ARCHITECTURE AND BUILDING BLOCKS FOR VME OPTICAL BACKPLANE BUS

Ray T. Chen
Microelectronics Research Center
Department of Electrical and Computer Engineering
University of Texas, Austin
Austin, TX 78759

ABSTRACT

To eliminate the intrinsic problem associated with electrical interconnects[1], processor-to-processor interconnects based on fiber optics are under development. The major bottleneck of fiber-based optical interconnects is their point-to-point characteristic which seriously limits their interconnectivity. As a result of the low interconnectivity, multiplexing/demultiplexing (e.g.., 1:1024 and 1024:1[2]) technology was employed to minimize optical channels. This approach significantly reduces the bandwidth of the data transfer rate. Moreover, it is incompatible with most of the IEEE standard buses (such as FASTBUS, VMEbus and Futurebus). The high density highly distributed channel waveguide array presented in this paper is the only feasible solution to provide a lithographically defined optical interconnection network with full compatibility of IEEE standard and special purpose high performance bus systems.

1.0 INTRODUCTION

The limitations of current computer data buses stem from their purely electronic interconnects. These limitations include wide interconnection time bandwidths, large clock skew and large RC time constants. Even 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 Mbit/sec. For present high-speed buses, the electronic limitations are even more pronounced. For example, the VMEbus serial bus (VSB) transfers data at 3.2 Mbit/sec, and its speed degrades to 363 kbits/sec when the two communication points are separated by 25 m [1].

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. 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

0-8194-1075-6/93/$6

Therefore, the RC time constant and thus the input charging time remain the same, independent of scaling. Given the RC parameters of VLSI, the signals will propagate at approximately 0.5% to 20% of the speed of light.

Figure 1 illustrates the interconnection hierarchy for various interconnection scenarios. The speed limit based on the previous discussion becomes more stringent as the interconnection distance increases. (For example, board to board and computer to computer interconnects) For example, the dispersion-limited 1 GHz speed limit is for an electrical interconnect length not longer than a few millimeters, and the 100 MHz speed limit holds for an interconnect length not longer than a few centimeters.

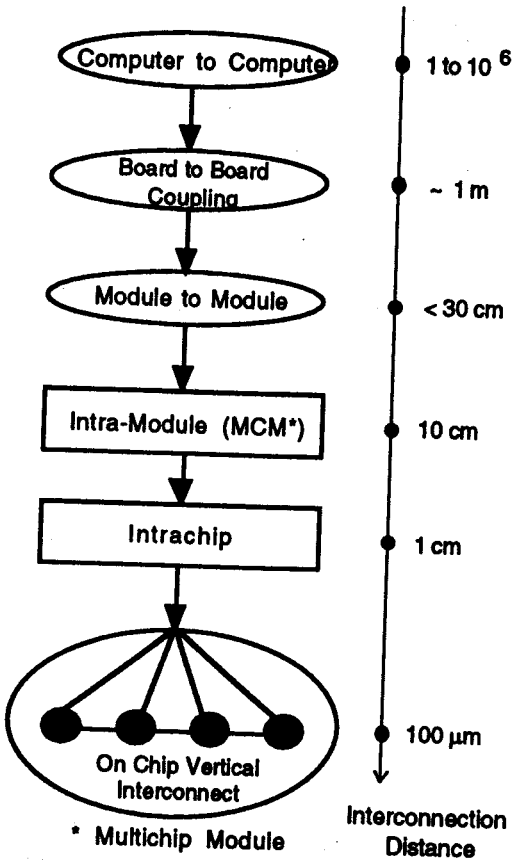

Figure 1
Interconnection Hierarchy

Computer to computer optical interconnections have already been successfully demonstrated using optical fibers. Interconnection distances from 1 m to 10^6 m have been realized. The present challenge is for interprocessor/memory interconnection, where the backplane bus represents the most serious problem for upgrading system performance. VMEbus is the leading standard 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 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 VME platform. VME is also an important military and instrumentation bus, especially the enhanced-protocol VXI 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. The electrical interconnection distance will be minimized in such a way that the full speed of the processors can be utilized.

2.0 FUNDAMENTAL LIMITATIONS OF THROUGHPUT OF ELECTRONIC BACKPLANES

Unlike the optical interconnect, the electrical interconnect on the backplane level has two serious problems that significantly limit the data transfer speed. They are signal propagation time delay and skew between parallel bus links. These problems become more stringent when the linear dimension of a backplane for multiprocessor machines increases.

The propagation time does not affect the maximum data rate of an uncompelled asynchronous block transfer (such as the source-synchronous block transfer (SSBLT) which is proposed for addition to the VME standard). However, it does limit other types of bus transactions: address transfers, handshaked single-word transfers, bus contention, and so on. Estimates have been made by Sweazey [3] of the sustained throughput, i.e., the data transfer rate averaged over a time which is long compared to the duration of a single transaction. Assuming that bus overhead is 200 ns per read and 100 ns per write operation, and assuming reads outnumber writes by 2 to 1, Sweazey calculated the sustained throughput as a function of block transfer speed (burst speed) and of the number of bytes per transfer. For 64-byte transfers, the calculated sustained throughput is 196 MB/sec for a burst rate of 400 MB/sec, and 384 MB/sec for infinitely rapid block transfers.

