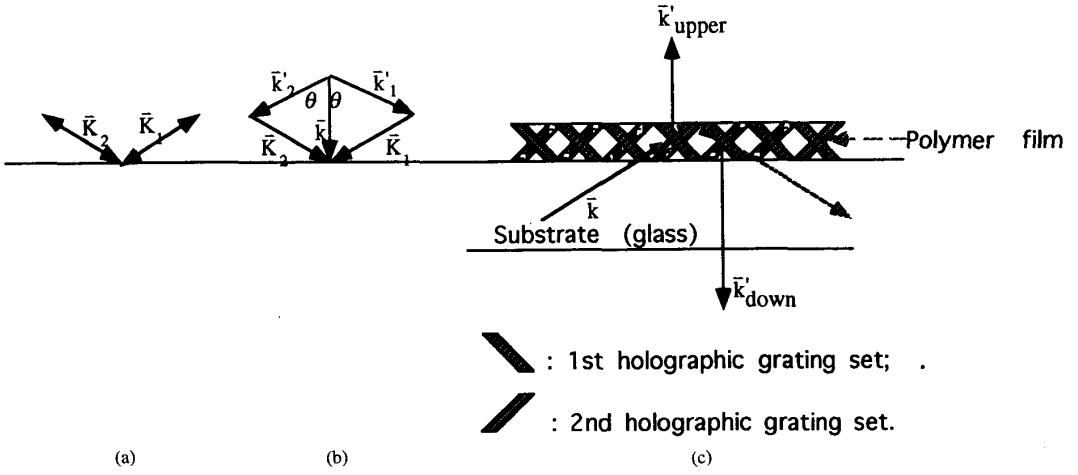

Fig. 9. Diffraction mechanism of light in by two sets of hologram gratings recorded recorded on a glass substrate. (a) Two grating vectors recorded; (b) Light diffraction geometry of the gratings; (c) Light transmission, diffraction, and reflection inside the bidirectional backplane bus.

$\vec{K}_1$ and $\vec{K}_2$, are formed in the film (Fig. 9(a)). A surface-normal light beam with a wave vector $\vec{k}$ is incident onto the holographic gratings. If the Bragg diffraction conditions are met, two diffracted light beams will be generated by the two holographic gratings and converted into substrate modes in two opposite directions (Fig 9(b)). During the subsequent transmission processes, whenever the light beams hit a grating and the Bragg condition is met, part of the light beams will be diffracted surface-normally out of the backplane, while the remaining beams continue to propagate in the waveguiding plate (Fig. 9(c)).

In relation to our experimental device, we consider the following bidirectional optical bus with 9 boards on one side of the substrate (as in Fig. 3). In our analysis we will assume that the diffraction efficiencies of the first set of hologram are $\eta_1, \eta_2 \cdots$, from left to right, respectively. Due to symmetric requirement for optimization, the diffraction efficiencies for the second hologram array η_1, η_2, $\cdots\cdots$ and η_9, from right to left. To provide the optimized power budget, we have to impose the follwing criteria: $\eta_1 = 1$and $\eta_9 = 0$, i.e., there is only one hologram at the 1st channel and the 9th channel. If we denote P_{ij} to be the output power at the jth channel when light is incident from the ith channel, the same holds for P'_{ij}, except that they have a reversed propagating direction. We thus have

$$P_{ij} = 0 \quad \text{whenever } i = j, \tag{5}$$

$$P'_{i1} = P'_{i9} = 0 \quad i = 1, \cdots, 9. \tag{6}$$

Based on the argument we made above, it's easy to write the expressions for P_{ij}'s and P'_{ij}'s as functions of η_1, $\cdots\cdots$ and η_9. Then, the optimization process is as following. The goal is to find an optimized distribution of diffraction efficiencies leading to a fan-out intensity distribution with a minimum power fluctuation and therefore an optimized power budget. For this purpose, we need to generate an objective function [11]. By optimizing the objective function, a well balanced fan-out distribution can be reached. For our problem, the minimum fluctuation nature requires us to establish an objective function of the form

