

High Speed, High Capacity Bused Interconnects Using Optical Slab Waveguides

Martin Feldman, Ramachandran Vaidyanathan, and Ahmed El-Amawy

Department of Electrical & Computer Engineering
Louisiana State University, Baton Rouge, LA 70803, USA
{feldman, vaidy, amawy}@ee.lsu.edu

Abstract. In this paper we consider the use of optical slab waveguides as buses in a parallel computing environment. We show that slab buses can connect to many more elements than conventional electrical or fiber optic buses. We also introduce a novel multiplexing scheme called mode division multiplexing that vastly increases the number of independent channels that a single slab can support. We show that optical slab waveguides have, in principle, capacities of over a million independent channels (distinguished by about 1000 “out-of-plane modes” and about 1000 wavelengths) in a single physical medium, with each channel capable of sustaining a load of over 1000. This becomes comparable to the high capacity of a free space optical system, but with the ability to broadcast each light source to many physically separated locations. Preliminary experiments on the “sawtooth slab bus” point to the feasibility of practical slab buses. We also present a bus arbitration example that uses the high capacity and loading of slab buses to achieve sublogarithmic arbitration time.

1 Introduction

Buses are essential components in most computer systems. Their uses range from buses within processors and host buses on a board, to backplane buses connecting boards and global shared buses in multiprocessor systems.

As interconnects, buses have several desirable features. Being shared resources, they can be used more efficiently and flexibly than dedicated connections. Bus-based systems support broadcasting easily and naturally, have reasonable costs, and admit simple layouts. Besides traditional multiple bus systems [3, 11, 14, 24, 26], buses have been used in a variety of computing structures including enhanced meshes [2, 4, 22], reconfigurable networks [8, 12, 17, 19] and in the context of emulating interconnection functions and topologies [6, 7, 13, 15, 16].

As the number of connections to a bus increases, the following problems arise:

Loading: In this paper we will use the term *loading* loosely to mean the maximum number of connections that can be made to a bus without significantly degrading the quality of the signal on the bus.

Connections to an *electrical* bus cause capacitive loading that limits the rate at which the signal switches state reliably (i.e., bus clock rate). In addition, cross talk problems arise from the close proximity of high frequency signals in adjacent buses. With existing technology an electrical bus operating at a few hundred MHz, can only support about 30 connections. A fiber optic bus can connect to a somewhat larger number of elements (about 100), still inadequate for moderately large systems.

Arbitration cost (for asynchronous systems): As the number of masters connected to a bus increases, so does the likelihood of multiple masters attempting to use the bus simultaneously. Such conflicts have to be resolved by an arbitration mechanism, whose complexity (cost, delay) depends directly on the number of bus masters.

In this paper we propose buses that use *optical slab waveguides*. Unlike optical fibers that have a single (or a few) transmission modes, slabs can support several thousands of modes. To a large extent, slab buses can overcome both of the above problems of traditional buses. We will show that, in principle, a single slab waveguide can admit over a million independent channels, each capable of supporting a loading of over 1000. This high capacity of slab buses can also be used to devise a fast bus arbitration scheme. Traditional (single) bus-based systems are also limited by bus bandwidth. Multiple bus systems, capable of overcoming this limitation are easily implemented on a single slab waveguide, whose independent channels can be treated as individual buses.

Much previous work on optical buses has centered on the adaptation of components used for optical fiber communication. This is understandable since optical fiber communications is an established technology. However, many of the considerations important to long haul optical fiber communications are not relevant to the communication between proximate processors. In order to operate over long distances at high frequencies, optical fibers are designed to have minimal attenuation, less than 1 db/kilometer, and minimal dispersion. Clearly, the low attenuation and dispersion required for transmission over many kilometers is totally unnecessary between processors at most a few feet apart. Moreover, the low dispersion is usually obtained by restricting the fibers to single mode operation. This requires the use of laser diode transmitters for efficient light coupling to the fiber. In addition, if the signals from several fiber inputs are passively coupled or fed into a single fiber (as in a bus), most of the available power will be reflected. This is because the optical power that is carried by a single mode fiber with input from a single source is determined by the brightness of the source. This power cannot be increased by the addition of other sources. Not only does this seriously reduce the power available in many optical bus configurations, it also couples some of the power back into the laser that emitted the light. This unwanted feedback can cause instabilities in the operation of the laser.

