Nationaal Lucht- en Ruimtevaartlaboratorium
National Aerospace Laboratory NLR

NLR-TP-98225

A novel approach for all-optical packet-switching in wide-area networks

G. Wedzinga, I. Chlamtac and A. Fumagalli

NLR-TP-98225

A novel approach for all-optical packet-switching in wide-area networks

G. Wedzinga, I. Chlamtac* and A. Fumagalli*

* *University of Texas at Dallas*

This report is based on a presentation held on the SYBEN '98 International Symposium on Broadband European Networks, Zürich, Switzerland, 18-20 May 1998.

Division:	Electronics and Instrumentation
Issued:	August 1998
Classification of title:	unclassified

A novel approach for all-optical packet-switching in wide-area networks

Imrich Chlamtac[a], Andrea Fumagalli[a], and Gosse Wedzinga[b]

[a]University of Texas at Dallas, E. Jonsson School of Engineering and Computer Science,
P.O. Box 830688, Richardson, TX 75083-0688, USA

[b]National Aerospace Laboratory NLR, Avionics Department,
P.O. Box 90502, 1006 BM Amsterdam, The Netherlands

ABSTRACT

All-optical Wavelength Division Multiplexing (WDM) networks are believed to be a fundamental component in future high speed backbones. However, while wavelength routing made circuit switching in WDM feasible the reality of extant optical technology does not yet provide the necessary devices to achieve individual optical packet switching.

This paper proposes to achieve all-optical packet switching in WDM Wide Area Networks (WANs) via a novel technique, called slot routing. Using slot routing, entire slots, each carrying multiple packets on distinct wavelengths, are switched transparently and individually. As a result packets can be optically transmitted and switched in the network using available fast and wavelength non-sensitive devices. The proposed routing technique leads to an optical packet switching solution, that is simple, practical, and unique as it makes it possible to build a WDM all-optical WAN with optical devices based on proven technologies.

Keywords: All-Optical Networks, Wide Area Networks, Wavelength Division Multiplexing, Photonic Slot Routing, Time Slot Assignment Algorithm, Optical Packet Switching

1. INTRODUCTION

Considerable attention is being dedicated to the development of all-optical Wide Area Networks (WANs). These networks offer a huge transmission bandwidth with bit error rates as low as 10^{-12}-10^{-15}, a contained network latency, data transparency and freedom from interference. Wavelength Division Multiplexing (WDM) technology offers a practical way to exploit the vast bandwidth of optical fiber of about 30 Thz.[1] WDM partitions the optical bandwidth into separate channels, each at a different wavelength, operating at transmission rates compatible with the electronics speed available today, to support transmission and reception at an aggregate bandwidth beyond any single channel system. Generally, one can distinguish between two classes of optical network architectures, single-hop and multihop.[2] In single-hop networks, optical signals travel from source to destination without encountering electronic regeneration. Single-hop networks are "transparent" for (optical) signals with different modulation formats. In multihop networks, connections consist of a sequence of single-hop paths that are concatenated by means of electronic switching. Multihop connections are therefore not inherently optically transparent.

Available technology provides today the necessary optical devices to achieve all-optical *circuit switching*.[3] Based on the *wavelength routing* concept, started by Refs. 4,5, node pairs can establish point-to-point paths of light (or lightpaths) for data exchange. Intermediate nodes along the lightpath optically route the signal, thus avoiding opto-electronic (O/E) and electro-optical (E/O) conversions of the transmitted signal. Single-hop transmission is possible between nodes connected by a lightpath. However, as each lightpath requires one wavelength and the number of wavelengths is finite, not every node pair can be connected via a lightpath. Nodes that cannot be connected by a lightpath use a multihop connection to communicate with one another.

Solutions for all-optical *packet switching* have been proposed for regular topologies, such as star,[6] ring,[7] and bus.[8] These topologies are inherently passive and node to node transmission is achieved using the "broadcast and select"

Other author information: (Send correspondence to G.W.)
I.C.: E-mail: chlamtac@utdallas.edu
A.F.: E-mail: fumagalli@utdallas.edu; A.F. is currently on leave from the Politecnico di Torino.
G.W.: E-mail: wedzing@nlr.nl; WWW: http://www.nlr.nl; Telephone: +31-20-5113285; Fax: +31-20-5113210

approach, according to which a transmitted packet is broadcast to all nodes, but only the intended destination selects and receives it. Due to the lack of fast and wavelength sensitive switches, however, no solution has been proposed yet for WANs which are typified by unrestricted topologies*.

