Photonic Slot Routing in All-Optical WDM Mesh Networks*

Hui Zang†, Jason P. Jue, and Biswanath Mukherjee
Department of Computer Science, University of California, Davis, CA 95616
zang@cs.ucdavis.edu, juej@cs.ucdavis.edu, mukherje@cs.ucdavis.edu
†Correspondence Author (Tel: (530) 752-5129; Fax: (530) 752-4767)

Abstract

Photonic slot routing has been proposed as an approach to implement an all-optical packet-switched network in a manner which is scalable and not overly complex. Although photonic slot routing has been studied in bus and ring networks, it has not been applied to general mesh topologies until now. In this study, we investigate some of the issues involved in implementing photonic slot routing in a mesh network. In particular, we propose an approach for assigning slot destinations in a photonic slot routing network, and we evaluate the performance of the proposed technique.

1 Introduction

Wavelength-division multiplexing (WDM) has been rapidly gaining acceptance as a means to handle the ever-increasing bandwidth demands of network users [1]. In addition to providing huge amounts of bandwidth, all-optical WDM networks also offer the benefit of high-speed data transmissions without electronic conversions at intermediate nodes. By transmitting the signal entirely in the optical domain, data transparency can be achieved. One such type of all-optical network is the wavelength-routed WDM network, in which all-optical lightpaths are set up on specific wavelengths between pairs of nodes [2, 3]. Some of the major challenges in designing wavelength-routed WDM networks include the complexity and scalability issues which arise from the need to demultiplex and individually route each wavelength at a node.

An emerging alternative to the wavelength-routed WDM network is a WDM network based on optical packet switching. One recently proposed approach to optical packet switching is Photonic Slot Routing (PSR) [4]. In PSR, time is slotted, and data is transmitted in the form of photonic slots which are fixed in length and span all wavelengths in the network. Each wavelength in the photonic slot may contain a single packet, and all packets in the photonic slot are destined to the same node. By requiring the packets to have the same destination, the photonic slot may be routed as a single integrated unit without the need for demultiplexing individual wavelengths. Thus, wavelength-insensitive components may be used at each node, resulting in less complexity, faster routing, and lower network cost [5].

Although PSR has been studied in bus and ring networks, it has not been considered in mesh networks until now. Also, no analytical model has been developed to measure the performance of a PSR network. In a mesh network, nodes may have multiple input and output ports, resulting in a higher degree of contention than in ring-based networks. In this work, we investigate various approaches for reducing the amount of contention, and for resolving contentions when they occur.

In Section 2, we describe the proposed PSR network architecture for arbitrary mesh networks. The protocols for transmitting packets and for resolving contentions are discussed in Section 3. We present numerical examples in Section 4. Section 5 concludes the paper and discusses areas for future research.

2 Network Architecture

We consider a mesh interconnected network in which each node consists of a wavelength-insensitive optical packet switch, optical delay-line buffers (to hold photonic slots at intermediate nodes), and electronic packet buffers (to hold packet arrivals at each source node). Each node (while acting as a source) maintains a separate packet buffer for each destination node. A diagram of the node architecture is shown in Fig. 1.

At each input fiber link, an optical splitter (or demultiplexer) is used to extract the header of each photonic slot. The header, which may be carried on a separate wavelength, contains information such as the destination of the slot, and which wavelengths in the slot are occupied by packets. The destination information is used to configure the switch setting so that the slot can be appropriately routed towards its destination. The slot occupancy information determines the extra allowable packet transmissions for each outgoing photonic slot. Because it takes some time to process the slot header and to configure the optical packet switch, delay lines may be required on each input fiber link as shown in Fig. 1. Alternatively, the slot header may precede the data payload of the slot by some fixed duration.

On a given output fiber link, the node may insert packets into existing photonic slots which are headed for the same destination or the node may transmit newly created photonic slots if no other slots are contending for the link. The packet insertion may be performed by an optical cou-

---

*This work has been supported in part by NSF Grant Nos. NCR-95-08239 and ANI-98-05285.