$$E = E_1 + E_2 \tag{7}$$

where

$$E_1 = \sum_{i=1}^{M}\left[\sum_{\substack{j=1\\ j\neq i}}^{N} W_{1ij}\left(\frac{P_{ij}}{\overline{P}} - 1\right)^2 + \sum_{j=2}^{N-1} W'_{1ij}\left(\frac{P'_{ij}}{\overline{P}} - 1\right)^2\right] \quad \text{for } P_{ij} \text{ and } P'_{ij} \geq \overline{P}, \tag{8}$$

$$E_2 = \sum_{i=1}^{M}\left[\sum_{\substack{j=1\\ j\neq i}}^{N} W_{2ij}\left(\frac{\overline{P}}{P_{ij}} - 1\right)^2 + \sum_{j=2}^{N-1} W'_{2ij}\left(\frac{\overline{P}}{P'_{ij}} - 1\right)^2\right] \quad \text{for } P_{ij} \text{ and } P'_{ij} < \overline{P}, \tag{9}$$

in which $\overline{P}$ is the average value of the fan-outs and, to effectively reduce the differences of each fan-outs from the average, we have added before each term the weight factors W_{1ij}, W'_{1ij}, W_{2ij} and W'_{2ij}, so that the fan-outs with bigger differences from $\overline{P}$ will be effectively dragged back. In our calculation, we used two different weights. One has an exponential form; the other has a nonexponential power form and the same results has been obtained. A detailed discussion of our calculation will be presented in a future publication. Here, we provide the calculation results. Minimization of the objective function E eventually gives an optimized diffraction efficiency distribution

$$(\eta_1, \eta_2, \cdots, \eta_9) = (1.0,\ 0.3341,\ 0.2005,\ 0.1435,\ 0.1548,\ 0.1665,\ 0.2492,\ 0.4984,\ 0.0) \tag{10}$$

TABLE III
SUMMARY OF THE FAN-OUT DISTRIBUTION OBTAINED FROM THEORETCIAL ANALYSIS (NOT INCLUDING P'_{ij} PORTION)

Input \ Output	1	2	3	4	5	6	7	8	9
1	0	0.498	0.083	0.033	0.022	0.015	0.015	0.015	0.015
2	0.498	0	0.083	0.033	0.022	0.015	0.015	0.015	0.015
3	0.083	0.083	0	0.033	0.022	0.015	0.015	0.015	0.015
4	0.033	0.033	0.033	0	0.022	0.015	0.015	0.015	0.015
5	0.022	0.022	0.022	0.022	0	0.022	0.022	0.022	0.022

The resultant fan-out power distribution is summarized in Table III.

For our backplane bus proposed in this paper, the sensitivity of a photodiode detector is a function of the input signal modulation speed, the bit error rate, and the wavelength of the signal carrier. For a PIN-FET photodiode with a quantum efficiency of 50% at 1.3 μm, if the data transfer rate is 1.2 Gbt/s and the required error probability after amplification of the detector signal is less than 10^{-9}, then the minimum modulated power required at the detector site is determined to be 5.08×10^{-5} W. From Table III, we see the ratio between the maximum and the minimum fan-out powers is about 47. This implies that the maximum power the detector should be able to handle before it saturates is around 2.39×10^{-3} W. The dynamic range of the detector can be employed to closely approximate the dynamic range of the receiver subsystem that will perform efficiently in conjunction with the optical backplane bus.

A bidirectional backplane bus with optimized power distribution, and thus hologram diffraction efficiency, is provided for the demonstrated system in Section III. Discrepancy between experimental and theoretical results is observed. An optimized system shall provide us with a diffraction efficiency equivalent to the theoretical value, for each holographic element involved. Such an arrangement can be made by individually recording each hologram along a backplane bus.

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 [12]. 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 coordination between devices that are sending and receiving data.