The propagation speed for the electronic bus is at present greatest for backplane-transceiver logic (BTL) backplanes such as Futurebus: about 0.18c, giving a 15 ns round trip time for a 40 cm backplane. This cannot decrease by much, since it is based on the extremely low driver capacitance of 5 pf/driver provided by BTL.

Uncompelled block transfers are limited by bus line skew. The principal cause of this is speed variations (time jitter) between transceiver chips. This jitter is at least 5 ns even for a well-designed set of transceivers. This means that there will be a total skew between data lines and strobes of up to 20 ns from transmission to receiver latching. In addition, there may be skew in the transmission lines themselves, due to unequal capacitive loading, unequal distances to ac grounds, or for some other reason. (Futurebus transmission lines are purposely skewed to ensure that data arrive before strobe.) These skews limit the attainable transfer rate to 40 megatransfers/sec or 160 MB/sec for a 32-bit bus.

Electronic bus lines are not typically terminated in a matched impedance, since this would require the drive currents to be too high. Therefore, the bus line will not settle until all end reflections have subsided (several round trip times). By contrast, polymer bus lines may be terminated in anti-reflection coatings, suppressing end reflections and reducing settling time to zero.

It is obvious up to this point that optical interconnects based on a polymer optical data bus have a number of advantages such as much shorter propagation time and skew and no settling time.

3.0 OPTICAL EQUIVALENT FOR ELECTRONIC BUS LOGIC DESIGN

Before discussing an optical backplane design in its entirety, we must present optical equivalents of necessary bus components, such as bidirectional transmission lines, stubs, transmitters and receivers. Further, optical equivalents of line voltage, logic levels and open-collector and tri-state line driving and receiving must be derived. Once these issues are resolved, the way will be clear for defining an optical backplane that is fully compatible with existing VMEbus protocol.

The optical equivalent of a PC board trace is a polymer-based optical waveguide. Polymer deposition techniques are very useful here, since surface waveguides can be fabricated on a wide variety of surfaces, including PC boards themselves. A very important consequence is the ability to provide modulation and demodulation for the same backplane.
An unloaded (no boards attached) PC board trace has a typical signal propagation speed on the order of 0.6 c. The speed drops to below 0.2 c for a fully loaded bus line. The polymer which forms the optical waveguide has an index of refraction $n \cong 1.5$. The optical signal propagation speed is $c/n \cong 0.67$ c, similar to that of the unloaded electronic bus line.

It is important to note, however, that there is no optical analogue to driver capacitance from attached boards, which causes loading of electronic bus lines. Therefore, the optical signal speed retains the same high value regardless of the presence or absence of line drivers in the system. This means that the optical bus round-trip delay time will be lower by a factor of three than that of the electronic bus. A connection to an electronic bus line takes the form of a stub or tee junction in the PC board trace; usually, such a stub connects to a line transceiver. The optical equivalent of a stub is high-efficiency unidirectional coupling, which allows light from a second waveguide to be coupled into the optical bus line, and low-efficiency coupling, which will be used to couple light out of the bus for detection. The key feature of the unidirectional coupling is that while light is injected from the stub with high (>90%) efficiency, light propagating in the bus line and passing the coupling is almost unaffected by it (<1% fan-out). This is not the case, however, for light propagating in the other direction which suffers high losses at the coupler.

Because of this, an optical waveguide with stubs attached is necessarily unidirectional. The optical equivalent of an electronic bus line thus involves two parallel optical waveguides, each carrying light in the opposite direction.

Optical waveguide signals can be detected at much lower levels than the level of transmission. For instance, a laser diode might generate 10 mW, while a photodetector (e.g., a p-i-n diode) can detect a 1 μW signal. This implies high fan-out capability, i.e., many receivers can be connected using low-efficiency couplings to a bus line driven by one transmitter.

Figure 2 shows the optical equivalent of a single bidirectional electronic bus line. The drive current provided by each electronic transceiver powers the corresponding laser diode, whose output is split and injected into both waveguides. Each photodiode detects light from either waveguide, since the low-efficiency couplings lead to waveguide segments which are merged with a unidirectional coupler. Each photodiode current powers the corresponding electronic receiver.

The scheme in Figure 2 may be considered fully equivalent to an electronic bus line driven by open-collector drivers, terminated in pull-up resistors, if the following identification is made:

- The state in which no light is present on either waveguide (no laser diode is operating) corresponds to the unasserted electronic line which is pulled high by the pull-up resistors; and
- The state in which there is light in both waveguides (one or more laser diodes are operating) corresponds to the asserted (low level) electronic line.

Note that there is no optical effect corresponding to the wire-OR glitch. When two diodes are on, and one is turned off, every detector continues to receive light from the other as long as it remains on. This is because optical powers from two diodes simply add in the waveguide; any detected nonzero power (above a characteristic noise threshold) corresponds to the asserted line state.