0-7803-5796-5/99/$10.00 © 1999 IEEE

pler. For example, in Fig. 1, a photonic slot departing on the top output fiber link contains packets on wavelengths λ_1 and λ_2. The node may then insert a packet into the photonic slot on wavelength λ_3.

In the given architecture, there is the possibility of contention between two or more photonic slots at each of the output fiber links. One approach for resolving contention is to optically buffer the photonic slots which lose a contention. In this case, optical delay line buffers, consisting of lengths of fiber, are required at each node [6]. Alternative approaches for dealing with contention include deflection routing and dropping photonic slots.

One of the challenges in implementing PSR in a mesh environment is maintaining the synchronization of photonic slots. Each of the incoming photonic slots must arrive simultaneously on all input fibers in order to be routed through the optical packet switch simultaneously. We address this problem by including adjustable delay elements at the input ports of each node.

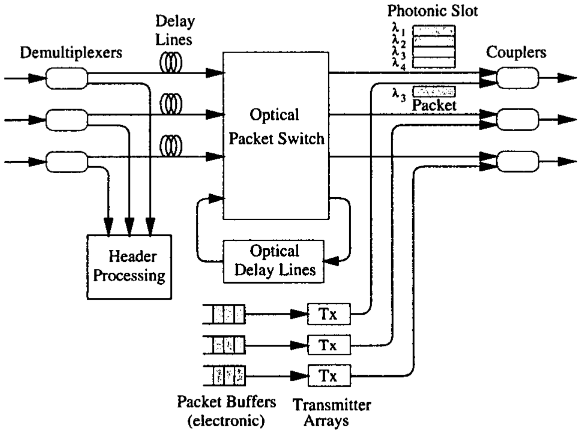

Figure 1: PSR node architecture.

3 PSR Protocols

When a node has a new packet to transmit, it may add the packet to any slot which is headed for the same destination and which is not full. It may also place the packet in an empty slot which does not have a destination assigned to it. In this case, the node must set the slot destination to the destination of its packet.

There are a number of slot-assignment policies for utilizing empty slots and for assigning destinations to these slots. These policies may have an effect on contention at intermediate nodes. We discuss two different policies and their performance in Sections 3.1 and 3.2. There must also be a policy for determining how many packets should be added to a slot. In this work, we will assume a greedy approach to filling slots, i.e., a node will place as many packets as possible into a slot, as long as the packets are headed to the same destination. Alternatively, we may limit the number of packets placed in a slot in order to provide downstream nodes in the network with an opportunity to also transmit packets in the same slot.

Nodes which receive forwarded slots from upstream nodes are referred to as *intermediate nodes.* The problem for an intermediate node is to route non-empty slots to the appropriate output link. Optimistically, the intermediate node should put the slot onto the proper link as determined by the routing algorithm[1]. If two or more slots arrive at the same time and require the same outgoing fiber link, then contention occurs. In this work, we will utilize optical slot buffers to resolve contention. When the buffers become full, slots are dropped.

In the remainder of this section, we study two types of slot-assignment policies. First, we consider a slot-assignment algorithm based on packet arrivals. We then consider a slot-assignment algorithm based on capacity allocation. In both schemes, a node which has packets to transmit can always add its packets to slots which are headed to the same destination node.

3.1 Packet Arrival Based Assignment

In the slot-assignment algorithm based on packet arrivals, no knowledge of the network traffic pattern is required. Upon receiving an empty slot, a node randomly chooses one queue from among its packet queues which are not empty. It then inserts a number of packets into the slot. Thus, if the queues for a link (if fixed routing is used, a packet queue always requires the same outgoing link) are not all empty at the beginning of a time slot, then the slot departing on this link will always have a destination assigned to it and will also be non-empty.

A disadvantage of this algorithm is that it results in a high probability of contention as well as some unfairness in resource allocation. Nodes located towards the edge of the network, which have little or no chance of being intermediate nodes, always receive empty slots in which to transmit their packets. For a (tagged) node located in the internal region of the network, the passing slots may already be occupied by packets from upstream nodes. Hence, the tagged node has relatively fewer chances to transmit its packets than the nodes at the edge of the network.