This paper proposes to achieve all-optical packet switching in WDM based WANs via a novel technique, called Photonic Slot Routing (PSR). PSR was originally proposed in Refs. 9,10 to provide scalability in regular topologies. According to this concept, packets transmitted simultaneously on distinct wavelengths form a *photonic slot* that is individually (and optically) routed at the switching nodes towards the intended destinations as a whole. The proposed solution can therefore handle *wavelength sensitive* data-flows using *wavelength non-sensitive* fast optical switches based on proven technologies. As a result, the fundamental technological problem of (per wavelength) switching limitations at the switching node is shifted to a solvable problem of finding effective solutions to organize packet transmissions at the source nodes.

This paper addresses the problem of constructing *correct* and *efficient* link frames in the PSR network to support given traffic demands. The correctness aspect of the problem must ensure that all the traffic demands are satisfied, while the efficiency aspect is to ensure that some objective is optimized, such as network throughput, or network delay. The construction of link frames requires the solution of three problems: determine how packets are organized into slots, referred to as the *slot composition problem*, determine the path for each slot, referred to as the *slot routing problem*, and determine when slots must be transmitted, referred to as the *slot transmission problem*. The solution of these problems shall guarantee contention free transmission of the slots throughout the network without requiring optical buffers at the switching nodes. Since the construction of optimal link frames can be proven to be an NP-hard problem,[11] a heuristic solution is proposed, that solves the composition, routing, and transmission problems separately, using distinct algorithms, whereby each algorithm performs optimization within its specific problem domain.

The proposed PSR architecture is simple, practical, and unique as it makes it possible to build WDM all-optical WANs using optical devices based on proven technologies. A PSR network offers single-hop transmission between any arbitrary pair of end nodes. In addition, it is possible to flexibly allocate the network bandwidth to accommodate different patterns of load offered to the network, a flexibility that cannot be reached in wavelength routing systems where the minimum amount of bandwidth allocated between two nodes is the bandwidth carried by one wavelength. Finally, due to its inherent optical transparency the PSR network is scaleable in the number of wavelengths, making it possible to gradually increase the network capacity by incorporating additional wavelengths in the system without having to replace the hardware and control of the switching nodes.

2. PHOTONIC SLOT ROUTING IN A WIDE AREA NETWORK

This section gives a description of the PSR network and its functions, followed by a formal formulation of the problem of constructing the link frames, and its three subproblems of slot composition, slot routing, and slot transmission.

2.1. The PSR Network

Figure 1 shows an example of an all-optical PSR network. It consists of *End Nodes*, *Access Nodes* and *Switching Nodes* interconnected by fiber optic links. Each link is bi-directional and actually consists of a pair of unidirectional links. End Nodes form the sources and destinations of the network traffic. They are connected with a single link to an Access Node. Access Nodes and Switching Nodes are responsible for routing the traffic from source to destination. Switching Nodes and Access Nodes are functionally identical; a Switching Node has however no End Nodes connected to it. Switching Nodes and Access Nodes are interconnected into a network with a mesh topology. Optical amplifiers may be placed on the links to compensate for fiber loss and the Access/Switching Nodes' insertion loss.

In the PSR network End Nodes communicate by means of single-hop connections in which optical signals travel from source to destination without encountering electronic conversion. E/O and O/E conversions take only place at the End Nodes, which can have one out of two types of optical transmitter and receiver: a single fixed transmitter (receiver) that is tuned to a certain wavelength channel, or a single transmitter (receiver) that is tunable to any of the available wavelength channels. For the sake of simplicity all End Nodes are assumed to have the same type of transmitter and receiver pair. The pair fixed transmitter-fixed receiver is not allowed since it guarantees full

---

*Clearly, the broadcast and select approach is not practical in the WAN scenario.

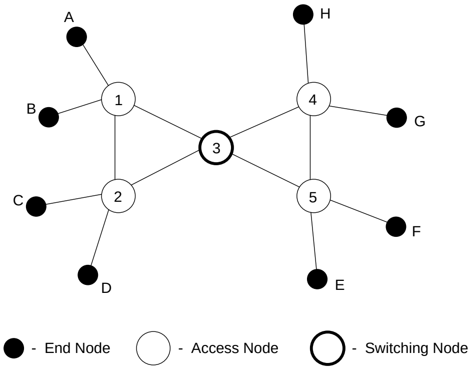

Figure 1. Example of an all-optical PSR wide-area network

connectivity only when one wavelength is used. End Nodes may either be single sources (destinations) of traffic, or represent (all-optical) subnetworks, such as LAN segments.

