An Extended Fiber-Optic Backplane for Multiprocessors

A. V. Ramanan, H. F. Jordan and J. R. Sauer

Department of Electrical and Computer Engineering
University of Colorado at Boulder
Boulder, CO 80309-0525

D.J. Blumenthal

School of Electrical Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0250

Abstract

This paper describes a parallel by wavelength, packet-switched, extended photonic backplane for multiprocessors. Packet-switching is implemented with minimal optoelectronic conversion at switching nodes, with the payload data maintained in optical format. Since, currently it is not possible to provide efficient asynchronous storage in optics, the deflection routing protocol is used to avoid static buffering at nodes. Space-time switching at nodes is utilized to reduce contention for spatial channels. The ShuffleNet topology yields packet flight times that grow logarithmically with size of the system. With current technology, our approach can support up to 100 bits in parallel at network cycle times of the order of 4ns. The architecture is capable of supporting hundreds to thousands of processors separated by distances up to a few hundred meters.

1 Introduction

Multiprocessing using hundreds to possibly thousands of high performance processors has been recognized as a route to teraflop computing. The barriers to this approach are several, and certainly include system and application software, heat removal from the still large-grain processors, latency tolerance between components in a necessarily distributed system, and a sufficiently flexible and high capacity interconnection network. Though each of these issues is important in itself, they are interlinked. For example, a greatly improved interconnect can ease the requirements in the other areas. Also, processor speeds have increased more than an order of magnitude over interconnect speeds in the past decade, affecting the general-purpose, shared-memory, architectural paradigm to a great extent [10].

The potentials and limitations of optics for communications are complimentary to those of electronics in many respects. The large bandwidth and low loss of optical fiber allows for high throughput using transmission of parallel wavelength channels simultaneously over long distances [1]. Optical interconnects consume less power, are immune to cross talk and interference, and are increasingly capable of multiplexing and switching. The main drawback is the lack of optical static storage. While bandwidth is cheap in fibers, other optical devices such as amplifiers, multiplexers, demultiplexers and switches are expensive resources. Hence, a combination of the mature electronic technology and the potentials offered by optics with the rapid development of new optical devices and components can be exploited to build intelligent interconnects. With the goal of teraflop computing in mind, we have concentrated our efforts towards innovations in optoelectronic interconnect architectures [15]. Our efforts are focussed on various aspects from the node architecture to the network architecture, and finally the system configuration. In this paper, we present the architecture and performance of an extended fiber optic backplane for multiprocessors. The building block switch nodes described here are implementable with available and foreseeable photonic devices and have the potential to outperform any all-electronic system in several respects.

2 Overview of Our Approach

Messages in shared memory multiprocessors are typically word sized. Also, path requirements of messages through any switching point changes frequently. For such single word transfers, the packet switching paradigm and local control provide fast, reduced complexity nodes. Our approach to exploiting the high bandwidth of fibers is to transmit a word in parallel over a single fiber using different wavelength channels , i.e. each bit in a word is transmitted at a different optical wavelength. The control is separated from payload

Proceedings of the Twenty-Seventh Annual Hawaii International Conference on System Sciences, 1994

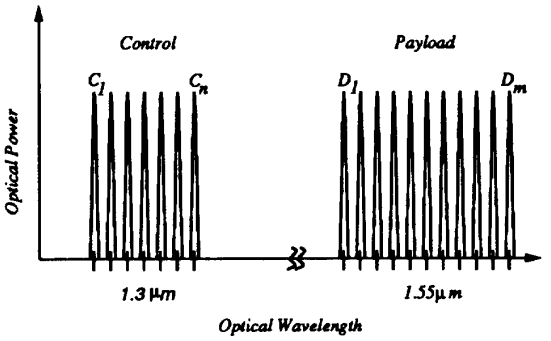

Figure 1: An example of Bit per Wavelength Encoding

by coding payload bits at one group of wavelengths (e.g. about $1.55\,\mu$m) and the control bits at another well-separated group of wavelengths (e.g. $1.3\,\mu$m). An example of the format in the optical wavelength domain is shown in Fig. (1) for n control bits and m payload bits. The cost of going to multiple wavelengths is in the fabrication of laser arrays with multiple sources, with each laser operating at a different stable wavelength. These devices are currently being realized in research labs [7].