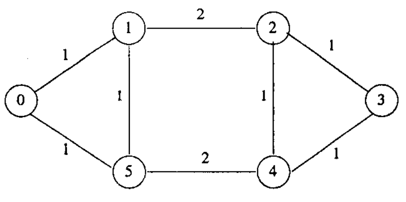

Figure 2: Topology for Network 1.

As an alternative to packet-arrival-based slot allocation, we can assign destinations to the empty slots in such a way that the upstream nodes will never overwhelm the downstream nodes.

---

[1] e.g., shortest-path routing, which is used throughout this study.

3.2 Capacity Allocation

Chlamtac et al. [4] proposed a transmission control protocol based on slot preassignment for photonic slot routing in optical networks with a ring topology. In that approach, a TDMA frame consists of L slots, and the source and destination of each slot in the TDMA frame is determined by a network-wide TDMA schedule. The number of slots in the TDMA frame assigned to each source-destination pair is determined in such a way that fair bandwidth distribution among the source-destination pairs, as well as contention-free slot routing at intermediate nodes, are achieved.

The objectives of our slot-assignment policy are to maximize throughput, to minimize contention, and to provide a fair allocation of bandwidth. It is difficult to achieve these objectives using TDMA frames. Fair allocation of bandwidth may not always be achievable due to the finite length of the TDMA frame and the fact that the number of allocated slots per frame is always an integer. Also, contention may not be fully avoidable because a network-wide TDMA schedule which also simultaneously satisfies the fair allocation requirement may not exist. An alternative to generating TDMA frames in a mesh network is to probabilistically assign destinations to arriving slots based on the capacity allocation results. The slot assignment consists of two steps:

1. the Capacity-Allocation step which determines the fraction a_{jl} of capacity (or slots) on each link l that a source node i should assign to destination j, and
2. the Slot-Assignment step in which a node assigns a destination to a slot based on the capacity allocation and a randomly generated number.

The capacity-allocation algorithm is based on the traffic matrix and the routing matrix of the network. Given traffic matrix T, with t_{ij} denoting the arrival rate to node i of packets destined to node j, and the capacity-allocation matrix C, with c_{ijl} denoting the fraction of slots going out on link l which are allocated to source-destination pair (i, j), c_{ijl} should satisfy:

$$c_{ijl_1} = c_{ijl_2} = c_{ij} > 0, \ c_{ijl_3} = 0. \tag{1}$$

if both link l_1 and link l_2 are on path $i \overset{*}{\to} j$ and link l_3 is not[2]. Equation (1) is called the *flow conservation constraint.*

For a node i, the probability that it assigns destination j to an empty slot going out on link l is given by:

$$Pr_{jl} = \frac{c_{ij}}{\sum_k c_{ik}}. \tag{2}$$

The sum in the denominator is over all nodes k such that path $i \overset{*}{\to} k$ goes through link l. That is, the probability equals the capacity allocated to path $i \overset{*}{\to} j$ divided by the sum of all capacities allocated to the paths starting at link l.

The algorithm to calculate c_{ij} for all sources i and destinations j is as follows:

- **Step 1**. For each link l, let A_l denote the available capacity on link l. Set $A_l = 1$.
- **Step 2**. For each link l, the fraction of capacity available to path $i \overset{*}{\to} j$ which contains link l is

$$\frac{t_{ij}}{\sum_{xy} t_{xy}}, \tag{3}$$

with the sum in the denominator being over all paths containing link l and which have *not* been allocated yet. Calculate the minimum available capacity over all the paths $i \overset{*}{\to} j$ containing link l, and assign this value, denoted by MC_l, to link l as the minimum available capacity to a path on link l that has not been allocated, i.e.,

$$MC_l = \min_{i,j} A_l \times \frac{t_{ij}}{\sum_{xy} t_{xy}}. \tag{4}$$

- **Step 3**. Let $c_{i'j'} = \min_l MC_l$, where $i' \overset{*}{\to} j'$ is the path containing link l' with the minimum MC_l over all links l, and path $i' \overset{*}{\to} j'$'s available capacity is the minimum on link l. Assign $c_{i'j'}$ to path $i' \overset{*}{\to} j'$.