Free space optical systems have also been studied as interconnects for parallel computing [9, 18, 21]. In a typical free space optical system (Fig. 1), lenses (not shown) image an array of light emitting diodes or lasers on a matching array of photodetectors.




Fig. 1. Structure of a typical massively parallel free space optical system.

They are capable of carrying an enormous amount of data. In addition, with “smart pixels” some processing can be done locally between each photodiode and its corresponding light emitter. Processing can also be done in additional optical components such as liquid crystal arrays or acoustooptic modulators inserted in the optical path. However, free space optical systems have limited applications as buses, where light from each of the input sources is broadcast to all of the detectors.

This paper reports some preliminary results with slab waveguides as buses. In the next section we will show why a single mode fiber optic bus can support a loading of only about 100. Slab buses, on the other hand, are shown to be capable of connecting to over 1000 elements per channel (Section 3). In Section 4 we introduce a novel multiplexing scheme called mode division multiplexing that is possible only on slab waveguides. This technique can be used in conjunction with other conventional methods such as wavelength division multiplexing to increase the number of independent channels in the bus by a factor of about 1000. In Section 5 we briefly outline some results from preliminary experiments on slab buses. In Section 6 we show how the high capacity of slab buses can be used in a fixed priority-based arbitration scheme that resolves conflict among N w -bit processors in $O\left(\frac{\log N}{\log w}\right)$ steps. The fact that this method can be used to resolve conflicts in access to any shared resource may be of independent interest. Finally in Section 7 we summarize our results and make some concluding remarks.

2 Loading Limitation of Fiber Buses

A typical optical fiber bus consists of an U-shaped optical fiber connected by directional couplers to light sources (on one end of the bus) and detectors (on the other) [5, 10, 20] (see Fig. 2). Normally, directional couplers are used to transfer power from one fiber to the other, while the direction of information flow along the fibers remains the same. The fraction of power transferred between the fibers (*coupling factor*) is a parameter of the directional coupler. To ensure that all detectors receive the same amount of power (regardless of the transmitting light source) different couplers need different coupling factors. Moreover, if one unit




Fig. 2. Fiber optic bus with directional couplers.

of power is transmitted by a source, each of the N detectors only receive $\frac{1}{N^2}$ units of power, as shown below.

For the following analysis we assume a single mode fiber bus and neglect transmission and insertion losses. Number the processors from 1 to N (left to right) as shown in Fig. 2. Let the coupling factors for the light source and detector of processor i be α_i and β_i , respectively. That is, if processor i transmits one unit of power through its light source, then α_i units are transferred to the bus. If processor i is not transmitting and the bus is carrying one unit of power, then α_i units of power are siphoned off from the bus. Likewise at the detector end, if one unit of power appears on the bus, then β_i units of power is transferred across the coupler to the detector of processor i . Clearly, $\alpha_1 = \beta_1 = 1$ as the first processor does not require a coupler.

Let each light source transmit one unit of power. A unit of power transmitted by the first source reduces to $(1 - \alpha_2)(1 - \alpha_3) \cdots (1 - \alpha_i)$ just after the i th source. Since the power available to the detectors must be independent of the transmitting source, $\alpha_i = (1 - \alpha_2)(1 - \alpha_3) \cdots (1 - \alpha_i)$, for each $i > 1$. This implies that $\alpha_i = \frac{1}{i}$. In a similar manner it can be shown that $\beta_i = \frac{1}{i}$. The power on the bus, just after the N th light source is $\alpha_N = \frac{1}{N}$. The detector of processor i receives $\frac{1}{N}(1 - \beta_N)(1 - \beta_{N-1}) \cdots (1 - \beta_{i+1})\beta_i = \frac{1}{N^2}$ units of power.

This inefficient power distribution limits the loading, or the maximum number of detectors that may be connected to the bus. Let N_f denote this maximum for a fiber bus. Also let λ (in microns) and τ (in nsecs) denote the wavelength of the light and the detection time, respectively. Assuming a 100% transmission efficiency, it can be shown that P , the number of photons per detector per mW of incident energy is given by

$$P = 5.04 \times 10^6 \frac{\tau \lambda}{(N_f)^2}.$$

With $\tau = 1$ nsec (corresponding to a Gigabit data rate) and $N_f = 100$, P ranges from about 300 photons in red light ($\lambda \approx 0.6$ microns) to about 500 photons at the limiting wavelength for silicon photodetectors (about 1.0 microns). These numbers are close to the minimum required for detection with ample error margins, especially in the presence of background noise and with more realistic transmission rates. Thus, the maximum loading for fiber buses is about 100.