In order to utilize the high bandwidth of photonic switches, asynchronous storage must be avoided. Purely optical memory with random access in time is not currently possible. Therefore, it is desirable that a photonic switch have a flow-through architecture so that the data is neither slowed down nor resynchronized by an internal clock. So, the payload data is delayed using a delay line while the routing processor computes the switch settings. The control processor is pipelined so that one control setting can be produced every network cycle [5]. Static buffering is avoided by using the deflection routing (hot-potato) protocol [3] to handle output port contention. Deflection routing minimizes the necessary processing and eliminates the storage overhead at each node by routing a packet out through an alternate output port rather than holding it until the desired output port is available. If deflection routing is not used, optical data must be converted into electronics at nodes for storage and reconverted to optics for transmission. Such electro-optic conversion at every node is expensive and will necessarily slow the network.

The main drawback of deflection routing networks is that packets losing contention take a necessarily longer path to their destination. The effect of this penalty for deflection tends to be severe as internode distance is increased. Our approach to minimize the

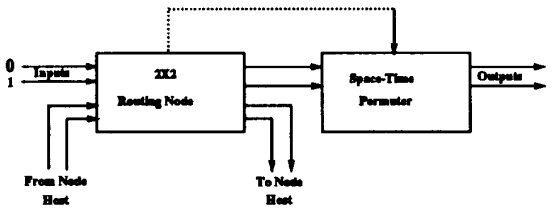

Figure 2: A Network Node

negative effects of deflection routing is to reorder packets in time, dynamically during flight, so as to reduce the contention for spatial channels [14]. We use limited temporal reordering of packets using 2-Space channel, 2-Time channel (2S2T) switch nodes in order maintain fast packet-switching.

Deflection routing requires a topology which provides a path from every node to all other nodes in the system. Because of the multiple hops packets make in the network, the available insertion bandwidth is smaller than that suggested by the network link bandwidth. We have chosen the ShuffleNet topology [8] which yields a packet flight time that grows slowly with size of the system. The topology of deflection networks make them suitable for distributed shared memory multiprocessors [9, 2]. Distributed shared memory systems have memory associated with each processor. Each processor's local memory is either wholly or partially shared among all processors. The processors are connected by means of some interconnection network such that all nodes can be reached from all other nodes. A variant of this architecture is one in which processors and shared memory modules are separate, but are distributed throughout the network. We consider a shared memory system in which each network node is attached to a processor-memory pair. The program and private data are local to a processor and are accessed through a dedicated connection between the processor and it's local memory. Only shared memory is accessed through the network. Hence, the insertion bandwidth needed by the processors is smaller than the network link bandwidth, and can be matched by the insertion bandwidth provided by our network.

3 The Node

Figure (2) illustrates a switch node in our system. The 2×2 routing node is a basic spatial node with additional capability to compute the control settings for the space-time permuter. In this configuration, the

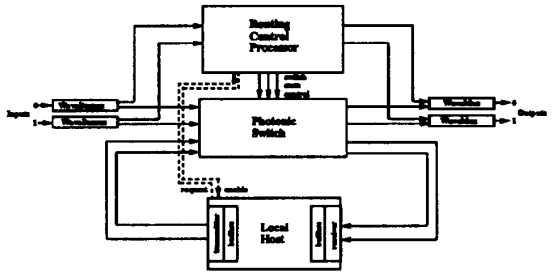

Figure 3: Block diagram of a switching node showing its functionality

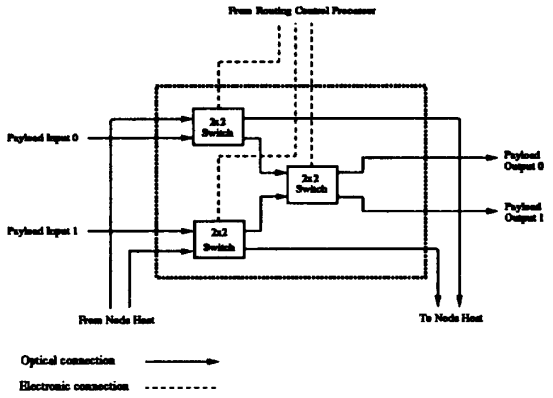

Figure 4: The Photonic Switch

2×2 deflection routing node will perform contention resolution between packets arriving simultaneously, enforcing the priority scheme built into it. The higher priority packet will be routed to it's preferred channel. In case of a deflection which is not repaired by the space-time permuter, the higher priority packet will still remain undeflected.