  For all links l'' containing path $i' \overset{*}{\to} j'$, $A_{l''} \longleftarrow A_{l''} - c_{i'j'}$, according to the *flow conservation constraint.*

  If all paths are allocated, stop. Else go to **Step 2.**

In the above algorithm, a path can never be allocated a capacity which causes the capacity allocated on any given link to exceed 1. Also, in this approach, a path consisting of only a single link is allocated the remaining capacity on that link after all of the other paths containing this link have been allocated. In this way, the capacity of each link is fully utilized and the capacity allocated to the one-link path is not limited by the traffic load.

Given the capacity allocation for all paths, we can then calculate the probability that a node i will assign destination j to an empty slot going out on link l by using Equation (2). Upon generating a new empty slot on a link, a node generates a random number uniformly distributed between $(0, 1)$ and decides which destination to assign to that empty slot based on the probabilities from Equation (2). If no packets are available for this destination, the node can either put the empty slot onto the link with the destination unassigned, or with the destination assigned, but with no packets in the slot.

For the network to be stable, the maximum arrival rate to source i of packets for destination j should be less than the capacity allocated for this path, c_{ij}, times the number of packets that a node can fill in a slot. Usually, this threshold is higher for the capacity-allocation approach than for the packet-arrival-based approach.

4 Numerical Examples

A detailed analytical model for a network implementing slot assignment based on capacity allocation was developed, but has been omitted due to space constraints. We consider two networks whose performance we illustrate via our analytical model and simulation. The first (shown in

---

[2] The notation $i \overset{*}{\to} j$ indicates the path from i to j which may contain multiple links, i.e., $i \overset{*}{\to} j \equiv i \to \cdots \to j$.

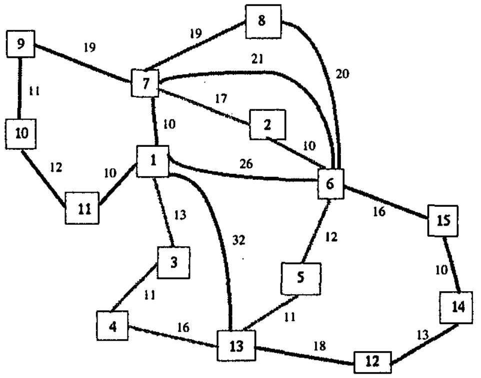

Figure 3: Network 2, a typical telecommunication network consisting of interconnected rings.

Fig. 2) is a small network with 6 nodes and 8 bidirectional links. The second (Fig. 3) is a larger network with 15 nodes and 21 bidirectional links. The lengths of the links range from 10 slots to 32 slots. Packets arrive to each node according to a Poisson process, and the packet rate is assumed to be λ for each source-destination pair.

Following are the default parameters used in the simulation:

- Packet size: 1,000 bytes.
- Bandwidth of a link: OC-48, i.e., 2.5 Gbps.
- The propagation delay on a 1 km link is 1 slot time.
- Number of wavelengths: $W = 4$ or 8.
- Packet buffer size at each node to each destination $B_p = \infty$ (unless otherwise stated).

We consider the following performance metrics:

- Average packet delay.
- Link utilization.
- Fairness index:

$$FI = \frac{(\sum x_{ij})^2}{N \times (N-1) \times \sum x_{ij}^2} \tag{5}$$

where x_{ij} is the throughput for source-destination pair (i, j). If the throughput for all the source-destination pairs are equal, $FI = 1$. Note that, due to finite buffering, the network throughput may be lower than the offered load and the fairness index may be lower than unity.

Figure 4 plots the average packet delay versus link utilization for $W = 4, 8$ in Network 2. The packet buffer size is assumed to be $B_p = 200$, and slot buffers are assumed to be infinite. We observe that the utilization is nearly the same for different number of wavelengths. Note the good agreement between the analysis and the simulation. Although the utilization in Fig. 4 is limited to around 0.5, this performance is not a limitation of photonic slot routing, but rather a limitation of the fixed shortest-path routing. By implementing better routing algorithms which more evenly balance the load across all links in the network, higher utilizations may be achieved.