Though the analysis presented above is for a particular arrangement of sources and detectors on a fiber bus, the essential ideas apply to other fiber bus configurations as well (such as couplers arranged in a binary tree with N leaves with each coupler having a 50% coupling factor). This is because the reciprocal property of directional couplers is a fundamental limitation: it implies that the power carried by a fiber cannot be extended indefinitely simply by adding more inputs. To do so would increase the power per mode, and hence the brightness in the fiber without limit, which would violate the second law of thermodynamics.

3 Slab Waveguides

The term *slab* is used in this paper to characterize waveguides for which the dimensions of the cross section are much larger than the wavelength of the light. Slab waveguides have many modes in which the light may be transmitted. The term *fiber* is used for waveguides that permit only a single (or a few) modes.

It can be shown [25, page 312] that the number of modes, m , in the plane of a slab of width d is given by

$$m = \frac{2d\sqrt{(n^2 - 1)}}{\lambda},$$

where n is the index of refraction of the slab in air, and λ is the wavelength of the light. For the slabs considered here, with d equal to a few mm, m is on the order of several thousands. In contrast with optical fibers, the total power carried by a slab can be proportional to the number of sources (as long as there are at least as many modes as sources).

In a slab bus with a carefully designed geometry, it is possible to divide the power from a laser diode or LED equally among all detectors (assuming no transmission and insertion loss). That is, if the light source transmits one unit of power, then each of N detectors receives $\frac{1}{N}$ units of power. This directly translates to a larger loading, N_s . Under the same assumptions used above to calculate the loading of a fiber bus, the loading, N_s , of a slab bus would approximately be equal to $(N_f)^2$; recall that N_f is the loading of a fiber bus. With even a conservative 10% transmission efficiency, N_s is around 1000, which corresponds to approximately 10 times as many photons per unit time, an order of magnitude improvement over fiber buses, in the number of photons and hence the number of detectors.

It should be pointed out that the parameters used here are for adequately bright light sources, such as laser diodes, and sufficiently sensitive detectors, for example avalanche diodes, that are well within the present state of the art.

Several slab geometries with various properties are possible. The slabs considered in this paper have a “sawtooth” shape (Fig. 3). In these waveguides light is injected and removed from ports spaced along the side. These locations could correspond to the locations of different circuit boards. There is a minimum distance between the last transmitter and the first receiver, so that no element is

shadowed. The light is nominally divided uniformly between the receivers, with minimal loss and reflection. The sawtooth shape may also be folded, so that each card in a rack may have both a transmitter and a receiver.



Fig. 3. The sawtooth slab waveguide; in this figure, the *plane of the slab* is the plane of this page; a view of this plane will be called the *side view*. The direction perpendicular to the plane of the slab will be termed the *top view*.

For our discussion, we classify the transmission modes of the light propagating in the slab into two groups.

In-plane modes (Fig. 4(a)), or modes in which the propagation of light is parallel to the plane of the slab but in general not collimated within that plane.

Out-of-plane modes (Fig. 4(b)), in which the light is collimated, but its propagation is not necessarily parallel to the plane of the slab.

Of course, most modes have components which are both in-plane and out-of-plane.

The in-plane multimode capacity of the slab waveguide is useful in mode mixing, so that the energy from each of the transmitters is uniformly distributed among the receivers. It is also important that the slab waveguide have minimal delay, dispersion, and echo. Echo is defined as light which does not reach a detector by a direct path, but returns after a relatively long delay, possibly interfering with the signal from the next transmission on the same channel. Although we have not yet measured echo in sawtooth guides, we have observed that the polarization of the propagating light is preserved. This implies that conventional polarization techniques, e.g. double passes through a quarter wave plate, could be used to minimize echoes, if necessary.