Synchronous, uncompelled asynchronous, and compelled asynchronous [12] are the three standard type of data transfer

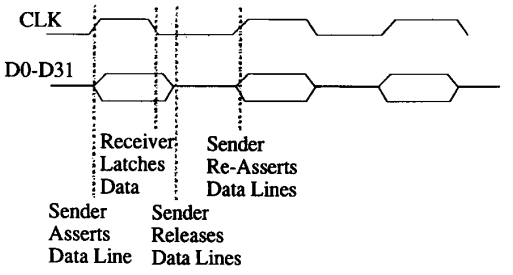

Fig. 10. Timing diagram for synchronous data transfer.

protocols used in electronic backplane buses. This section provides an insight into these three protocols and our demonstration shown in the previous section provides the hardware needed to optoelectronically interconnect these backplane systems.

A. Synchronous Data Transfer

In general, the term "synchronous" means in step with an externally imposed clock signal. In the present context, a synchronous bus provides a bus clock line whose rising and/or falling edges are signals for data transfers to take place. (See Fig 10.) In the figure, it is assumed that both devices are already enabled; for example, device A might be asserting a line labeled "ready to send," and device B might be asserting a line called "ready to receive." Upon receipt of the rising edge of the clock CLK, device A asserts the data on the data lines D0–D31. Device B, upon receipt of the falling edge of CLK, latches the data into its receivers. The cycle repeats with each clock cycle.

B. Uncompelled Asynchronous Data Transfer

As will be shown, this type of protocol is potentially the fastest for block data transfers. The bus timing is controlled by the data sender rather than by the bus itself. The timing is just as in Fig. 10, except that the bus clock CLK is replaced by a strobe driven by the sending board.

As in the synchronous transfer, the receiving board does not have the opportunity to confirm the receipt of the data. (This is the meaning of "uncompelled": The receiver cannot force the sender to slow or halt the transmission.) The transfer speed is left up to board designers. In general, no transfer can take place more rapidly than the slowest board in the system can handle, unless special arrangement is made: A fast sender might have two sending modes, one at normal speed and one at a high speed, with selection based on some system-defined flags encoding the receiver type.

Since the timing signals originate from the sending board, they travel the same backplane distance to the receiver as the data signals. To the extent that transmission skew can be eliminated, the timing signals thus arrive at the same time as the data, independent of location on the backplane. The sender need not wait for confirming signals from the reciever, so the transfer rate is not limited by absolute propagation time but by skew.

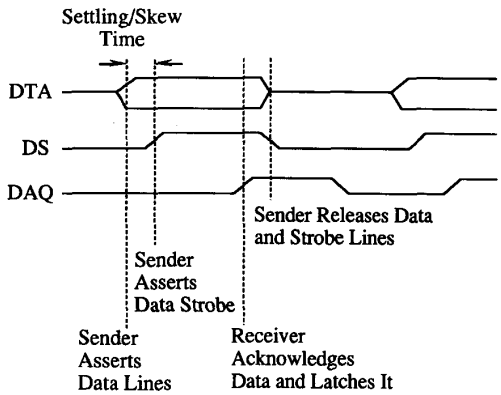

Fig. 11. Timing diagram for a compelled asynchronous data transfer.

C. Compelled Asynchronous Data Transfer

Compelled protocols specify an exchange of signals (handshake) between sender and receiver upon each transmission (see Fig. 11). The sender must continue to assert the data lines until it receives the data acknowledge signal DAQ from the receiver. The transfer speed is thus controlled by both sender and receiver; either can operate as slowly as it needs to. Only if both boards are fast will the transfer rate be high.

This feature is very convenient for system and board designers, since the bus may be designed to a maximum speed much higher than present boards can use, while each board can be designed for any speed without constraint by faster or slower boards. The buses described below—FASTBUS, Futurebus [13], and VMEbus—all employ compelled protocols (FASTBUS employs an uncompelled asynchronous protocol for high-speed block transfers and a compelled protocol for all other transfers).

The maximum transfer speed of any compelled protocol is limited by absolute propagation delay, since the data must be asserted during the time of propagation to the receiver plus the propagation of the acknowledge signal back to the sender. Therefore, compelled transfers take place in a time whose lower bound is twice the propagation time between sender and receiver. This lower bound is largest for transfers between boards at opposite ends of the backplane.

The main advantage of the bidirectional optical bus that we have designed and demonstrated is that equivalent data transfer protocols can be designed for the optical backplane bus. bidirectional 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 [14]. The bidirectional 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 bidirectional optical backplane bus.

