



NLR-TP-99478

Upgradable photonic slot routing networks to cope with increasing capacity demands

G. Wedzinga, I. Chlamtac, A. Fumagalli and J.B. Carruthers



NLR-TP-99478

Upgradable photonic slot routing networks to cope with increasing capacity demands

G. Wedzinga, I. Chlamtac*, A. Fumagalli* and J.B. Carruthers**

* University of Texas at Dallas

** Boston University

This report is based on a presentation held at SPIE's International Voice, Video and Data Communication Symposium, Boston, USA, 19-22 September 1999.

The contents of this report may be cited on condition that full credit is given to NLR and the authors.

Division: Avionics
Issued: November 1999
Classification of title: unclassified



Contents

1	Introduction	3
2	Enhanced photonic slot routing network	4
2.1	End nodes	4
2.2	Switching nodes	5
2.3	Transmission control	6
2.4	Synchronization aspects	7
3	Transmission frame construction for the E-PSR network	8
3.1	The transmission frame construction problem	8
3.2	A transmission frame construction algorithm	9
4	Performance analysis	10
4.1	Performance bounds	10
4.2	Uniform capacity demand	11
4.3	Non-uniform capacity demand	12
5	Conclusion	13



Upgradable photonic slot routing networks to cope with increasing capacity demands

Imrich Chlamtac^a, Andrea Fumagalli^a, Gosse Wedzinga^b and Jeffrey B. Carruthers^c

^aCenter of Advanced Telecommunications Systems and Services, University of Texas at Dallas,
P.O. Box 830688, Richardson TX 75083-0688, USA

^bAvionics Department, National Aerospace Laboratory NLR,
P.O. Box 90502, 1006 BM Amsterdam, The Netherlands

^cDepartment of Electrical and Computer Engineering, Boston University,
8 Saint Marys Street, Boston MA 02215, USA

ABSTRACT

Recently, Photonic Slot Routing (PSR) has been proposed as a practical solution for achieving optically transparent communication in Wavelength Division Multiplexing (WDM) wide-area networks. PSR makes use of available and simple wavelength-non-sensitive optical devices to switch data flows at different wavelengths, but in the same time slot, as a single unit. It thereby trades switching node complexity for network performance. With continuing traffic growth, the need to increase the network performance will, however, inevitably arise.

As optical device technology progresses, it will become feasible to implement more complex switching nodes that can perform Individual Switching of the Wavelengths within a time slot (IWS). While IWS nodes can be expected, at least in the beginning, to be more costly than PSR nodes, they offer increased switching flexibility, thereby potentially improving the network performance. This paper presents a solution that supports a gradual migration from a PSR-based network to an IWS-based network, by replacing selected PSR nodes with IWS nodes.

We show that with a controlled migration towards an IWS-based network it is possible to gradually increase the capacity of the network to cope with increasing traffic demands. An important property of the solution is that upgrading a PSR node to an IWS node does not require any changes to the other nodes, or the network layout. In addition, using the NSFNET backbone topology as a benchmark, the paper shows that with an 80% PSR and 20% IWS network almost 50% of the performance reduction incurred by slot routing can be recovered.

Keywords: All-Optical Networks, Wavelength Division Multiplexing, Photonic Slot Routing, Wavelength and Time Slot Assignment Algorithm

1. INTRODUCTION

Recently, the authors have proposed Photonic Slot Routing (PSR) as a practical solution to achieve optically transparent communication in Wavelength Division Multiplexing (WDM) wide-area networks.⁴ The concept is equally applicable for local- and metropolitan-area networks with irregular mesh topologies. Using the PSR concept, data flows transmitted simultaneously on distinct wavelengths are grouped to form a *photonic slot* that is optically routed as a single unit at the switching nodes towards the intended destinations. The PSR approach is unique in that it can handle *wavelength-sensitive* data flows using *wavelength-non-sensitive* fast optical switches based on proven technologies.¹³ When compared to the wavelength routing concept,^{2,3} in which the entire wavelength bandwidth can be assigned to one traffic demand only, the PSR concept offers the possibility to multiplex each wavelength bandwidth among different capacity demands by means of *optical* traffic grooming. Optical transparency is thereby retained, which is a crucial property for coping with the existing proliferation of protocols.¹⁰ Due to the wavelength-non-sensitive characteristic of the optical switches, however, slot routing imposes some constraints on the scheduling

Further author information: (Send correspondence to G.W.)

I.C.: E-mail: chlamtac@utdallas.edu; Telephone: +1-972-883-6295; Fax: +1-972-883-6204

A.F.: E-mail: andrea@utdallas.edu; Telephone: +1-972-883-6853; Fax: +1-972-883-2710

G.W.: E-mail: wedzing@nlr.nl; WWW: <http://www.nlr.nl>; Telephone: +31-20-5113285; Fax: +31-20-5113210

J.B.C.: E-mail: jbc@bu.edu; Telephone: +1-617-353-0204



of transmissions at the source nodes. Under certain traffic patterns, e.g., unbalanced traffic, these constraints may limit the system capacity. In short, the PSR approach trades optical layer complexity for network performance.