4 Mode Division Multiplexing

In this section we describe a novel multiplexing scheme called *mode division multiplexing (MDM)*, that significantly increases the number of channels that a single slab can support. We will show that more than 1000 independent channels can coexist on a slab. Since each of these channels can utilize other multiplexing schemes, MDM can increase the number of channels available through these schemes by a factor of 1000.




Fig. 4. In-plane and out-of-plane modes of a slab waveguide; for simplicity only a portion of the slab is shown.

In this section the in-plane modes and the out-of-plane modes are treated very differently: the in-plane modes are well mixed to insure uniform illumination at the detectors, while the out-of-plane modes are well separated to provide multiple communication channels (Fig. 5).




Fig. 5. A section of the slab, showing three out-of-plane modes. These modes are well separated as they propagate in different angles when viewed in a direction perpendicular to the plane of the slab (top view). In the plane of the slab itself, the modes are well mixed.

If light is propagated through the slab waveguide parallel to the plane of the slab, then the light leaving the slab will also be collimated in a direction parallel to the plane of the slab. This is true at each of the slab's exit ports. Additional modes may also be introduced in the slab, in which the light is again collimated, but at various angles to the plane of the slab. The light from these modes exits the slab in a parallel beam at an angle, θ_{air} , to the plane of the slab that is given by

$$\sin \theta_{\text{air}} = n \sin \theta_{\text{slab}}, \quad (1)$$

where n is the index of refraction and θ_{slab} is the angle of the beam within the slab. Because it may undergo either an odd or an even number of reflections, each mode exits at both positive and negative angles, even though it may have been excited by light which only entered the slab at either a positive or a negative angle. These properties have two important consequences: First, light introduced at various angles may be detected separately by different detectors, so that multiple channels of communication are possible within the same waveguide. Second, light may be intentionally introduced at both positive and negative




Fig. 6. A single out-of-plane mode entering and exiting a slab. The figure shows a top view of the slab and, for clarity, only one input port and one output port have been shown.

angles by different transmitters, and then be detected by the same detector. This permits operation of the two transmitters in an “OR” configuration, if that is desired. Note that a mirror may be used to fold the light exiting the slab, so that although it exits at both positive and negative angles a single detector is all that is required for each mode.

The principle of mode division multiplexing is illustrated in Fig. 6, which shows the light from a laser diode being collimated into a parallel beam and filling the slab. Note that Fig. 6 shows only one input port and one output port; similar effects occur at all input and output ports. It is important that the light not be collimated in the in-plane direction, in order to insure good mixing and uniform illumination of the detectors. The collimation could be performed with a cylindrical lens whose axis is perpendicular to the plane of the laser diodes. Such a lens is most simply formed by shaping the input face of the slab.

The light leaving the slab is focused by a second lens, either spherical or cylindrical, to a spot of size about $\frac{\lambda f}{t}$, where f is the focal length of the second lens and t is the thickness of the slab.

Additional laser diodes may also be arranged at angles θ_N to the plane of the slab (Fig. 7). In this figure the input lens has been placed a focal length away




Fig. 7. Multiple out-of-plane modes entering and exiting a slab. The figure shows a top view of the slab. For clarity, only light entering from one input port and leaving through a one exit port of the slab have been shown.

from the LED's and also a focal length away from the slab. This permits the LED's to emit light perpendicular to their plane, facilitating their fabrication on a single chip, while also collimating their light in the slab. The light from each of these additional laser diodes propagates within the slab at an angle given by equation (1). It then exits at angles of $\pm\theta_N$ and is focused by the output lens into two lines. The images are lines because the light is not collimated in the orthogonal direction. As indicated in Fig. 7, a mirror can be used to overlap the two images for improved light collection; it is important to use all the light because the light can potentially be shared between a large number of detectors.

If we take the total useful range of the focused output light as $\pm f$, then the number of resolved spots is $\frac{f}{(\lambda f/t)} = \frac{t}{\lambda}$. The number of useful modes is somewhat less, because of the necessity of providing clearance between channels and because of possible degradation of the modes during transmission through the slab. Nevertheless, for a slab a few mm thick and visible light the number of distinct out-of-plane modes is on the order of 1000. (Note that in addition to these out-of-plane modes, the slab may have several thousands of in-plane modes that are used to support a large loading.)