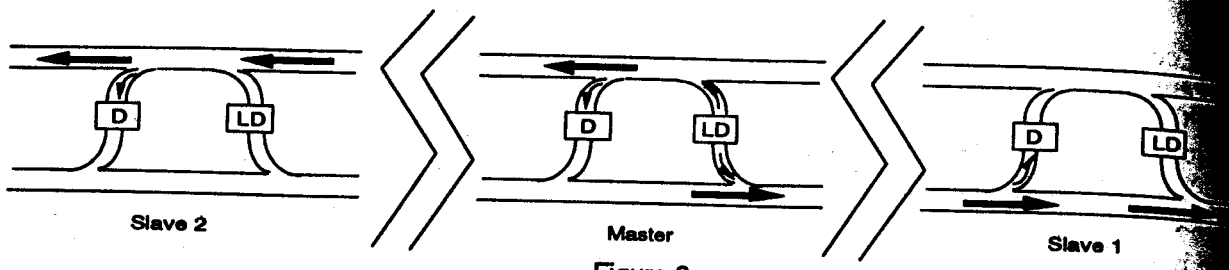

Figure 2
The Optical Equivalent of a Bidirectional Electronic Bus Line Driven by Open-Collector Drivers (3-D Coupling Involving TIR Holograms is not Shown). Communication between one master (e.g., processor) and two slaves (e.g., memories) is clearly indicated (boards not shown). Commands from the master to slaves 1 and 2 are carried out using the bottom and the top polymer buslines, respectively. In this specific scenario, the Master is broadcasting signals that are received by Slaves 1 and 2, and the high power margin of the operation is preserved.

The analogue to resistive transmission line termination is treatment of the ends of the optical waveguides that the reflected power is zero. This is done by implementing an antireflection coating or by fabricating the waveguide ends at the Brewster angle, with an absorptive beam dump outside the waveguides. If this optical isolation is done, settling-time effects are removed. Note that the optical bus line just described has exactly two states: light present (low) and light absent (high). This suffices, as just described, to represent a two-state open-collector-driven line exactly. However, it is insufficient to represent a tri-state-driven line; a tri-state driver has an asserted-high state which is distinguishable from the disconnected state. Where the tri-state line must be exactly emulated, the corresponding optical bus line can consist of a pair of the lines just described, that is, four optical waveguides, with a separate laser diode and a separate photodiode for each pair. See Figure 3. In this case, the state with no light in any waveguide represents the disconnected state (no device asserting) as before, while the state with light present in the top pair (for example) indicates the

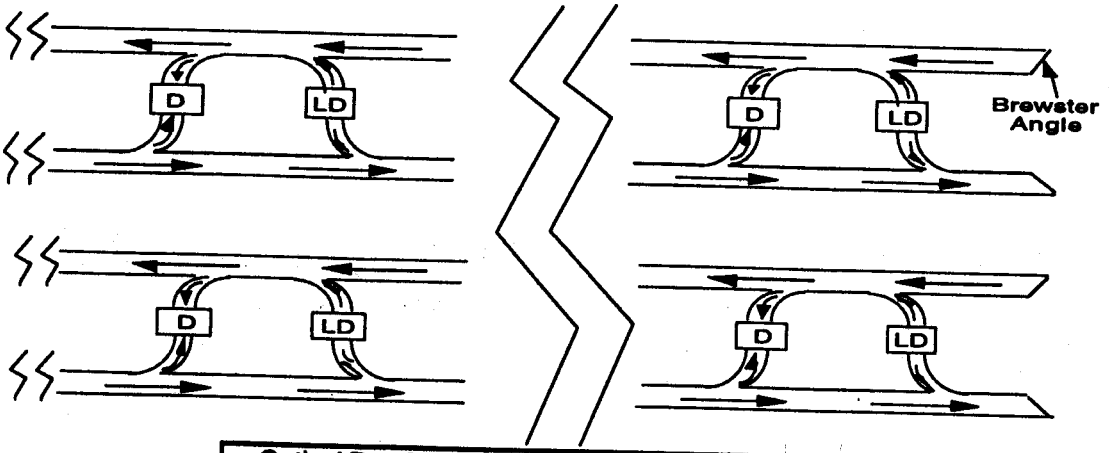

Optical Bus State	Tri-State Logic State
Light in Upper Waveguide Pair	Asserted High
Light in Lower Pair	Asserted Low
No Light	Disconnected (High Impedance)
Light in Both Pairs	No Equivalent

Figure 3
The Optical Equivalent of a Tri-State Bi-Directional Electronic Bus Line

asserted-high state, and that with light in the lower pair, the asserted-low state. Note that this scheme over-represents the tri-state line, as there is a fourth state, that with light in both pairs of waveguides.

As in Figures 2 and 3, laser diodes and photodetectors are located either on the associated cardboards or on the backplane itself. In both cases, optical waves can be easily coupled into and out of the optical bus by using a total internal reflection hologram . The phase-matching diagrams for coupling in and out of the ODB are shown in Figure 4.