3.1 The Basic Node

An individual switch node is designed to permute a set of optical network inputs plus a local host output to a set of optical network outputs plus a local host input as illustrated in Fig. (3). Each node consists of a photonic switch, a routing control processor (RCP), and a local host [6]. Solid lines indicate fiber optic connections and dashed lines indicate electronic connections. The photonic switch consists of three 2x2 switches, as shown in Fig. (4), in order to provide the full connectivity between the network and host. Incoming data can be routed to the node host via one of the first two switches, or routed to one of the two output-ports via the final 2x2 switch. Blocking states within the switch are arbitrated by the control processor in real time. Switching elements must route wideband optical data at relatively high reconfiguration speeds. Typical switch bandwidths should be compatible with optical amplifier bandwidths (e.g. 25-100 nm). Reconfiguration rates should be on the order of 1 ns. Switch crosstalk must be kept to a minimum since data is amplified at each node output and may travel through several nodes before reaching its destination.

Control of the switch is handled by a parallel pipelined processor which operates on optical control information extracted from each input. The individual control bits are separated at the control processor using a demultiplexer capable of resolving the individual wavelengths. Since photonic switches can have very high bandwidths, the complete parallel payload portion can be switched to one of the two output ports with one device. Optical data signal levels are restored by optical amplifiers at the node outputs. A single optical amplifier is used for each payload. The optical control information is regenerated by the control processor and reinserted on the network at the node outputs. The maximum processing rate or throughput is limited by the slowest processing element in the pipeline. Pipelining of the routing requests allows packets to be served as they enter the node without elastic buffering in the data path. The control processor has simultaneous access to requests from both network inputs and both node host output queues. An example lookup table type RCP is shown in Fig. (5). The processor arbitrates access to the network from the node host by disabling one or both of the host output queues when full packets are present at the node inputs. The processor also arbitrates contention for the two outputs to the network. The control processor is responsible for detecting and generating both electronic and optical control signals. We use electronically controlled photonic switches since they are readily available. Electronic signals are used to set the switch state and regenerated optical control signals are passively combined with the output data to form new packets which are transmitted to the next node. Thus only the minimal routing information is electro-optically regenerated while the full word-wide optical payload is routed through the switch while remaining in optical format. A complete treatise on the multiwavelength node design can be found in [4].

3.2 The Space-Time Permuter

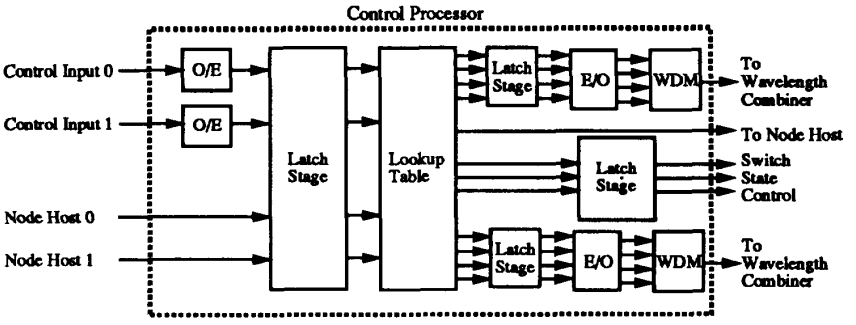

Figure 5: Block diagram of a parallel pipelined switch control processor

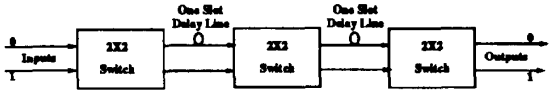

Figure 6: The 2S2T Space-Time Permuter

The space-time permuter shown in Fig. (2) is a 2S2T time-slot interchanger (TSI) which reroutes a deflected packet through it's preferred output channel, if possible. Figure (6) illustrates the details of the space-time permuter. During every network cycle it considers 4 packet slots, 2 slots that have just arrived at it on either spatial channel and slots that are immediately ahead of them, and routes them to the node's output channels after permuting them as necessary. The controls for the permuter are computed by the routing control processor which knows the routing requirements of the packets in the four slots. In this configuration, the TSI needs to perform only 3 permutations out of the 24 permutations possible in an universal 2S2T TSI. The first permutation is an identity permutation, while the other two require swapping of leading packet slot in one of the two spatial channels with the trailing packet slot in the other spatial channel. This organization of the space-time permuter considers a moving time window allowing a deflection to be repaired either by the pair of packet slots preceding it or by those following it.