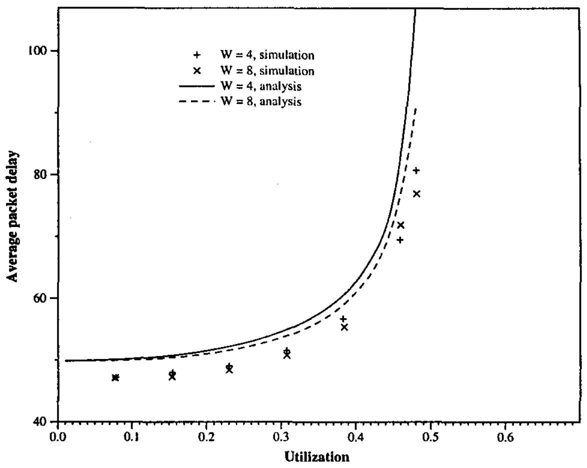

Figure 4: Average delay versus link utilization in Network 2, for varying number of wavelengths ($W = 4, 8$).

Figures 5 and 6 compare the performance of the two slot-assignment schemes. For each scheme, we evaluate the performance of Network 1 with $W = 8$.

Figure 5 plots the fairness index versus λ for the two slot-assignment algorithms. In both schemes, each node has a finite packet buffer of size 200 for each destination, and a finite slot buffer of size 10. Therefore, when the load increases, both packets and slots may be dropped. We observe that, when the load increases, the fairness index decreases faster for the packet-arrival-based slot-assignment scheme than for the capacity-allocation-based slot-assignment scheme. Under the packet-arrival-based scheme, source nodes may allocate slots in a manner which may prevent intermediate nodes from receiving enough slots for a given destination. Also, source nodes may completely fill each slot, leaving very little capacity for downstream nodes. By allocating slots according to the capacity-allocation scheme, each node is guaranteed to have a certain number of empty slots available for each destination, resulting in a higher degree of fairness.

Figure 6 plots the average packet delay versus λ for the two slot-assignment algorithms. Infinite packet buffers and slot buffers are assumed. For low load, the algorithm based on packet arrivals performs better than the algorithm based on capacity allocation. In the packet-arrival-based scheme, an arriving packet is immediately transmitted in the next empty slot. Under low loads, empty slots are fairly plentiful, thus a packet will not have to wait long before being transmitted. On the other hand, in the capacity-allocation scheme, even when there are empty slots, there is a certain probability that a packet will not be transmitted if these slots are assigned to some other destinations. Thus, a packet may have to wait longer for a slot to the correct destination while letting empty slots pass by. Under higher loads, the capacity-allocation scheme offers better performance than the packet-arrival-based scheme. In the packet-arrival-based scheme, since

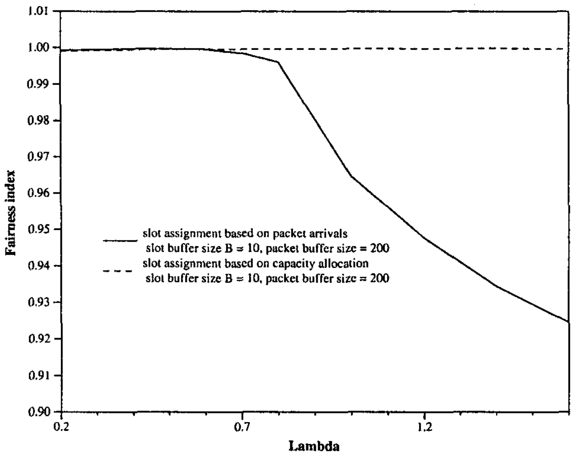

Figure 5: Fairness index versus packet arrival rate (λ) in Network 1, for $W = 8$ and different slot-assignment schemes based on packet arrivals and capacity allocation.

no restrictions are placed on the source nodes, there is the potential for a severe load imbalance. This load imbalance is created when the source nodes generate too many slots for a particular destination, and these slots overload some link in the network. In the capacity-allocation scheme, by limiting the rate at which source nodes generate slots for each destination, the problem of load imbalance is eliminated and a higher potential throughput can be achieved.