For a fixed wavelength, the deviation due to wavelength spreading is given by

$$\ell \Delta\phi = \ell \left[\frac{\cos \delta'}{\lambda \cos\phi} \right] \Delta\lambda \qquad (1)$$

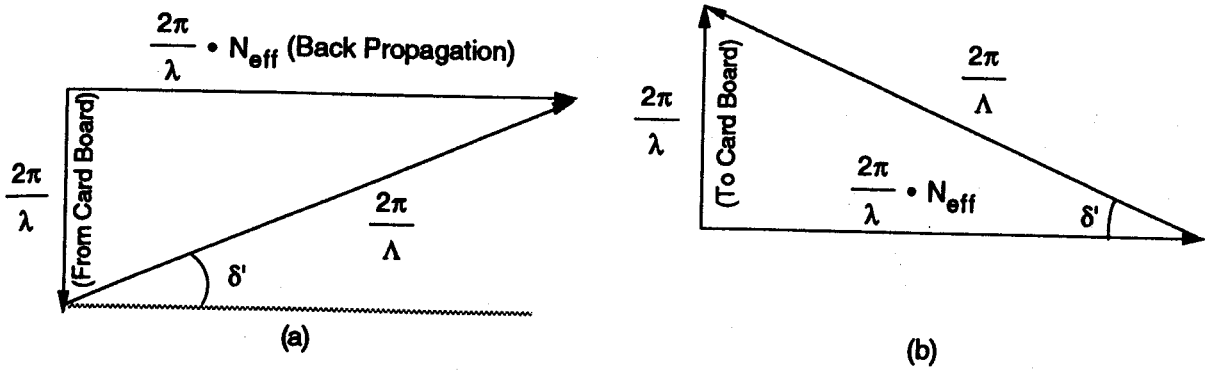

Figure 4
Phase-Matching Diagram for (a) Coupling Into and (b) Coupling Out of the Polymer-Based Optical Data Bus (ODB)

where ℓ is the distance from the TIR hologram to a detector or to a laser diode, ϕ is the vertical 3-D board-to-board coupling angle ($\phi = 0$ for vertical coupling shown in Figure 1), and δ' is the TIR grating tilt angle. For a setup with $\Delta\lambda = 4\ \mu m$, $\delta' = 60°$, $\ell = 10\ \mu m$, and $\lambda = 1\ \mu m$, we have $\ell\Delta\phi \approx 32\ \mu m$. A photodetector with 32 μm linear dimension of active area is capable of providing a multi GB/sec demodulation speed. Therefore, implementation of a real time signal demodulation scheme similar to that of Figure 3 is enough to provide the speed requirement. Therefore, the high-efficiency coupling from SEL to backplane polymer-based optical bus can be easily realized.

The previous section demonstrates the ability of optical technology to emulate present electronic bus technology. This section considers physical details of implementation, with particular attention paid to preserving compatibility with present boards developed for some particular bus.

Figure 5 shows a circuit with which an electronic transceiver utilizing an open-collector driver can connect to optical bus line. Such a circuit is necessary to allow the use of present boards, since only one line connects the transceiver (located on the plug-in board) to the bus line through the board connector. It is essential for this circuit that the buffer driven by the photodiode be able to pull down the receiver line below threshold of the lower voltage level (e.g., 0.8 V for TTL and negative for ECL) but not sink enough current to cause the laser diode to reach threshold. Fortunately, this is not difficult because of the strongly nonlinear current-voltage characteristic of the laser diode.

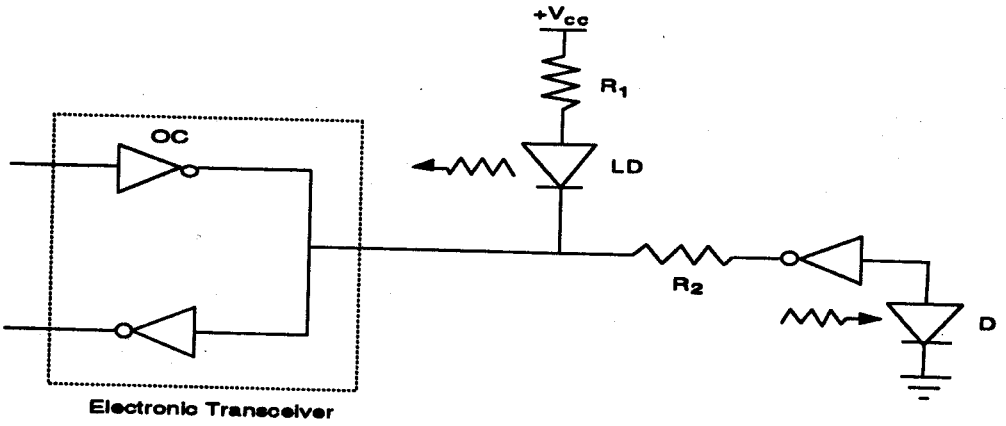

Figure 5
An Interface Circuit Allowing a Single-Line Transceiver to Drive the Optical Bus Line. Resistors R_1 and R_2 are chosen so that the driver can power the laser diode, but the detector buffer cannot.