The nodes described in the previous section are connected to form a ShuffleNet. The ShuffleNet is a logarithmic interconnect topology, which consists of k columns of 2^k nodes with nodes in each column connected to those in the next by a perfect shuffle connection. This topology provides near minimum distance between any source destination pair. The wrap around nature provides a path from every node to every other node in the network and is thus suitable for implementing the deflection routing protocol.

Figure (7) shows a 3 column 24 node ShuffleNet. A host is attached to each switch node by means of two bidirectional ports. Potentially, a host can inject and/or receive two packets into/from the switch node every network cycle. A newly generated packet can enter the network, if either one of the incoming links is free or if one of the incoming packets is for the host. Until then the packet is buffered in the electronics at the the input to the network. It is assumed that enough resources are available to send a packet that has reached its destination node immediately into the host. Hence, the latency of packets has two distinct components: one is the wait time in the input buffer to the node, and the other is the flight time through the network.

The nodes as well as the links are pipelined so that internode transit times will be an integral multiple of network cycle time, and a number of packets can be in transit between two nodes. If all node-link pairs are of constant length, the latency of packets in flight through the network will be proportional to the the number of hops made, i.e. node-link pairs traversed, by the packet and the flight latency distribution of packets as a function of internode transit time will be independent of internode link length. Whenever there is a conflict for an outgoing link in a deflection routing network, one of the contending packets is deflected. The response of a network at any given time is a dynamic function of a number of factors including the topology, size, routing protocol and the nature of the load on the system. The effect of deflections is two-fold: Each deflection will cause a packet to go around the network one more time thereby increasing the latency by the time required for k hops. Also, deflected packets utilize network resources that would

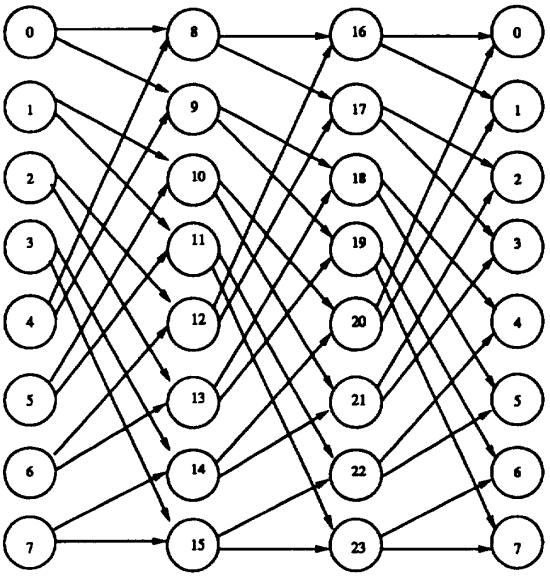

Figure 7: A 24 node ShuffleNet interconnect

otherwise be available for new packets waiting to enter the network. As a result, the network access bandwidth available to the hosts will decrease and packets will have to wait longer in the input buffer.

3.3 The Network Model

We have formulated a simple model of the ShuffleNet [11] that makes use of the specific property that in a ShuffleNet packets require correct routing at the last $min(l,k)$ hops of its flight through the network. Here, l is the number of node-link pairs in the minimum distance path between the current node and the destination of the packet.

Consider an ideal abstraction of the network in which packets encounter no contention and hence no deflection. Packets inside such a network will always reach their destination in minimum time. If packets are generated at constant rate with uniform address distribution, then the average flight latency, in hops, is given by

$$\langle E_0 \rangle = \left(\frac{N}{N-1}\right)\left(\frac{3(k-1)}{2} + \frac{1}{2^k}\right) \qquad (1)$$

The maximum number of node-link pairs traversed is $2k-1$.