A PSR network is a time slotted system allowing one fixed length packet to be transmitted in each slot per wavelength. Slots are synchronized across the wavelengths in order to form groups of aligned packets, i.e., the photonic slot. The basic concept of PSR is to limit the complexity of both optical hardware and electronic control of the Access/Switching Node. This is achieved by allowing only simple functions to be performed at the Access/Switching Node that operate on a per slot basis, rather than on a per packet basis. These functions are:

- Slot Switching: slots arriving on any input port (link) can be switched individually to any output port (link). Output ports selected by simultaneously arriving slots must be mutually exclusive to prevent output link contention at the node.
- Slot Merging: slots arriving on a number of input ports can be switched to the same output port, thus overlapping with one another to form one single slot leaving the node via that port. Clearly this operation is only possible when the merged slots are compatible, i.e., they do not carry packets on the same wavelength channel.
- Slot Splitting: a slot arriving on an input port is duplicated and switched to two or more output ports.

To be realized, these basic functions require passive couplers, splitters, and wavelength non-sensitive space switches with high switching rate, i.e., devices based on proven technologies. If necessary, the basic functions can be combined to perform more complex operations. For example, slots from two input ports are merged and the result is split and switched to three output ports. In the present study, it is assumed that neither optical time slot exchange (by means of fiber delay lines), nor wavelength exchange (by means of wavelength converters) are possible at any node.

2.2. The Link Frame Construction Problem

The Slot Merging function of the Access/Switching Node allows merging of packets (with different destinations) into a single slot. Once packets have been merged, they cannot be separated while they are transmitted due to the limited functionality of the Access/Switching Node. To allow each individual packet to reach its destination, the Slot Splitting function of the Access/Switching Node is used to create duplicate slots, which are then switched to the different destinations of the packets in the slot. At each Destination End Node adequate optical filtering is used to extract the packet intended for it.

The slot composition, routing, and transmission problems are solved by assigning to each packet that needs to be transmitted, a route through the network, a wavelength channel, and a transmission time slot. Slot Merging takes place at the Access/Switching Nodes where the paths of packets in the same time slot join and Slot Splitting takes place at the Access/Switching Nodes where the paths of packets in the same time slot diverge. To avoid contention at the output links of the Access/Switching Nodes, all packets present in the same time slot on a certain link, shall have different wavelengths. Generally, one would be interested in an optimal assignment of routes, wavelength channels and time slots. In this paper, the total network throughput is used as the objective to be optimized. Clearly, the optimal solution depends on the characteristics of the offered traffic and the network topology.

This paper addresses the Routing, Wavelength, and Time slot Assignment (RWTA) problem for the case that Time Division Multiplexing (TDM) is employed to allocate the network capacity to packet transmissions. With TDM, the transmission pattern of each link consists of fixed frames with equal number of time slots, that are repeated in a cyclic fashion. Packet transmissions between a specific source and destination node are associated with a communication connection, which is assigned a wavelength channel and a time slot in the frames of the links that the connection is routed through. Each packet transmission uses the route, the wavelength channel and the time slot of the connection it is associated with. Depending on the traffic demands more than one connection can be established between node pairs.

Given a network topology and the set of connections to be established, the problem is to construct transmission frames for each link in the network, while maximizing the total network throughput. Each connection occupies exactly one time slot in the transmission frames of the links in its path. Assuming a framelength of K slots and a bitrate B, the throughput of a connection is B/K (ignoring any time guards for transmitter/receiver tuning, etc.). If the total number of connections to be supported by the network is Z, the total network throughput will be: $Z \cdot B/K$. Since Z and B are considered fixed, the objective is to construct link transmission frames of minimum length K, that can accommodate all the required connections. This problem is referred to as the *Static PSR Frame Construction* problem. It can be considered as a combination of three subproblems: 1. the routing of each connection from its source node to its destination node, 2. the assignment of a wavelength channel to each connection, and 3. the assignment of a time slot to each connection. Once these three subproblems are solved the construction of the transmission frame for each link is straightforward. A formal definition of the Static PSR Frame Construction problem is defined in set theoretic terms in the following.

STATIC PSR FRAME CONSTRUCTION PROBLEM

INSTANCE: A PSR network is represented by a directed graph $G(V, E)$, with V, the set of vertices, representing the network nodes, and E, the set of edges, representing the unidirectional optical links. $V = D \cup A \cup S$, with D the set of End Nodes, A the set of Access Nodes, and S the set of Switching Nodes. Two types of transmitter $T \in \{FT, TT\}$, and two types of receiver $R \in \{FR, TR\}$ can be used at the End Nodes, with FT (FR) fixed transmitter (receiver), and TT (TR) tunable transmitter (receiver). A set C of connections exists, associated with a source node function $s : C \rightarrow D$, and a destination node function $d : C \rightarrow D$. A positive integer W, represents the number of available wavelengths.