The circuit in Figure 5 can be integrated into a compact three-pin surface-mount optoelectronic integrated circuit (OIC) as shown in Figure 6. Part (a) shows the package layout, and (b) shows the optical waveguide connections to the OIC. Each bus line consists of two optical waveguides, one OIC per plug-in board, couplings and stub waveguides, and power and ground lines (metal traces) to power the OICs.

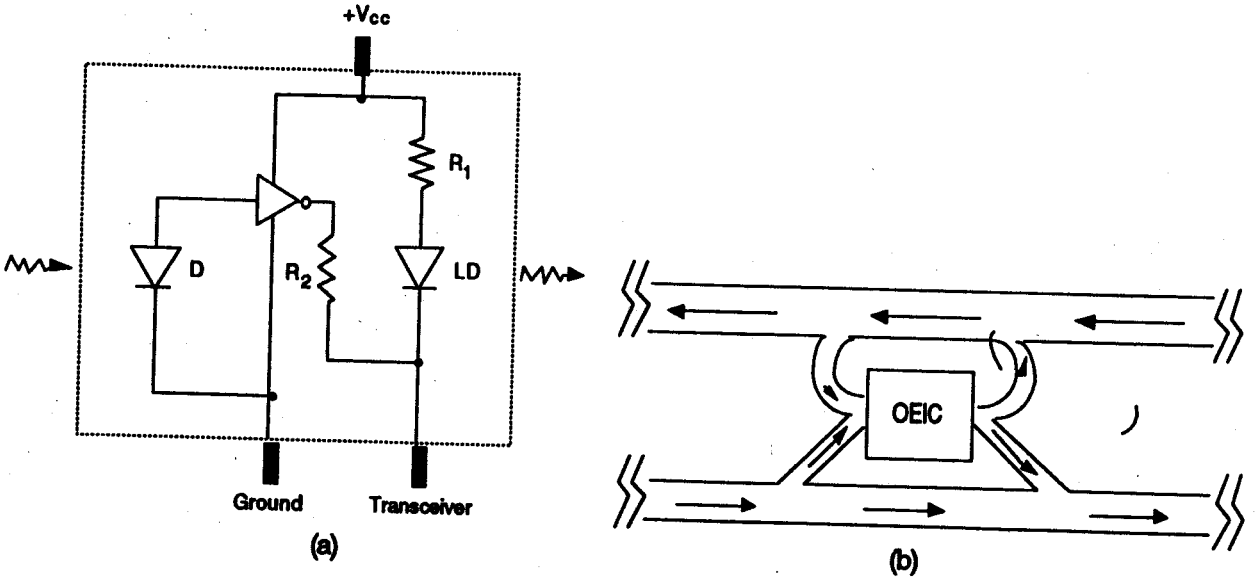

Figure 6
An Integrated Circuit Version of the Interface of Figure 5. (a) Package layout; (b) optical waveguide connections.

4.0 INTEGRATION OF DETECTOR ARRAYS

SEL arrays are clearly the best choice for the proposed optical VME backplane bus. Implementation of a single optoelectronic circuit for a single VME optical bus line was introduced in the previous subsection. The currently developed laser diode array works from 800 to 900 nm with up to 24 GHz modulation speed and 0.3 W CW power in the main lobe. Typical arrays are 80 μm wide and 1000 μm long. For single element (~ 2 μm wide) devices, theoretical 3 dB modulation bandwidths are in the 30-35 GHz range. The best experimental result is 24 GHz. Other commercially available laser diodes with TTL and ECL driving circuits at a 1 Gbit/sec are currently available through a number of venders including HP, Sharp and Honeywell.

For demodulation, hybrid technologies combining both III-V and Si technologies are the best approach for high-speed receivers. High speed photodetectors working at the 800 nm region, e.g., Si and InGaAs/InP, can be integrated either monolithically or flip chip onto an Si-based amplifier circuit. The same GaAs chips can contain the laser diode arrays. A typical photodetector circuit using CMOS technology for Si is shown in Figures 7.

The silicon photodetector is a P^+N junction diode and shares the same N^- Well as its PMOS transistor load. The output of the first stage (source of PMOS transistor load and PD cathode) is connected to the input stage of a CMOS inverting amplifier consisting of NMOS and PMOS transistors. The amplifier/inverter produces digital pulses at V_{out} corresponding to light ON/OFF at PD.

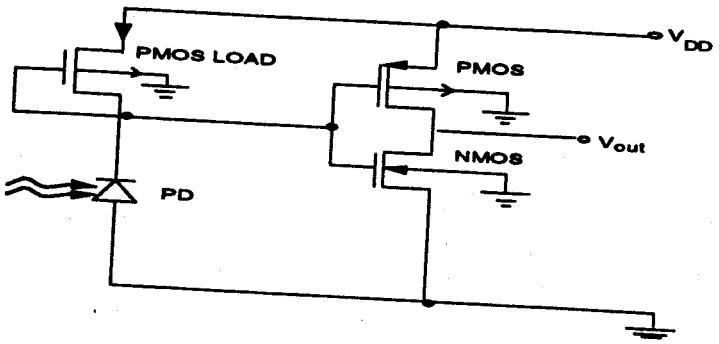

Figure 7
High Speed Photodetector Using CMOS Technology, Interconnection between backplane optical bus and silicon detector and CMOS invertor/amplifier