Another quantity of interest is the average number of nodes $\langle C_0 \rangle$ at which the packets will care about the output port, i.e. require a particular output port. Packets which can reach their destination in $i <= k$ hops will care about their output port assignment at every node they visit, while packets that require more than k hops will care about their output port assignment only during the last k hops. Thus,

$$\langle C_0 \rangle = \frac{((k^2-2)2^k + k + 2)}{(N-1)} \qquad (2)$$

For large N, $\langle E_0 \rangle$ tends to $3k/2$, while $\langle C_0 \rangle$ tends to k. So, as the size of the ShuffleNet grows, about 2/3 of packets entering a node will be care packets.

A N node network has $2N$ internode links. Hence, on an average 2 packet slots are available for occupancy by packets from each host. Since a packet stays in the network for $\langle E_0 \rangle$ hops, the average network access bandwidth available to each host γ_0 is given by

$$\gamma_0 = \frac{2}{\langle E_0 \rangle} \qquad (3)$$

The system will be able attain a steady state as long as the average injection rate of the hosts is less than this limit.

Consider a real network in which the hosts generate uniform requests at constant rate. Further, we assume contention resolution at nodes to be random, i.e. packets have no priorities attached to them. Under these assumptions, the probability of deflection of a packet p_d will be independent of the state of the packet and will be the same at all nodes. The average number of hops made by the packets turns out to be

$$\langle E \rangle = \langle E_0 \rangle + \left(\frac{N+1}{N-1}\right)\frac{k(1-(1-p_d)^k)}{(1-p_d)^k} \qquad (4)$$

Thus, the average number of node-link pairs traversed by the packets is the sum of the minimum average expected due to the topology and the overhead incurred due to deflections. As the probability of deflection approaches zero, the contribution of the second term vanishes. As will be seen later, the probability of deflection is reduced drastically in our space-time node as opposed to a purely spatial node.

Similarly, $\langle C \rangle$, the number of nodes at which an average message cares about the output port assigned by the router is found to be

$$\langle C \rangle = \left(\frac{N+1}{N-1}\right)\frac{(1-(1-p_d)^k)}{p_d(1-p_d)^k} - \frac{(2^{k+1}(1-p_d)^k - 2)}{(N-1)(1-p_d)^k(1-2p_d)}. \qquad (5)$$

As p_d tends to zero, $\langle C \rangle$ tends to $\langle C_0 \rangle$.

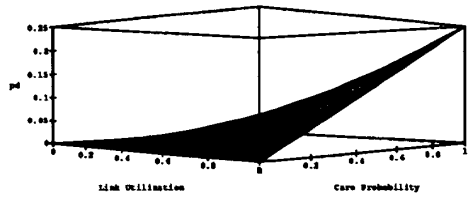

Figure 8: Deflection Probability in a Spatial Node

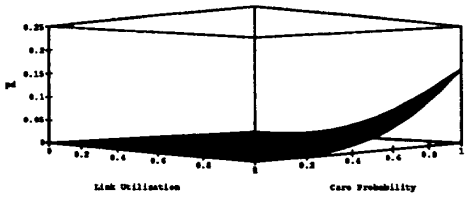

Figure 9: Deflection Probability in a Space-Time Node

The number of packet slots occupied at any given time is dependent on the injection rate of the hosts and the flight latency of the packets. In the steady state, the average throughput of the nodes will be equal to the average injection rate of the hosts. Hence, the link utilization α, the throughput or offered load γ and flight latency $\langle E \rangle$ are related as

$$\alpha = \frac{\gamma}{2}\langle E \rangle \tag{6}$$

β, the care probability, is

$$\beta = \frac{\langle C \rangle}{\langle E \rangle} \tag{7}$$

The probability of deflection of a packet, p_d is the conditional probability that a packet cares about its output port and gets deflected. The probability that a packet will ultimately get deflected is given by

$$p_d = \frac{p_{cont} p_r p_r^{'}}{2} \tag{8}$$

where p_{cont} is the probability that the other packet also requires the same output port and the probabilities p_r and $p_r^{'}$ are the probabilities that the preceding and following packet slot combinations do not resolve the contention [12]. The last two probabilities are given by

$$\begin{aligned} p_r &= \frac{\alpha\beta(1-\alpha\beta/4)}{1-(\alpha^2\beta^2/4)(1-\alpha\beta/2)^2} \\ p_r^{'} &= \alpha\beta(1-\alpha\beta/4) \end{aligned}$$