QUESTION: Find

- the routing $r : C \rightarrow P$, with $P = \{p : p \subseteq E\}$,
- the wavelength assignment $w : C \rightarrow \{1, \ldots, W\}$, and
- the time slot assignment $t : C \rightarrow \{1, \ldots, K\}$,

for the minimum value of K (the framelength), subject to the following conditions:

1. $\forall c \in C$, the set of edges $p = r(c)$ shall constitute a *tree*[12] in G with $s(c)$ the root of the tree and $d(c)$ the end vertex of one of the leaves of the tree. (Due to the Slot Splitting function the route may not be a simple path, but become a tree.)
2. $\{(c, c') : c, c' \in C, c \neq c', r(c) \cap r(c') \neq \emptyset, (t(c), w(c)) = (t(c'), w(c'))\} = \emptyset$. (Two connections sharing at least one link cannot be allocated both the same time slot and the same wavelength channel.)
3. $\{(j, k) : j, k \in C, j \neq k, r(j) \cap r(k) = \emptyset$, and $\exists i : i \in C, r(i) \cap r(j) \neq \emptyset, r(i) \cap r(k) \neq \emptyset, t(i) = t(j) = t(k)$, and $w(j) = w(k)\} = \emptyset$. (If two connections do not share any link, but share at least one link with a third connection, and all three connections are assigned the same time slot, then the two original connections cannot be assigned the same wavelength channel.)
4. $|\{c : c \in C, s(c) = n, t(c) = k\}| \leq 1$, $\forall n \in D$, and $\forall k \in \{1, \ldots, K\}$. (A transmitter can transmit on at most one connection in each time slot.)

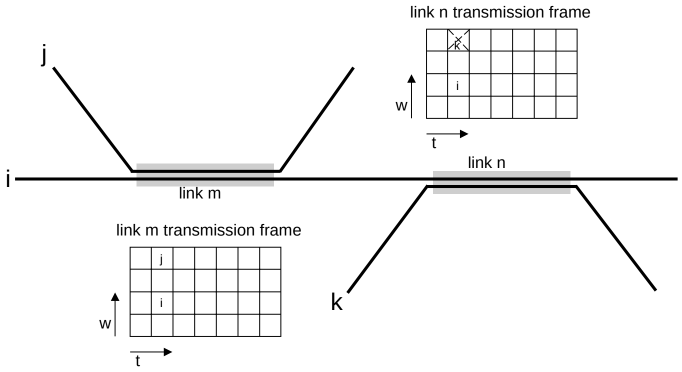

Figure 2. Triple connections condition

5. $|\{c : c \in C, d(c) = n, t(c) = k\}| \leq 1$, $\forall n \in D$, and $\forall k \in \{1, \ldots, K\}$. (A receiver can receive at most from one connection in each time slot.)

The third condition is required because the Access/Switching Node can only split slots, i.e., create duplicate slots, and not separate the packets in a slot to form new slots. This is illustrated in Figure 2, which shows three connections i, j, and k, two physical links n, and m, and their respective transmission frames. Connections i and j share link m, whereas connections i and k share link n. Connections j and k do not share any link, so condition 2 allows them to be assigned the same time slot and wavelength channel. If connections j and k are assigned to the same time slot as connection i, then according to condition 3, connections j and k cannot have the same wavelength channel. As a result of the Slot Splitting function, connection j remains in the slot with connection i, beyond the point where the paths of connections i and j diverge. So, the wavelength channel of connection j cannot be reused by connection k.

In case of fixed transmitters or fixed receivers, the wavelength channels are fixed and therefore the wavelength assignment function w is known a priori and can be removed from the question part of the problem definition.

The Static PSR Frame Construction problem is an NP-hard problem.[11] This can be proven e.g., by first restricting the number of wavelength channels W to 1, assuming the routing to be known, and then transforming the resulting problem to the Static Lightpath Establishment (SLE) problem that was proven to be NP-hard in Ref. 5.

3. ROUTING, WAVELENGTH, AND TIME SLOT ASSIGNMENT ALGORITHMS

Since the Static PSR Frame Construction problem is NP-hard, it is necessary to search for a heuristic solution for all but trivial sets of connections. For this study, it was decided to determine *sub-optimal* solutions under the assumption that routing, wavelength assignment, and time slot assignment are treated independently. The algorithms used for each of the three subproblems are described in separate subsections.