Computer to computer optical interconnections have already been successfully demonstrated using optical fibers. Interconnection distances from $1–10^6$ m have been realized. The present challenge is for inter-processor/memory interconnection, where the backplane bus represents the most serious problem for upgrading system performance. VMEbus is the leading standarad for single-board processors and multiprocessors. The bus crates themselves are widely available, while bus communication chipsets are available to streamline the introduction of any new VME product. Mainstream computer companies such as Sun, Silicon Graphics, and Digital Equipment Corporation use VME back planes in some of their machines, while smaller-market, high-performance vendors such as Heurikon, CSPI, and Mercury generally use the VMS platform. VME is also an important military and instrumentation bus, especially the enhanced-protocol XVI standard (VME Extensions for Instrumentation). VME products, including interface chipsets, are generally available to full military specifications (MIL-STD-883B). The polymer-based optical bus successfully developed can be employed to upgrade the interconnection speed and distance. High parallelism will be realized by implementing a linear optical waveguide array. Modulation (electrical to optical) and demodulation (optical to electrical) will be realized at the card boards and can be reduced to 1 psec due to the availability of multi-Gbt/s transceivers. The electrical interconnection distance will be minimized in such a way that the full speed of the processors can be utilized.

The bidirectional 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 [12].

VI. Factors Influencing the Commercial Viability of the Bi-Directional Optical Backplane Bus

Fig. 8 shows that the angular full-width at half-maximum diffraction efficiency is $\pm 2°$, which translates to a large angular tolerance of $\pm 1°$ before the diffraction efficiency drops significantly. The angular misalignment of the input beam θ, will cause 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 [15]

$$\delta L = 2D[\tan(\phi) - \tan(\phi \pm \delta\phi)] \tag{11}$$

where D is the substrate thickness. For our device $D = 3000\ \mu$m, $\theta = 0°$, and $\phi = 45°$. If we consider angular misalignments of $1°$ and $0.1°$, they correspond to $\sim 200\ \mu$m and $\sim 2\ \mu$m 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 bidirectional optical backplane bus to multichip modules (MCM's) [15]. 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 [16]. 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. A situation never arises where there is no optical signal in a channel at a particular time. In fact cross-coupling is an essential part of device operation. A theoretical treatment of the optimum energy distribution acorss the device has already been presented in a preceding section. 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 bidirectional 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 appplications. As already discussed at the outset, this optical bus does not have the constraints of existing electronic bus architectures where RLC effects, electromagnetic interference, and impedance mismatches restrict system performance. 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 holographic optical elements must be perfectly aligned to one another [17]. The source, which can be a semiconductor laser and the detector, can be flip-chip bonded to the glass substrate with the holograms recorded on it. The source and the detector can be aligned to the substrate by means of alignment references that can be photolithographically transferred [17] to the optical backplane bus. The whole system can be bonded in a ceramic housing with a conductive epoxy [17]. Misalignment between the various components is a serious packaging limitation, and it can occur due to assembly errors and operating temperature shifts causing thermal expansion mismatch between components. The wavelength shift of the source laser due to temperature shifts can also be an important contributor to misalignments during and after packaging. An efficient feedback mechanism can be used to reduce this temperature shift, which in turn will minimize the wavelength shift. Packaging problems are bound to occur if the solder joints are not properly placed during flip-chip bonding. The component weight and substrate-component separation need to be taken into consideration during packaging [17]. In effect, if alignment between components is done with care, packaging limitations can be considerably reduced.

VII. Conclusion

We have discussed the advantages and applications of a bidirectional 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. We have described our bidirectional optical backplane bus, operating at 632.8 and 1300 nm. 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 Gbt/s 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. A theoretical analysis aimed at minimizing intensity fluctuations along both directions of the backplane bus, has also been presented.

The polymer-based bidirectional 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 Mbt/s data highways) as long as the appropriate protocols governing the rules of data transaction are provided. Note that all the existing protocols for electrical buses can still be used for the optical backplane reported herein. The only difference is that settling time and propagation delay are much smaller for optical interconnects. The data propagation speed of the polymer-based optical bus is $\sim 0.67\ c$ (c is the speed of light in a vacuum) while there is no settling time due to mismatched impedance. This technology shall provide us with an open bus architecture for all existing and future high performance buses.