The probability of contention p_{cont} is given by

$$p_{cont} = \frac{\alpha\beta}{2} \tag{9}$$

Hence, the probability of deflection p_d is

$$p_d = \frac{\alpha^3\beta^3}{4} \cdot \frac{(1-\alpha\beta/4)^2}{1-(\alpha^2\beta^2/4)(1-\alpha\beta/2)^2} \tag{10}$$

while, it is $\alpha\beta/4$ for a purely spatial node. Eqn. (10) shows that at full link utilization, the probability of deflection in a single minimum distance path network falls from 0.25 in a spatial switch to 0.15 in a space-time switch. The decrease in probability of deflection is much more in multiple minimum distance path networks. Figures (8) and (9) show the variation of probability of deflection with link utilization α and care probability β, in a spatial node and space-time node respectively obtained through the preceding analysis.

Equation (4) for flight Latency $\langle E \rangle$, Eqn. (6) for link utilization α and Eqn. (10) for probability of deflection p_d when solved together with Eqn. (5) gives the flight latency as a function of offered load γ. A complete derivation of the equations presented in this section can be found in [11], [12], and [13].

3.4 The Network Behavior

In this section we present analytical as well as simulation results on the performance of the network. The model presented above has been verified both for spatial network and space-time network using simulation for network sizes from 64 to 2048 nodes and extended to predict the performance up to 10240 nodes. The model assumes that hosts inject traffic that is uniformly distributed in space at a constant rate. Even at highest possible link utilization, the space-time technique exhibits a better behavior as shown in Figs. (10) and (11). Though the flight latency in a space-time network is higher than the ideal, the disparity does not increase with network size. The reduction in network latency leads to consistent improvement in user throughput. Simulations show that the wait buffer latency is negligible in a uniformly loaded system up to moderately high link utilization. The flight latency component increases sharply with load in a network with purely spatial switch. As expected, the extremely low probability of deflection in a space-time switch keeps the number of deflections down so that, the flight latency is only slightly above the ideal with no deflections in the network as shown in Fig. (12). Since the transit time of packets is less, links are less sparsely populated than in a spatial network under