Wavelength division multiplexing may also be introduced, by using sources at different wavelengths. The different wavelength signals may then be directed to different detectors, for example by using dichoric mirrors, which reflect some wavelengths while transmitting others. In principle this can be done at each of the output ports, and for each of the propagating modes, so that the total number of channels at each exit port would be the product of the number of wavelengths and the number of modes. A more interesting concept is to use a dispersive device, such as an efficient blazed diffraction grating, to direct the different wavelengths to corresponding detectors (Fig. 8). This could also be done in conjunction with mode division multiplexing, again greatly expanding the total number of channels available for communication. In principle, the light signals at different wavelengths could be spatially separated into more than 1000 different channels, resulting in an array of 1000×1000 elements. Each case requires a detector for each channel, but these may be arranged in an array on a single chip (or a few chips).

It is interesting that inputs may also be arranged at negative angles, as shown in Fig. 7 for one angle, $-\theta_N$. Since light launched from these locations arrives at exactly the same outputs as light from the positive angle inputs, this presents the possibility of a permanently "built in" OR function. It also forms the basis for the simplest form of wavelength division and polarization division multiplexing, since it provides convenient input ports for two different wavelengths or orthogonal polarizations.

Finally, it is important to note that time division multiplexing may also be used, further extending the capacity of the slab mode. This is especially significant in its possible application as a pipelined bus, a widely studied model of computation [10, 20].




Fig. 8. Using a blazed grating to demultiplex a WDM signal. The figure shows the spatial separation of two wavelengths (indicated by solid and dotted lines)

5 Experimental Results

Preliminary experiments have been performed with sawtooth optical slabs machined from Plexiglas and polished with 1500 grit sandpaper. The slabs were 3 mm thick, and had 8 input ports and 8 output ports, generally arranged as in Fig. 3. The cross sectional area of each port was $3 \times 3 \text{ mm}^2$, and the ports were spaced 25 mm apart. The separation between the input ports and the output ports was varied between 200 mm and 800 mm.

The source of light was a red light emitting diode (LED) operated in a constant current mode. The detector was a photodiode operating in the photovoltaic mode into a resistive load under conditions which limited the output voltage to 50 mV to insure linearity. No optics was used to match the LED to the input ports, or the photodiode to the output ports. A digital voltmeter was used to perform dc measurements of the photocurrent.

We observed an insertion loss of approximately 10 db, largely attributable to coupling losses at the input and output ports. The uniformity between pairs of input and output ports was approximately 1.5 db, full range between the best




Fig. 9. A portion of the slab showing two dimensional separation of modes and wavelengths at a “tooth” of the slab.

pair and the worst pair. Additional measurements across a cross section of the slab showed that the light distribution became uniform to within a factor of 2 after just a few cm.

In a second experiment light from a HeNe laser beam at 6328 \AA was passed through a $25 \text{ mm} \times 300 \text{ mm} \times 3 \text{ mm}$ Plexiglas slab. The beam did not fill the slab. The angle by which the light was inclined to the plane of the slab was varied up to approximately 60° . Fig. 10 is an oscilloscope trace which shows the angular distribution of the intensity in the exit beam at an angle to the slab of 18° . The angular distribution of a similar beam which did not pass through the slab is also shown. Although the exact shape of the profile is not yet understood in detail, it is certainly comparable to the profile from the comparison beam.

6 Bus Arbitration

The discussion in the previous sections indicate that a single slab waveguide has the potential to provide over a million independent channels (corresponding to 1000 out-of-plane modes and 1000 colors used in wavelength division multiplexing). Moreover, each of these channels could support over 1000 loads. Although we do not expect such an enormous capacity to be utilized on a single bus, at least initially, a slab waveguide can be used to implement multiple bus systems with much larger loads than previously possible. This is comparable to the high capacity of a free space optical system, but with the ability to broadcast each light source to many physically separated locations.




Fig. 10. Angular distribution of reference beam (top trace) and beam through slab (bottom trace). The full scale is 8 mrad

As an example of what slab waveguides can offer in the parallel processing context, we take up the topic of bus arbitration. In this section we show how the bus arbitration scheme of [1, 23] can be adapted to exploit the high capacity and loading of slab buses.