eature introduced in the proposed optical backplane is that each element of the detector array is nected via a guided wave optical bus. That is, in addition to the conventional optical I/O, each contains an integrated waveguide coupled to the optoelectronics associated with that optical data bus uch an architecture, while perhaps being a natural outgrowth of current photonic and electronic presents significant challenges to device and systems engineers. In particular, it requires a level tronic integrated circuit technology which is still beyond our present capabilities. Hence, hybrid combining both III-V and Si technologies present an excellent opportunity for the realization of ensity backplane optical interconnects and indeed the hybrid approach may be a valuable or future military and civilian applications.

ement shown in Figure 7 requires a significant modification of the interface. Modulation and n speed, however, are much faster (multiGbit/sec/busline).

de-off between speed and detection limit (sensitivity) in PD devices. For instance, in PIN tector junction capacitance increases linearly with the detector active area:

$$C_D \propto A_D \tag{2}$$

This implies that the maximum intrinsic speed (f_{max}) of a photodetector device is inversely proportional to the detector area

$$f_{max} \propto \frac{1}{A_D} \tag{3}$$

Optimum PD device designs on InGaAs/InP should maximize the ratio of active PD area to total PD device area. This will help with the PD figure-of-merit which is speed × sensitivity. The Si PD device designs can be made compatible with advanced CMOS technologies.

Massively-parallel processors may employ optical interconnects with wavelength division multiplexing (WDM) or time division multiplexing (TDM). In this application, III-V-based PDs may be preferred over the Si-based detectors. This is due to the fact that TDM demands very high speed MUX/DEMUX circuitry along with very fast PDs.

5.0 BUS PROTOCOLS

This section will discuss the general features of bus data transfer protocols; the goal is to present the major types of protocols in order to be able to assess the maximum potential data transfer speeds of each. To make the ODB transparent to the current existing VMEbus subsystem, equivalent protocols will be used for the optical bus.

We will concentrate on the mechanism of synchronization of data transfers between one board and another, and not on the many other tasks taken care of by the protocols, such as arbitration of bus requests and of interrupts. For high-data-rate applications, the block transfer of large amounts of data must be made as rapidly as possible. (It must, however, be noted that this is not the only problem. Even if block transfers were infinitely fast, the overall throughput would remain finite due to bus contention, addressing, and so on as addressed in the study by Sweazy discussed previously.)

Assume that device A is continuously sending data to device B. The bus must provide a mechanism for both devices to base their timing. There are three general types of transfer protocols: synchronous, uncompelled asynchronous, and compelled asynchronous.

5.1 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 Figure 8. 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.

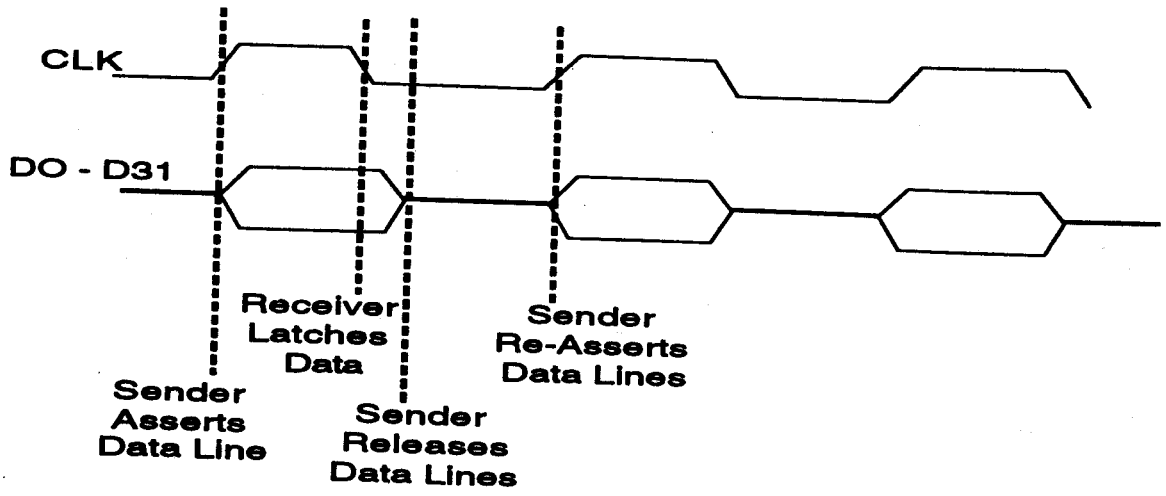

Figure 8
Timing Diagram for Synchronous Data Transfer

The clock half-period τ is chosen to allow for settling time, if any, of the data lines and for clock skew. An important feature of every synchronous bus protocol is that the transfer speed is determined once and for all by the bus specification. For example, the NuBus microcomputer bus has a 10 MHz clock. When a synchronous bus is introduced, the clock speed must be compatible with then-current technology; the bus cannot 'grow' with improving board speeds. For this reason, synchronous buses are avoided in high-data-rate applications.