In principle, slot merging and splitting can take place at any Access/Switching Node in the network. Since partially filled slots do not fully utilize the available transmission capacity, slot merging should preferably take place early after slots have been transmitted. Similarly, slot splitting results in duplicate slots, which partially waste available transmission capacity. Slot splitting should therefore preferably take place as late as possible before a slot reaches its destinations. To simplify the Static PSR Frame Construction problem, an additional constraint in the form of a slot merging/splitting policy is introduced. The slot merging/splitting policy specifies which connections are allowed to be combined into a single slot and where slot merging and splitting can take place. The slot merging/splitting policy applied here allows only compatible slots containing connections passing through the **same** Destination Access Node to be merged at the Access/Switching Node where their paths join. When the slot reaches its Destination

Access Node, it is split (duplicated) to all connected End Nodes. The Destination End Nodes select the appropriate wavelength channel from the received slot.

3.1. Routing

The result of the routing step is the routing function $r(c)$, which specifies for each connection the set of links constituting the path of the connection. The routing algorithm needs to satisfy condition 1 given in the problem definition (Section 2.2). For reason of simplicity we use Dijkstra's routing algorithm[13] to find a path with the minimum number of fiber links. This allows all connections to be routed independently of each other, but it may be at the cost of some network throughput, since the network congestion is not necessarily minimized.

In case of equal path lengths an arbitrary path is chosen, but always the same selection is made for the path between two arbitrary nodes a and b. This guarantees that connections passing through the same Destination Access Node follow the same path, after their paths join at some point. Therefore, these connections may be merged into a single slot by the time slot assignment algorithm conform the slot merging/splitting policy.

The slot merging/splitting policy is implemented in the routing algorithm by copying connections from the Destination Access Node to all connected End Nodes. The route of a connection is thereby modified from a simple path into a tree.

3.2. Wavelength Assignment

The result of the wavelength assignment step is the wavelength allocation function $w(c)$, which specifies for each connection the wavelength channel that is used. The wavelength assignment problem becomes trivial in the cases of **fixed** transmitter or **fixed** receiver configurations. In these cases the wavelength allocated to a connection is determined by the wavelength channel that the transmitter of the source or receiver of the destination is tuned to. The following equations give the wavelength assignment used in case of fixed transmitters and fixed receivers, respectively:

$$w(c) = s(c) \bmod W + 1, \tag{1}$$

$$w(c) = d(c) \bmod W + 1, \tag{2}$$

where w represents the assigned wavelength channel number, s the source (address) of connection c, d the destination (address) of connection c, and W the number of available wavelength channels.

For the transmitter/receiver configurations that allow free wavelength selection, an algorithm is used to assign wavelengths such that the conditions for merging connections passing through the same Destination Access Node are optimized. The algorithm is based on the round robin principle, and uses a counter Q_i for each Access Node i. Let function $a(e)$ specify the Access Node to which End Node e is connected. Then the algorithm can be represented in pseudo-code as follows:

```
for  $i \in A \rightarrow \quad Q_i \leftarrow 0$ rof;     {Counter initialization}
for  $c \in C \rightarrow$
     $i \leftarrow a(d(c))$;                   {Destination Access Node}
     $w(c) \leftarrow Q_i + 1$;                {Wavelength channel assignment}
     $Q_i \leftarrow (Q_i + 1) \bmod W$;       {Counter update}
rof;
```

When the number of wavelength channels (W) is greater than or equal to the maximum number of End Nodes (N) connected to an Access Node, the performance of the TT-TR configuration is the same as the performance of the TT-FR configuration. (The proof is beyond the scope of this paper.) So, instead of using the round robin algorithm, in the case of $W \geq N$, the fixed assignment of Equations (1) and (2) guarantees that the maximum performance can be achieved.

3.3. Time Slot Assignment

The result of the time slot assignment step is the time slot allocation function $t(c)$, which specifies for each connection the time slot that is used. Time slot assignment is carried out using a graph coloring method introduced in Ref. 9, which is extended in this paper to handle general topology networks. Using the routing function $r(c)$, the wavelength allocation function $w(c)$, and the problem conditions from Section 2.2, a *Slot Constraints* graph $\tilde{G}(\tilde{V}, \tilde{E})$ is constructed, whose vertex coloring represents a time slot assignment for all connections. The graph $\tilde{G}$ is constructed in the following steps:

1. Let $\tilde{V} = C$, i.e., each communication connection constitutes a vertex in the graph $\tilde{G}$.
2. Implement condition 2 by placing an edge between each pair of vertices associated with connections that have the same wavelength and are present on the same link. (Two vertices cannot be assigned the same color if they are connected by means of an edge. Hence, the associated connections cannot be allocated to the same time slot.)
3. Implement condition 4 by placing an edge between each pair of connections that originate from the same source.
4. Implement condition 5 by placing an edge between each pair of connections that have the same destination.
5. Implement the merging/splitting policy by placing an edge between each pair of vertices associated with connections that are present on the same link and do not pass through the same destination Access Node.