Consider a bus with N processors numbered $0, 1, \dots, N - 1$. Without loss of generality, assume that $N = 2^k$, for some integer k . For any integer b (where $2 \leq b \leq k$), the k -bit processor addresses can be expressed in base 2^b , simply by grouping the k address bits into $\lceil \frac{k}{b} \rceil$ groups, each with at most b contiguous bits. For convenience, assume that $\frac{k}{b}$ is an integer. Thus the address of processor i can be written in base 2^b as $a_{i,1}, a_{i,2}, \dots, a_{i, \frac{k}{b}}$, where for each $1 \leq j \leq \frac{k}{b}$, $a_{i,j}$ is the j^{th} b -bit digit in the base 2^b representation of the address of processor i .

The aim is to determine the highest indexed processor that wants to write on the bus. Suppose we have $2^b - 1$ (one bit) buses available for the arbitration process. Let these buses be indexed $1, 2, \dots, 2^b - 1$. We will also assume that each processor has word size w and can, in constant time, determine the leftmost (or rightmost) one, if any, in a w -bit word. (This operation can be built into the processor at a cost comparable to that of a w -bit adder.) For our purpose, $b \geq w$.

At the first step, each processor wanting to write on the bus participates in the arbitration algorithm. Let processor i be one such processor. If $a_{i,1} > 0$, then this processor indicates its presence by writing a 1 on bus $a_{i,1}$. Several processors may be writing on this bus, but they will all be writing a 1. Next, each processor i that wants to access the bus reads from all $2^b - 1$ arbitration buses, and determines if there are any positions larger than $a_{i,1}$ for which a processor wants the bus. That is, is there a higher priority processor vying for the bus? If so, processor i withdraws from the rest of the arbitration process. If $a_{i,1} = 0$, then processor i checks to see if any processor has written on any

arbitration bus. Processors that find no higher priority processor vying for the bus continue to subsequent iteration. This step requires, $O(\frac{2^b}{w})$ time.

By applying this method to all $\frac{k}{b} = \frac{\log N}{b}$ digits of the addresses, the highest indexed processor attempting to write on the bus can be selected.

The entire procedure runs in $O\left(\frac{2^b \log N}{wb}\right)$ steps. If $w = \Theta(2^b)$, then the time needed is $O\left(\frac{\log N}{\log w}\right)$. This is an improvement over the conventional $O(\log N)$ time priority resolution (that proceeds one address bit at a time) and is achieved without extra hardware resources. The large number of channels available on a single physical slab admits a large value of b , and hence a lower arbitration time.

7 Conclusions

Slab waveguides hold immense potential for implementing high capacity buses for connecting proximate processors. Preliminary results indicate that a single slab waveguide has, in principle, the potential to provide over a million independent channels (corresponding to 1000 out-of-plane modes and 1000 colors used in wavelength division multiplexing). Moreover, each of these channels can support around 1000 loads. This is comparable to the large capacity of a free space optical system, but with the ability to broadcast each light source to many physically separated locations.

More modest systems, with smaller waveguides and less complex optics, are quite practical with existing technology and still would provide enormous advantages in bused interconnections. The general question of how best to exploit features of slab buses in a computing environment is open.

Acknowledgments

The authors wish to thank Diedre Reese and Yih Chyi Chong for assisting in the experimental work, and James Breedlove for constructing the optical slabs.

References

1. Alnuweiri, H.: A Fast Reconfigurable Network for Graph Connectivity and Transitive Closure. *Parallel Processing Letters* **4** (1994) 105–115
2. Bokka, V., Gurla, H., Olariu, S., Schwing, L., Wilson, L.: Time-Optimal Domain Specific Querying on Enhanced Meshes. *IEEE Trans. Parallel & Distributed Systems* **8** (1997) 13–24
3. Chen, W.-T., Shue, J.-P.: Performance Analysis of a multiple Bus Interconnection Network with Hierarchical Requesting Model. *IEEE Trans. Computers* **40** (1991) 834–842
4. Chen, Y., Chen, W., Chen, G., Sheu, J.: Designing Efficient Parallel Algorithms on Mesh Connected Computers with Multiple Broadcasting. *IEEE Trans. Parallel & Distributed Systems* **1** (1990) 241–246
5. Chiarulli, D.M., Levitan S., Melhem, R.G.: Optical Bus Control for Distributed Multiprocessors. *J. Parallel Distributed Computing* **10** (1990) 45–54