References

[1] P. Sweazey, "Limits of performance of backplane buses," in *Digital Bus Handbook*, J. DeGiacomo, Ed. New York: McGraw-Hill, 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] S. Heath, *VMEbus User's Handbook*. Boca Raton, FL: CRC Press, 1989, ch. 3.

[4] R. T. Chen, S. Tang, M. Li, D. Gerald, and S. Natarajan, "1-to-12 surface normal three-dimensional optical interconnects," *Appl. Phys. Lett.*, vol. 63, no. 20, pp. 1883–1885, 1993.
[5] H. Herzig and R. Dandliker, "Diffractive components: holographic optical elements," in *Perspectives for Parallel Optical Interconnects*, P. Lalanne and P. Chavel, Eds. New York: Springer-Verlag, 1993.
[6] R. T. Chen, M. Lee, S. Natarajan, C. Lin, Z. Z. Ho, and D. Robinson, "Single-mode Nd^{3+}-doped graded-index polymer waveguide amplifier," *IEEE Photon. Technol. Lett.*, vol. 5, no. 11, 1993.
[7] S. Tang and R. 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.
[8] R. T. Chen, H. Lu, D. Robinson, M. Wang, G. Savant, and T. Jannson, "Guided-wave planar optical interconnects using highly multiplexed polymer waveguide holograms," *J. Lightwave Technol.*, vol. 10, no. 7, pp. 888–897, 1992.
[9] J. P. G. Bristow, "Recent progress in optical backplane," presented at OE/LASE '94, Los Angeles, CA, Jan. 1994, pp. 22–29.
[10] H. Kogelnik, "Coupled wave theory for thick hologram gratings," *Bell Syst. Tech. J.*, vol. 58, no. 9, pp. 2909–2947, 1969.
[11] W. E. Pinney and D. B. Mcwilliams, Eds., *Management Science: An Introduction to Quantitative Analysis for Management.* New York: Harper and Row, 1982.
[12] Ray T. Chen, "Architecture and building blocks for VME optical backplane bus," *Proc. SPIE.*, vol. 1849, pp. 196–208, 1993.
[13] IEEE Futurebus^{+} P869.1: Logical Layer Specifications, IEEE, New York, 1990.
[14] S. Natarajan, R. T. Chen, S. Tang, and R. A. Mayer, "High speed optical backplane bus with modulation and demodulation capabilities," *Proc. SPIE*, vol. 2153, pp. 344–351, 1994.
[15] S. Tang, R. T. Chen, D. Gerold, M. M. Li, S. Natarajan, J. Lin, N. Chellappan, and M. Peskin, "Design considerations for high packaging density optical bus array," *Proc. SPIE*, vol. 2153, pp. 227–235, 1994.
[16] R. T. Chen and H. Lu, "Polymer-based 12-channel single-mode wavelength division multiplexer on GaAs substrates," *Proc. SPIE*, vol. 179, pp. 44, 1992.
[17] S. K. Patra, J. Ma, V. H. Ozguz, and S. H. Lee, "Alignment issues in packaging for free-space optical interconnects," *Optical Engineering.*, vol. 33, no. 5, pp. 1561–1570, May 1994.

Srikanth Natarajan received the B.E. degree in electronics and communication engineering from the University of Madras, India, in 1989, and the M.S. degree in electrical engineering from the University of Texas at Austin in 1991. He is currently working toward a doctoral degree in the Department of Electrical and Computer Engineering of the University of Texas at Austin. His research interests include the design and fabrication of polymer-based holographic devices for optoelectronic interconnects with emphasis on holographic backplane bus design for computing environments.

Chunhe Zhao received the B.S. degree in applied physics from the National University of Defense Technology, Changsha, P.R.C., in 1988. Currently he is a graduate student working toward his Ph.D. degree in electrical and computer engineering, the University of Texas at Austin. His research interest includes optical interconnects based on holographic gratings, optical backplane bus, and holographic neural networks.

Ray T. Chen, photograph and biography not available at the time of publication.