For the near future the proposed PSR solution is expected to be more than adequate, since the large number of WDM channels provides a system capacity that will be far beyond the capacity of any system in use today. However, with traffic continuing to grow over time, there will be a need to increase the network capacity at some point in time. One way to increase the network capacity is by increasing the number of wavelengths in the system. Although this affects the end nodes of the network, no upgrade of the switching nodes is required, due to their wavelength-non-sensitivity. Assuming, that the capacity increase obtained in this way cannot cope with the traffic growth, the network can also be upgraded by adding fiber optic links between switching nodes, where needed. However, when no dark fiber is available, new cables have to be put in place. In addition, the number of input and output ports of the switching nodes involved needs to be increased. The extremely high cost of this upgrade may be prohibitive. Therefore, an alternative solution is explored in this paper, which entails upgrading of selected network nodes to provide a more efficient switching technique. By replacing the wavelength-non-sensitive optical switches with the more complex wavelength-sensitive switches, the upgraded nodes provide better multiplexing granularity. The upgraded node, termed Individual Wavelength Switching (IWS) node, is able to switch data flows on each wavelength individually and independently from the data flows on the other wavelengths. Since IWS nodes are more complex than PSR nodes, in the optical part as well as in the electronic control, they are expected to be more costly than the latter. In a mixed PSR-IWS network, a trade-off can thus be made between network capacity and cost.

In this paper we study the increase in network capacity that can be obtained by gradually replacing selected PSR nodes by IWS nodes. We will call such a mixed PSR-IWS network *Enhanced Photonic Slot Routing* network, or E-PSR network for short. Existing algorithms for constructing transmission frames in PSR networks, only address the two extreme cases where all switching nodes are either PSR nodes, or IWS nodes.⁴ A new solution needs therefore to be developed that is able to handle the mixed PSR/IWS case as well. We develop a suboptimal transmission scheduling algorithm that maximizes the network throughput in an E-PSR network under given capacity demands. The study indicates that by replacing a relatively small percentage of PSR nodes with IWS nodes, already a significant increase in network throughput can be achieved. A gradual introduction of IWS nodes in order to cope with gradually increasing capacity demands is a cost effective solution, since only the nodes that are upgraded are effected and no other part of the network needs to be changed.

2. ENHANCED PHOTONIC SLOT ROUTING NETWORK

An Enhanced Photonic Slot Routing (E-PSR) network consists of *end nodes*, and *switching nodes* interconnected by *fiber optic links*. Each link is bi-directional and actually consists of a pair of unidirectional links. The end nodes (Section 2.1) form the sources and destinations of the network traffic, and are connected with a single link to exactly one switching node (Section 2.2). The switching nodes are responsible for routing the traffic from source to destination. Switching nodes exist in two types, i.e., Photonic Slot Routing (PSR) nodes and Individual Wavelength Switching (IWS) nodes. There are no constraints on the choice of the node type at a switching node location. Switching nodes are interconnected into a network with a mesh topology. Figure 1 shows an example of an E-PSR network configuration.

The E-PSR network operates on the basis of the Photonic Slot Routing concept. It is a time-slotted WDM system, in which time slots are synchronized across the wavelengths in order to form groups of aligned data flows, i.e., photonic slots. PSR nodes treat photonic slots as indivisible units that are switched as single, transparent units of information. Thus, PSR nodes can handle wavelength-sensitive data flows using wavelength-non-sensitive devices, such as the space digital switch. IWS nodes are capable of independently switching the individual wavelength channels in photonic slots, thereby effectively reorganizing the photonic slots. Therefore, IWS nodes have to be implemented with wavelength-sensitive devices. Details of transmission control in the E-PSR network are described in Section 2.3. Some synchronization aspects of the PSR concept are discussed in Section 2.4.

2.1. End Nodes

The end nodes in the E-PSR network generally represent gateways to subnetworks (e.g., ATM based networks). Such a subnetwork may be all-optical, in which case the end node acts as a bridge. End nodes may also be high performance single node systems. Each end node is modeled as having an arbitrary number of tunable transmitters and receivers, independently capable of accessing any wavelength channel; tuning is assumed to be performed instantaneously. The number of transmitters and receivers at an end node may depend on the capacity demands.

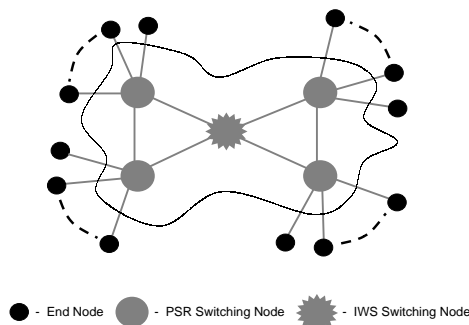


Figure 1. Example of an all-optical Enhanced PSR network

2.2. Switching Nodes

Switching nodes in the E-PSR network exist in two types, i.e., Photonic Slot Routing (PSR) nodes (Section 2.2.1) and Individual Wavelength Switching (IWS) nodes (Section 2.2.2). Note that when all switching nodes in an E-PSR network are PSR nodes, the original PSR network for irregular mesh topologies⁴ is retained. The two types of switching node have in common that they handle incoming data flows in the optical domain, whereby the switch operations are controlled by electronic circuitry. To keep the optical hardware and the electronic control simple and inexpensive, switching nodes do not provide any form of optical buffering, or wavelength conversion.