Due to the merging/splitting policy, which only allows connections passing through the same Destination Access Node to be merged in a single slot and which splits slots at the Destination Access Node to all connected End Nodes, condition 3 of the problem definition does not need to be implemented, since connections that are merged do not have diverging paths.

After the graph $\tilde{G}$ is constructed, its vertices are colored by some graph coloring algorithm. By assigning a unique integer t from the set $\{1, \ldots, K\}$, with K the number of different colors, to each color, a slot number $t(c)$ is obtained for each connection c. The algorithm by Welsh and Powell[14] is used to color the Slot Constraints graph. The exact coloring algorithm by Randall Brown[15] has also been considered. Some experimentation with this algorithm revealed that for low network loads the performance improvement (i.e., the reduction in number of time slots needed) compared to the algorithm of Welsh and Powell is insignificant, whereas for larger network loads the execution time of the Randall Brown algorithm becomes prohibitive. In addition it is noted that the results of the experiments are meant to be compared with each other. For these reasons, it was decided to use the Welsh and Powell algorithm for all experiments.

Using the routing function r (Section 3.1), the wavelength allocation function w (Section 3.2), and the time slot allocation function t (Section 3.3), transmission frames can be determined for each link in the PSR network operating in TDM mode.

4. PERFORMANCE ANALYSIS OF A BENCHMARK NETWORK

This section presents the results of a number of experiments to investigate the performance that the proposed PSR network architecture can achieve, by using the algorithms outlined in Section 3 to construct link transmission frames. The total network throughput is used as the performance measure. The benchmark network used in the experiments is similar to the network shown in Figure 1. Instead of two End Nodes, now $N = 8$ End Nodes are connected to each of the four Access Nodes. The interconnection of the Access Nodes and the Switching Node is unchanged. All End Nodes in the network are numbered sequentially, by numbering subsequently all the End Nodes connected to the same Access Node, for each Access Node.

In the experiments, the network load is determined by the number of connections that have to be routed and the sources and destinations of these connections. For each connection c, the source $s(c)$ is selected randomly from the set of End Nodes, and the destination $d(c)$ is selected randomly from the set of remaining End Nodes, using discrete uniform distributions. In all experiments, the results are obtained with 3% (or smaller) confidence intervals at 99% confidence level.

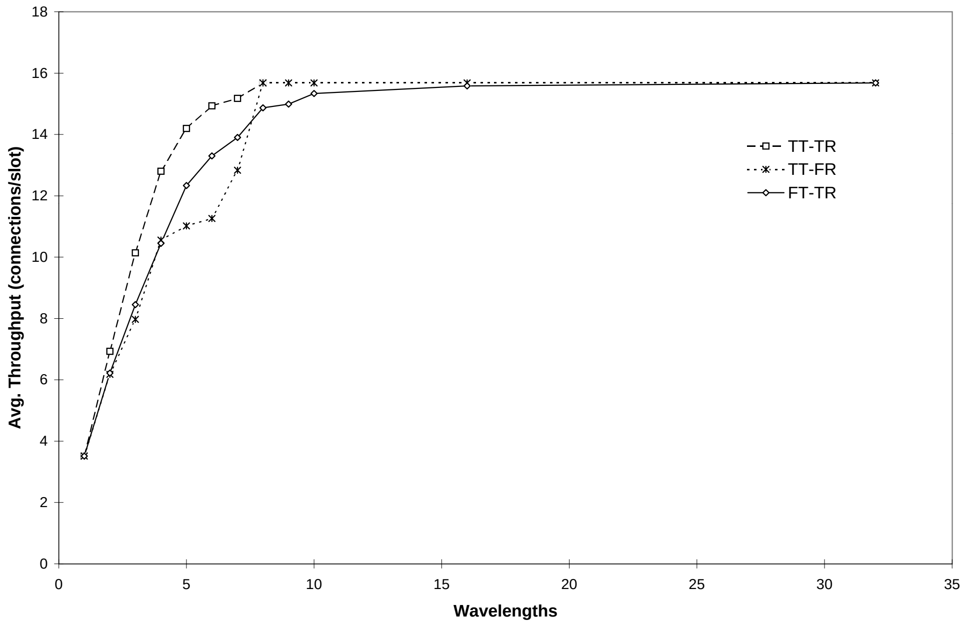

Figure 3. Throughput of TT-FR, FT-TR, and TT-TR configurations as a function of the number of wavelength channels W, for a fixed network load of 200 connections

4.1. Transmitter/Receiver Configurations

The performance of the different transmitter/receiver (T/R) configurations is investigated first. As described in Section 2.1, the choice of the transmitter and receiver type is homogeneous over all End Nodes in the network. There are two possible types of transmitter, i.e., single tunable transmitters (TT), and fixed tuned transmitters (FT). Similarly, two types of receiver have been defined: single tunable receivers (TR) and fixed tuned receivers (FR). Out of the four (2^2) possible T/R configurations, the combination of fixed tuned transmitters with fixed tuned receivers is the only one not allowed. Figure 3 shows the average network throughput of the three feasible T/R configurations as a function of the number of available wavelength channels W for a fixed network load of 200 connections. The network throughput is expressed as the number of connections supported by the network, divided by the framelength.