Another advantage of the capacity-allocation scheme is that it decreases contention remarkably. With capacity allocation, it is possible to use a small optical buffer for slots which cannot be routed properly, and still maintain a very low slot-drop rate.

5 Conclusion

Photonic slot routing has been proposed as a solution to fast routing in all-optical packet-switched WDM networks. In this approach, packets which have the same destination and which are transmitted on different wavelengths by different source nodes may form a photonic slot which is routed through the network as an integrated unit. This study demonstrated how the concept of photonic slot routing can be applied to a mesh network, introduced an analytical model to evaluate the performance of such networks, and discussed different slot-assignment algorithms, contention-resolution schemes, and slot-filling policies. It was found that the slot-assignment algorithm based on capacity allocation can lead to better performance in terms of average packet delay, network throughput, and fairness in the allocation of network resources. Some possible areas for future research include the development of analytical models for networks in which deflection is used, the study of capacity allocation algorithms which result in a higher degree of fairness, and the effect of physical constraints such as slot dilation.

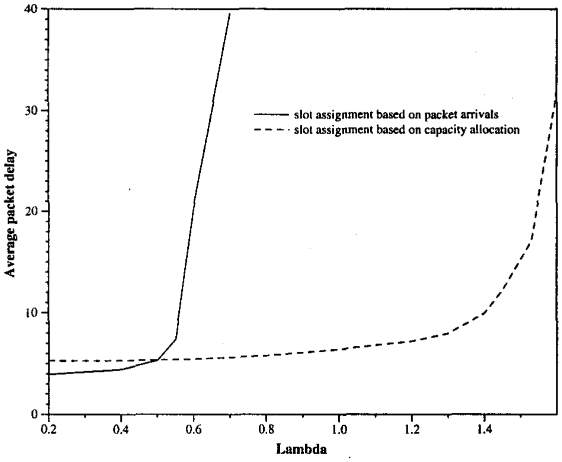

Figure 6: Average delay versus packet arrival rate (λ) in Network 1, for $W = 8$ and different slot-assignment schemes based on packet arrivals and capacity allocation.

References

[1] B. Mukherjee, *Optical Communication Networks*, McGraw-Hill, New York, 1997.

[2] I. Chlamtac, A. Farago, and T. Zhang, "Lightpath (wavelength) routing in large WDM networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, pp. 909-913, June 1996.

[3] B. Mukherjee, D. Banerjee, S. Ramamurthy and A. Mukherjee, "Some principles for designing a wide-area WDM optical network," *IEEE/ACM Transactions on Networking*, vol. 4, no. 5, pp. 684-696, Oct. 1996.

[4] I. Chlamtac, V. Elek, and A. Fumagalli, "A fair slot routing solution for scalability in all-optical packet-switched networks," *Journal of High Speed Networks*, vol. 6, pp. 181-196, 1997.

[5] I. Chlamtac, V. Elek, A. Fumagalli, and C. Szabo, "Scalable WDM network architecture based on photonic slot routing and switched delay lines," *Proceedings, IEEE INFOCOM '97*, Kobe, Japan, vol. 2, pp. 769-776, April 1997.

[6] I. Chlamtac, A. Fumagalli, L. G. Kazovsky, et al., "CORD: contention resolution by delay lines," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, pp. 1014-1029, June 1996.

[7] C. Qiao and R. Melhem, "Reducing Communication Latency with Path Multiplexing in Optically Interconnected Multiprocessor Systems," *IEEE Transactions on Parallel and Distributed Systems*, vol. 8, no. 2, pp. 97-108, Feb. 1997.

[8] C.-C. Yu and S. Bhattacharya, "Dynamic Scheduling of Real-Time Messages Over an Optical Network," *Proceedings, Sixth IEEE International Conference on Computer Communications and Networks*, Las Vegas, NV, pp. 336-339, Sept. 1997.