5.2 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 Figure 8, 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 timing signals from the receiver, so the transfer rate is not limited by absolute propagation time but by

COMPELLED ASYNCHRONOUS DATA TRANSFER

Compelled protocols specify an exchange of signals (handshake) between sender and receiver upon each transmission. See Figure 9. The sender must continue to assert the data lines until it receives the data acknowledge signal DAQ from the receiver.

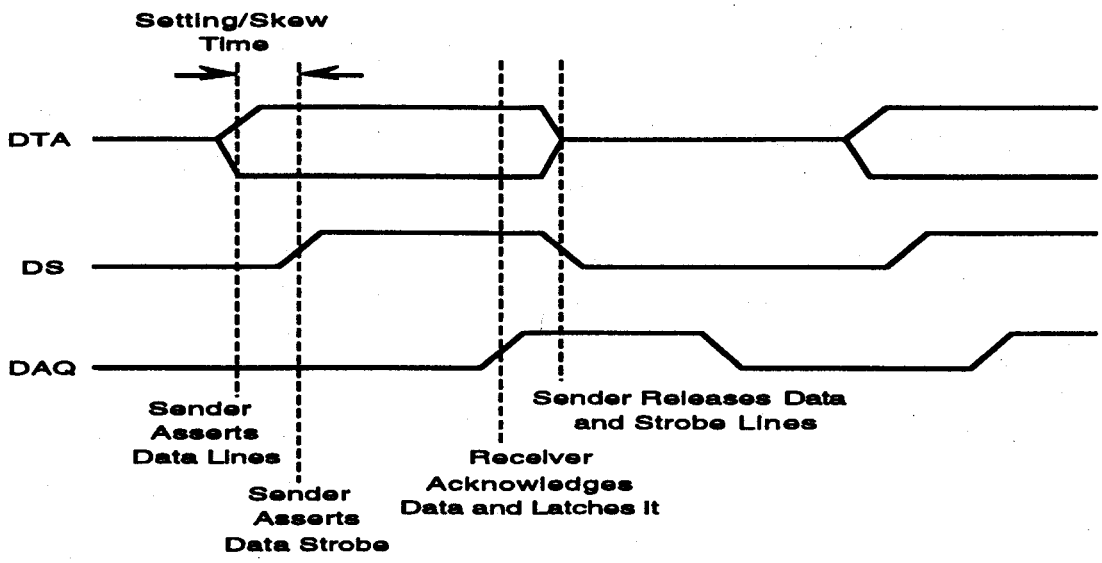

Figure 9
Timing Diagram for a Compelled Asynchronous Data Transfer. The transfer is initiated by the sender, which asserts the data lines DTA, waits a settling/skew time, and asserts the data strobe DS. The receiver, seeing DS, latches the data and asserts the data acknowledge DAQ. The transmitter releases the data lines upon receiving DAQ.

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 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.

5.4 VMEbus

Like Futurebus, VMEbus evolved as a 32-bit computing standard bus. In 1981 it was placed into the public domain by Motorola, and in 1987 it became IEEE/ANSI Standard 1014--1987. VMEbus products are more generally available than FASTBUS or Futurebus because of the wide acceptance of the standard by mainstream computer companies, such as Sun, Silicon Graphics and DEC. Therefore, implementation of an optical VMEbus will be beneficial for physics research and for both military and civilian applications.
As has been described, it is the VMEbus upon which we propose to demonstrate our optical backplane technology. A more detailed description of the electrical, optical, mechanical and protocol specifications has been given. Only the general features of the bus will be described here.

Unlike the previous two buses described, VMEbus does not multiplex address and data lines. This is because VMEbus developed from the different assumption that the majority of bus transactions would be single-word reads or writes to memory, rather than block transfers. By the use of non-multiplexed lines, a bus master initiates a memory write by presenting address and data simultaneously, which is faster than the necessary procedure for multiplexed lines, in which the data can only be sent following the address transfer.

Because of the large number of lines, the VMEbus consists of a backplane with two sets of 96 lines each. TTL and ECL open-collector or tri-state drivers are used. A VMEbus card thus has two 96-pin connectors, P1 and P2. 128 of these lines are described by the standard, while 64 are left to the system designer. (For 16-bit data transfers, a board needs to use only connector P1; half of the data lines are on connector P2.) The bus lines are terminated to provide pull-up for open-collector drivers or disconnected tri-state drivers; because of drive current limitations, the termination is too large to provide optimal transmission-line termination, and there is a settling time as described previously. The standard TTL VMEbus gives 35 ns as the standard settling interval.

6.0 CONCLUDING REMARKS

As far as other instrumentation buses, such as CAMAC (IEEE583), FASTBUS (IEEE960), and other special-purpose high speed backplane buses are concerned, the physical layer developed is also adaptable to these buses. The polymer-based optical bus demonstrated recently 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)[9] 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. 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 ~0.67 c (c is the speed of light in a vacuum) [10] while there is no settling time due to mismatched impedance. The projected features of the proposed VME optical bus for highly parallel, system-wide communications are summarized in Table 1 with the existing electrical VMEbus as a reference. The technology developed in this program shall provide us with an open bus architecture for all existing and future high performance buses.