The highest performance is achieved by the TT-TR configuration. This can easily be verified by observing that any solution (i.e., set of link frames) for a network with the FT-TR configuration also represents a solution for a network with the TT-TR configuration, where each transmitter is always tuned to the same wavelength channel. Likewise, any solution for the TT-FR case also represents a solution for the TT-TR case, where each receiver is always tuned to the same wavelength channel. Two regions can be distinguished in the figure for the number of wavelength channels, i.e., the region $W < N$, i.e., 8, and the region $W \geq N$. In the $W \geq N$ region the performance of the TT-FR and TT-TR configurations are equal and do not depend on the number of wavelength channels. The performance of FT-TR is slightly lower in this region, but approaches the performance of TT-FR, and TT-TR for large numbers of wavelength channels. In the $W < N$ region, the performance of the TT-FR configuration is generally lower than the performance of the FT-TR configuration, but for certain values, e.g., $W = 4$, TT-FR performs better slightly than FT-TR. It can be shown that the relative performance of FT-TR and TT-FR in this region depends on the value of $N \bmod W$. If $N \bmod W = 0$, then TT-FR generally performs slightly better than FT-TR; if $N \bmod W \neq 0$, FT-TR performs better. Overall, the best performance is obtained with the TT-FR and TT-TR configurations at $W = N$ wavelength channels. From a system cost point of view, preference is given to the TT-FR configuration, since fixed receivers are considered to be cheaper than tunable receivers.

4.2. Slot Routing Versus Optical Packet Switching

In a second investigation the performance of the PSR network is compared with the performance achievable by an ideal solution (OPS) in which Optical Packets are individually Switched without taking into account the functional

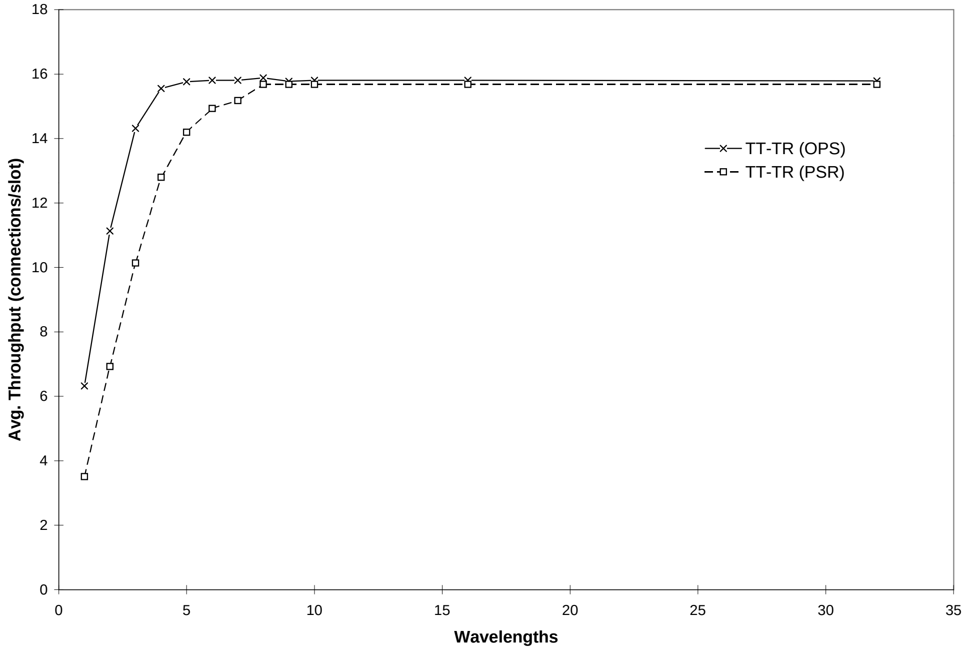

Figure 4. Throughput of PSR and OPS networks for the *TT-TR* configuration as a function of the number of wavelength channels W, for a fixed network load of 200 connections

limitations of extant optical devices. Although not currently feasible, OPS can be used as a reference solution for assessing the PSR efficiency. Only the *TT-TR* configuration is considered in this experiment, since it achieves the highest performance of all T/R configurations. Figure 4 shows the average throughput as a function of the number of available wavelengths W for a fixed network load of 200 connections, for the *TT-TR* configuration in case of the PSR network and the optical packet switching (OPS) network.