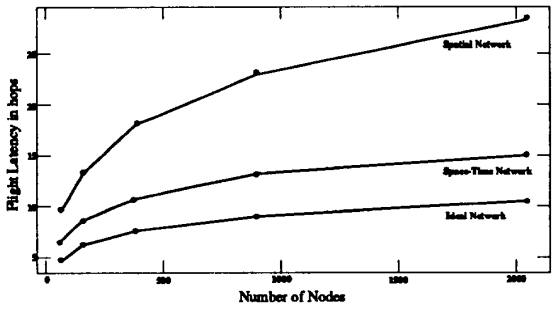

Figure 10: Variation of Maximum Flight Latency with Size

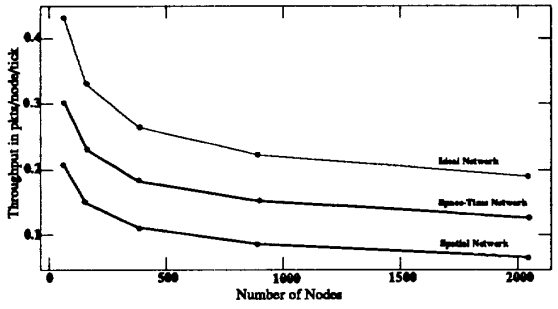

Figure 11: Variation of Maximum User Throughput with Size

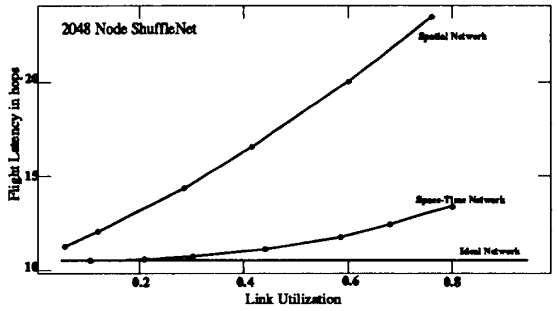

Figure 12: Flight Latency vs. Link Utilization for a 2048-Node ShuffleNet

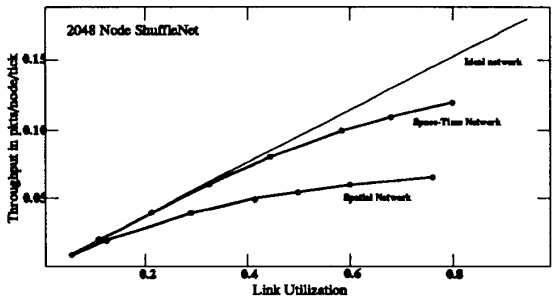

Figure 13: Throughput vs. Link Utilization for a 2048-Node ShuffleNet

same load. As a result the bandwidth available to the hosts increases, as shown in Fig. (13). No error bars are shown on any measured quantity, in this curve as well as all others, since the errors computed at 95% confidence interval were too small to be indicated.

An important improvement in the network behavior due to time-domain switching is in the flight latency distribution of packets. In a deflection routing network, the tail of the delay distribution tends to be long if random contention resolution is used at the nodes. Any priority scheme helps to shorten the tail of the distribution. We use the age priority scheme, which provides preference for the colliding packet that has been deflected more. Simulations show that though the tail is shortened, the average behavior of the network does not change significantly. Figures (14) and (15) show the latency distribution histogram for a 2048-node ShuffleNet with spatial nodes and space-time nodes under the same injection rate and the age-priority scheme. In each case, the ideal distribution, under no deflections, has also been plotted for comparison. Apart from reducing the average latency, a space-time node also shortens the tail of the latency distribution, i.e. it reduces the variations among the latency of packets.

The network is currently being studied under loads that are representative of shared memory multiprocessors. A study under temporal and spatial non-uniformities in traffic shows that the network adapts well to temporal bursts as well as to hot-spots [13].

4 Summary and Conclusions

We have presented a multi-hop, fiber optic, packet-switched direct network that can form an extended backplane for multiprocessors. It utilizes both space and time rearrangements, at nodes, using a deflection

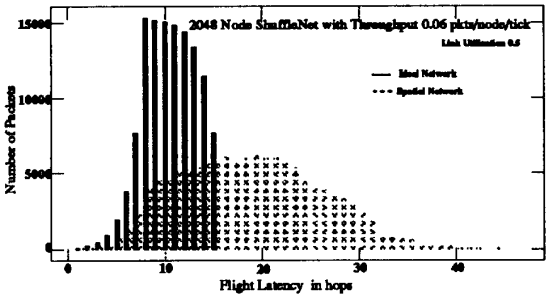

Figure 14: Flight Latency Histogram for a 2048-Node ShuffleNet with Spatial Nodes

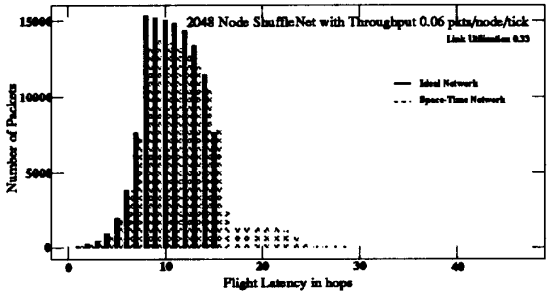

Figure 15: Flight Latency Histogram for a 2048-Node ShuffleNet with Space-Time Nodes

routing protocol. Simultaneous space and time switching make the network nearly as efficient, with reasonable network parameters, as an ideal network without contention. Limited temporal reordering greatly reduces the contention for output spatial channels at the switching points. Since the required switching logic is simple and pipelined, the network links can be very fast. Space-time switching involves a small additional delay at switch nodes. This overhead is offset by the decrease in latency due to reduced number of deflections even at short internode distances, while the advantage of increased throughput is maintained.