7.0 REFERENCES

1. Steve Heath, Chapter 3 of VMEbus User's Handbook, CRC Press, Inc. Florida, 1989.
2. T. A. Lane, J. A. Quam, B. O. Kahle and E. C. Parish, "Gigabit Optical Interconnects for the Connection Machine," SPIE OE/FIBER, Vol.1178 (1989).
3. P. Sweazey, "Limits of performance of backplane buses," in Digital Bus Handbook, edited by J. De Giacomo (McGraw-Hill, New York, 1990).
4. Ray T. Chen, et al., "Reconfigurable Optical Interconnection Network for Multi-Mode Optical Fiber Sensor Array," in press, Opt. Eng., May 1992.
5. A. Yariv, Introduction to Optical Electronics, Chapter 11, second edition, McGraw-Hill, 1976.
6. R. T. Chen et al., Appl. Phys. Lett., vol 60, 536-538 (1992).
7. D. Botez, et al., Appl. Phys. Lett., vol 54, 2183 (1989).
8. S. Pri-Tal, "VMEbus," in Digitalbus Handbook, Edited by De Giacomo (McGraw-Hill, New York, 1990)
9. IEEE Futurebus+ P869.1: Logical Layer Specifications, Published by IEEE, New Yor, 1990.
10. Ray T. Chen, et al., "Polymer Microstructure Waveguides on Al_2O_3 and BeO Substrates for Optical Interconnection," Appl. Phys. Lett., vol. 56, 709 (1990).
11. S. Pri-Tal, "VMEbus," in Digital Bus Handbook, edited by J. DeGiacomo, McGraw Hill, New York (1990).
12. R. V. Balakrishnan, "Transceiver technology and design," in Digital Bus Handbook, edited by J. De Giacomo (McGraw-Hill, New York, 1990).

Table 1 Projected Features of Polymer-Based Optical VMEbus with the Existing VMEbus as a Reference

	Feature	VME Polymer-Based Optical Data Bus (ODB)	Electrical VMEbus
1.	Settling Time	0 sec [a]	35 nsec [a] [11]
2.	Propagation Speed	0.67 c [b]	0.05 c - 0.1 c [b]
3.	Bus Line Skew (ns)	~2 nsec [c]	20 nsec [3]
4.	Peak Operating Speed	240 Gbytes/sec [d]	40 Mbytes/sec [12]
5.	Transfer Protocol	Compelled asynchronous [6]	Compelled asynchronous [12]
6.	Packaging Density (# of bus lines/cm)	1250/cm [e]	~10/cm
7.	Interconnectivity	High [f]	Low [f]
8.	Data Width	to 32 bit [g]	32 bit [g]
9.	Address Width	to 32 bit [g]	32 bit [g]
10.	Number of Parallel Bus Lines for Transmission	96 [g]	96 [g]
11.	Transceiver Types	Open-Collector Driver [h] Tri-State Driver	Open-Collector Driver [h] Tri-State Driver
12.	Logic Family	ECL and ECL [i]	TTL and ECL [i]

a For electrical VMEbus line, due to impedance mismatch, bus line won't settle to a new steady state for several round trip reflection times. However, the polymer-based optical bus does not have this problem. Optical Signal can be 100% terminated at the end of the bus line using antireflection coating.

b Polymer-based optical bus has an effective refractive index around 1.5 which gives a propagation speed of 0.67 c while electrical bus lines have a distributed RLC which limits the signal propagation speed to 0.05 c to 0.1 c.

c Since polymer-based optical data bus (ODB) has propagation speed ~1 order of magnitude higher than electrical bus, its bus line skew can be reduced by at least one order of magnitude.

d With 32 bit data width, we have 240 Gbyte/sec (60 Gbit x 32/sec[6], see Objective 6 for details); analog-to-digital conversion ratio is 1 in this case. Limitations shall be imposed by the processors.

e Previously demonstrated result[6], two orders of magnitude better than electrical packaging .

f Daisy chain architecture is implemented for electrical VMEbus. This architecture has the lowest interconnectivity. Due to high packaging density of the ODB (1250 channels/cm), other architectures such as hypercube and star, can be implemented.

g Current VMEbus uses a minimum of 96 lines, including 32 bits wide in data, 32 bits wide in addressing and other control lines. However, ODB can utilize wavelength division multiplexing (WDM) to reduce the number of lines (if needed) as well as multilayer 3-D optical packaging; see Notes c and e.

h The open-collector driver has two states, disconnected (logic level 1) and low voltage (logic level 0). The tri-state driver has three states, disconnected, asserting low voltage (logic level 0) and high voltage (logic level 1).

i TTL is transistor/transistor logic, ECL is Emitter Controlled Logic. ECL-based VMEbus lines have lower charging time and are thus faster than TTL-based bus. For our ODB, ECL will be implemented to reduce the charging time for modulation (electrical to optical) and demodulation (optical to electrical).