Despite the fact that the performance of the *TT-TR* configuration is limited and becomes independent of the number of wavelength channels W for $W \geq N$, the PSR network achieves almost the same performance as the OPS network. The same experiment was carried out with a number of different topologies, showing similar results.

5. CONCLUSIONS

This paper proposed a novel approach, called Photonic Slot Routing (PSR), to achieve optical packet switching in WDM based WANs that relies on extant optical technology. According to PSR, entire slots, each carrying multiple packets on separate wavelengths, are routed at the switching nodes “transparently” as single units. By not requiring to route individual packets (on separate wavelengths) wavelength sensitive data flows can thus be handled using wavelength non-sensitive devices based on proven technologies. In addition, due to its inherent optical transparency the PSR network is scaleable in the number of wavelengths.

A heuristic was presented that computes composition, routing and transmission time of the photonic slots to maximize network throughput given the network topology and the traffic demands. The throughput measured in a benchmark network making use of PSR was compared with the throughput achievable by an ideal solution in which optical packets are individually switched without taking into account the functional limitations of extant optical devices. Under the assumption that end nodes make use of one transmitter and one receiver, hardly any performance penalty was found using PSR when compared to the ideal packet switching scenario.

Overall, it can be concluded that while fast wavelength sensitive devices are not yet available to realize pure optical packet switching, the photonic slot routing concept provides an alternative solution to the realization of packet switching that is practical and sufficiently efficient.

ACKNOWLEDGMENTS

This research is sponsored by NSF under contracts # NCR-9628189 and # NCR-9596242.

REFERENCES

1. C.A. Bracket, "Dense Wavelength Division Networks: Principles and Applications," *IEEE Journal on Selected Areas in Communications*, vol. 8, August 1989.
2. B. Mukherjee, "WDM-Based Local Lightwave Networks Part I: Single-Hop Systems," *IEEE Network*, May 1992.
3. IEEE Communications Magazine, vol. 36, no. 2, February 1998.
4. I. Chlamtac, "Rational, Directions and Issues Surrounding High Speed Computer Networks," *IEEE Proceedings*, vol. 78, no. 1, pp. 94-120, January 1989.
5. I. Chlamtac, A. Ganz, G. Karmi, "Lightpath Communications: A Novel Approach to High Bandwidth Optical WAN-s," *IEEE Transactions on Communication*, vol. 40, no. 7, July 1992.
6. M.-S. Chen, N.R. Dono, R. Ramaswami, "A Media-Access Protocol for Packet-Switched Wavelength Division Multiaccess Metropolitan Area Networks," *IEEE Journal of Selected Areas in Communication*, vol. 8, no. 6, pp. 1048-1057, August 1990.
7. I. Chlamtac, A. Fumagalli, L.G. Kazovsky, P.T. Poggiolini, "A Contention/Collision Free WDM Ring Network for Multi Gigabit Packet Switched Communication," *Journal of High Speed Networks*, vol. 4, no. 2, pp. 201-219, 1995.
8. S. Banerjee, B. Mukherjee, "Fairnet: A WDM-based Multiple Channel Lightwave Network with Adaptive and Fair Scheduling Policy," *Journal of Lightwave Technology*, vol. 11, no. 5-6, pp. 1104-1111, May/June 1993.
9. I. Chlamtac, V. Elek, A. Fumagalli, Cs. Szabó, "A Scalable Optical Network Based on Folded Bus Architecture," *Proc. of EUROPTO - European Symposium on Advanced Imaging and Network Technologies (EOS-SPIE)*, Berlin, FR Germany, October 1996.
10. I. Chlamtac, V. Elek, A. Fumagalli, "A Fair Slot Routing Solution for Scalability in All-Optical Packet Switched Networks," *accepted in the Journal of High Speed Networks.*
11. M.R. Garey, D.S. Johnson, *Computers and Intractability, A Guide to the Theory of NP-completeness*, W.H. Freedman and Company, San Francisco, 1979.
12. L.R. Foulds, *Graph Theory Applications*, Springer-Verlag, New York, 1992.
13. E.W. Dijkstra, "A Note on Two Problems in Connexion with Graphs," *Numerical Mathematics*, vol. 1, pp. 269-271, October 1959.
14. D.J.A. Welsh, M.B. Powell, "An Upper Bound for the Chromatic Number of a Graph and its Application to Timetabling Problems", *The Computer Journal*, vol. 140, pp. 85-86, 1967.
15. J. Randall Brown, "Chromatic Scheduling and the Chromatic Number Problem," *Management Science*, vol. 19, no. 4, pp. 456-463, December 1972.