The network does not require optoelectronic conversion at nodes for asynchronous storage. Since, the control processor logic is simple and fully pipelined, the network cycle time is limited only by the packet size. With the current technology, we hope to achieve 4ns packets, i.e. a link rate of 250Mpkts/sec. For, internode separations up to a kilometer, the network architecture can support up to 100 bits in parallel. For longer internode distances, it will take longer to service a memory request. The processors have to be multithreaded, so that they can be kept busy while waiting for the requests of suspended threads to be satisfied. Hence, the main consideration in choosing the internode distance is the total number of threads that need to be supported to keep the processor busy, while the network is servicing shared memory access requests. If the processors and memories are fully pipelined at the same rate as the network, a network with 2048 nodes needs about 15 contexts at an internode distance of 5m and about 150 at a distance of 50m. These numbers are based on the assumption that a thread will require a remote memory access once in about 20 processor cycles. Thus, with proper choice of system parameters and with adequate support from software, the network we have presented can provide an efficient extended backplane for multiprocessors.

Acknowledgements

This work is supported in part by the US Army grant number DASG 60-91-C-0143, by the AFOSR grant number 91-0392, by the NSF Engineering Research Center grant number ECD 9015128 and by the Colorado Advanced Technology Institute.

References

[1] Special issue on dense wavelength division multiplexing techniques for high capacity and multiaccess communication systems. *IEEE Journal on Selected Areas in Communications*, 8, Aug 1990.

[2] R. Alverson, D. Callahan, D. Cummings, B. Koblenz, A. Porterfield, and B. Smith. The Tera Computer System. In *International Conference on Supercomputing*, pages 1–6, June 1990.

[3] P. Baran. On Distributed Communication Networks. *IEEE Transactions on Communication Systems*, 12:1–9, 1964.

[4] D. J. Blumenthal. *Multiwavelength Photonic Packet Switched Interconnects.* PhD thesis, University of Colorado at Boulder, Department of Electrical and Computer Engineering, 1993.

[5] D. J. Blumenthal, K. Y. Chen, J. Ma, R. J. Feuerstein, and J. R. Sauer. Demonstration of a Deflection Routing 2x2 Photonic Switch for Computer Interconnects. *IEEE Photonics Technology Letters*, 4(2):169–173, 1992.

[6] D.J. Blumenthal and J.R. Sauer. Multiwavelength information processing in Gigabit photonic switching networks. In *Proceedings of the SPIE, Conference on Multigigabit Fiber Communications*, volume 1787, Boston, MA, September 1992.

[7] C. J. Chang-Hasnain, M. W. Maeda, J. P. Harbison, L. T. Florez, and C. Lin. Monolithic Multiple Wavelength Surface Emitting Laser Arrays. *IEEE Journal of Lightwave Technology*, pages 1655–1673, Dec 1991.

[8] M. G. Hluchyj and M. J. Karol. ShuffleNet: An application of Generalized Perfect Shuffles to Multihop Lightwave Networks. In *Proceedings of IEEE INFOCOM '88*, pages 379–390. IEEE Computer Society Press, Mar 1988.

[9] J. S. Kowalik, editor. *Parallel MIMD Computation: HEP Supercomputer and Its Applications.* Cambridge, MA: MIT Press, 1985.

[10] E. P. Markatos and T. J. LeBlanc. Shared-Memory Multiprocessor Trends and the Implications for Parallel Program Performance. OCS Technical Report No. 420., Compuer Science Dept., Univ. of Rochester, Rochester, NY 14627, May 1992.

[11] A. V. Ramanan, H. F. Jordan, and J. R. Sauer. Performance of ShuffleNet under Deflection Routing. OCS Technical Report No.93-21, OCS Center, University of Colorado, Boulder, CO 80309-0525, 1993.

[12] A. V. Ramanan, H. F. Jordan, and J. R. Sauer. Space-Time Switching in Deflection Routing Networks. OCS Technical Report No.93-20, OCS Center, University of Colorado, Boulder, CO 80309-0525, 1993.

[13] A. V. Ramanan. *Ultrafast Space-Time Networks for Multiprocessors.* PhD thesis, University of Colorado at Boulder, Department of Electrical and Computer Engineering, 1993.

[14] A. V. Ramanan, H. F. Jordan, and J. R. Sauer. Space-Time Switching in Fiber-Optic Packet Switched Networks. In *Proceedings of the SPIE, Conference on Multigigabit Fiber Communication Systems*, volume 2024, San Diego, CA, July 1993.

[15] J. R. Sauer, D. J. Blumenthal, and A. V. Ramanan. Photonic Interconnects for Multicomputer Communications. *IEEE LTS: Special Issue on Gigabit Networks*, pages 12–19, Aug 1992.