6. Dharmasena, H.P., Vaidyanathan, R.: An Optimal Multiple Bus Network for Fan-in Algorithms. Proc. Int. Conf. on Parallel Processing (1997) 100–103
7. Dighe, O.M., Vaidyanathan R., Zheng, S.Q.: The Bus-Connected Ringed Tree: A Versatile Interconnection Network. J. Parallel Distributed Computing **33** (1996) 189–196
8. Fernández-Zepeda, J.A., Vaidyanathan R., Trahan, J.L.: Scaling Simulation of the Fusing-Restricted Reconfigurable Mesh. IEEE Trans. Parallel & Distributed Systems **9** (1998) 861–871
9. Gourlay, J., Yang, T.-Y., Dines, J., Walker, A.: Development of Free-Space Digital Optics in Computing. IEEE Computer **31** (Feb. 1998) 38–44
10. Guo, Z., Melhem, R.G.: Embedding Binary X-Trees and Pyramids in Processor Arrays with Spanning Buses. IEEE Trans. Parallel & Distributed Systems **5** (1994) 664–672
11. Jiang, H., Smith, K.C.: PPMB: A Partial-Multiple-Bus Multiprocessor Architecture with Improved Cost Effectiveness. IEEE Trans. Computers **41** (1992) 361–366
12. Jang, J.-w., Nigam, M., Prasanna, V.K., Sahni, S.: Constant Time Algorithms for Computational Geometry on the Reconfigurable Mesh. IEEE Trans. Parallel & Distributed Systems **8** (1997) 1–12
13. Kamath, S.T., Vaidyanathan, R.: Running Weak Hypercube Algorithms on Multiple Bus Networks. Proc. ISCA Int. Conf. on Parallel & Distributed Computing Systems (1997) 217–222
14. Karim, Md.N.: Design and Analysis of Reduced Connection Multiple Bus Systems: A Probabilistic Approach. Ph.D. dissertation, Louisiana State University, Baton Rouge, 1996
15. Kulasinghe, P., El-Amawy, A.: On the Complexity of Bussed Interconnections. IEEE Trans. Computers **44** (1995) 1248–1251
16. Kulasinghe, P., El-Amawy, A.: Optimal Realizations of Sets of Interconnection Functions on Synchronous Multiple Bus Systems. IEEE Trans. Computers **45** (1996) 964–969
17. Miller, R., Prasanna-Kumar, V.K., Reisis, D., Stout, Q.: Parallel Computations on Reconfigurable Meshes. IEEE Trans. Computers **42** (1993) 678–692
18. Mitkas, P.A., Betzos G.A., Irakliotis, L.J.: Optical Processing Paradigms for Electronic Computers. IEEE Computer **31** (Feb. 1998) 45–51
19. Nakano, K.: A Bibliography of Published Papers on Dynamically Reconfigurable Architectures. Parallel Processing Letters **5** (1995) 111–124
20. Qiao, C., Melhem, R.G.: Time-Division Optical Communications in Multiprocessor Arrays. IEEE Trans. Computers **42** (1993) 577–590
21. Raksapatcharawong, M., Pinkston, T.M.: An Optical Interconnect Model for k -ary n -cube Wormhole Networks. Proc. Int. Parallel Processing Symp. (1996) 666–672
22. Serrano, M.J., Parhami, B.: Optimal Architectures and Algorithms for Mesh-Connected Parallel Computers with Separable Row/Column Buses. IEEE Trans. Parallel & Distributed Systems **4** (1993) 1073–1080
23. Trahan, J.L., Vaidyanathan R., Subbaraman, C.P.: Constant Time Graph Algorithms on the Reconfigurable Multiple Bus Machine. J. Parallel Distributed Computing **46** (1997) 1–14
24. Wilkinson, B.: On Crossbar Switch and Multiple Bus Interconnection Networks with Overlapping Connectivity. IEEE Trans. Computers **41** (1992) 738–746
25. Wilson, J., Hawkes, J.F.B.: Optoelectronics. Prentice Hall, 1989.
26. Yang, Q., Bhuyan, L.N.: Analysis of Packet-Switched Multiple-Bus Multiprocessor Systems. IEEE Trans. Computers **40** (1991) 352–357