2.2.1. PSR Nodes

PSR nodes perform the following functions:

- Slot copying: a slot arriving on an input port is duplicated. A copy of the slot is created for each output port.
- Slot switching: A copy of a slot is either switched to its designated output port, or it is purged.
- Slot merging: slots switched to the same output port are merged, 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 data on the same wavelength channel.

To be realized, these functions require passive splitters, combiners, and wavelength-non-sensitive switches with high switching rate; all devices based on proven technologies. The simplicity of the optical hardware of the PSR node is demonstrated in Figure 2 with a possible architecture for an $N \times N$ switch. Each input port is connected to a 1:N splitter, whose output branches are connected to an on/off switch. The outputs of the on/off switches are recombined with an N:1 combiner for each output port, as shown in the figure. Switching operations are achieved by appropriate (electronic) control of the on/off switches (not shown in Figure 2). The architecture requires a total of $N \times N$ on/off switches, which can for instance be realized by gain-clamped Semiconductor Optical Amplifiers (SOA), with possible switching speeds below 1 ns.^{7,15} Examples of 4×4 integrated optical amplifier gate switches are described in Refs. 11,1.

2.2.2. IWS Nodes

Contrary to PSR nodes, IWS nodes are capable of individually switching all wavelengths in an incoming slot to the output port required by each wavelength. IWS nodes are generally more complex than PSR nodes, both in optical hardware, and electronic control, and are therefore assumed to be more costly. IWS nodes perform the following functions:

- Wavelength separation: The wavelength channels in a received photonic slot are separated into the individual wavelength channels.




Figure 2. Architecture of the optical part of a PSR node with N input ports and N output ports

- **Wavelength switching:** The individual wavelength channels are switched to an output port. To avoid output contention, two wavelength channels occupying the same wavelength cannot be switched to the same output port.
- **Wavelength merging:** The different wavelength channels destined to the same output port are merged into a photonic slot.

Due to the current lack of fast wavelength-sensitive devices, IWS nodes have to be based on the well-known combination of wavelength demultiplexers, fast space switches, and wavelength multiplexers. First, wavelength demultiplexers separate the wavelengths in space. Then, space switches route each wavelength from the incoming fiber to the designated outgoing fiber. Finally, multiplexers combine the different wavelengths destined for the same outgoing fiber.

2.3. Transmission Control

The purpose of transmission control is to provide orderly transmission under the constraints imposed by the E-PSR architecture, thereby fairly allocating the network capacity to the capacity demands and maximizing the network throughput. This can be achieved by a suitable scheduling of transmission, switching, and reception at the source nodes, the switching nodes, and the destination nodes, respectively. Since the switching nodes do not provide any form of optical buffering for resolving contentions, Time Division Multiplexing (TDM) is employed in this study to allocate the network capacity to data transmissions. With TDM, the transmission pattern of each link consists of frames with equal number of time slots, that are repeated in a cyclic fashion. To exchange information, node pairs make use of communication connections, which provide a transmission capacity of one wavelength channel for the duration of one time slot per frame. A connection uses a fixed wavelength channel and a fixed time slot in the transmission frames of the traversed network links. The capacity demand, i.e., the number of connections required between each pair of end nodes, is assumed to be fixed and predetermined. The TDM transmission frames will therefore be static over time.

The key to making efficient use of the network's capacity is to create well utilized slots that are transferred through the network. In general, end nodes will not have sufficient transmitters to transmit simultaneously on all available wavelength channels. Therefore, the slots produced by the end nodes will normally be under-utilized. Slot utilization can be increased by merging compatible slots (which may contain data flows with different destinations) at some switching node. Once data flows have been merged into a single slot, they can only be separated at IWS nodes by the wavelength separation function. To allow each individual data flow to reach its destination in the absence of IWS nodes, the slot copying function at PSR nodes creates duplicate slots, which are subsequently switched to the different destinations of the data flows in the slot. Adequate optical filtering is used at each destination end node




Figure 3. Unusable transmission capacity due to slot merging and slot copying

to extract the data intended for it. Besides increasing the slot utilization, slot merging and slot copying at the PSR nodes also result in unusable transmission capacity, as shown by an example in Figure 3. Here, a connection c_1 on wavelength 1 traverses links, m , o , and r , while another connection c_2 on wavelength channel 2 traverses links n , o , and q . Due to slot merging by PSR switching node A, wavelength 1 is *blocked* for use on link n . (Use of wavelength 1 on link n would result in a collision on link o .) Due to slot copying by PSR switching node B, connection c_1 is copied to link q , thereby *wasting* wavelength channel 2. Similarly, wavelength channel 2 is blocked for use on link m and wasted on link r . In general, an unintended blocked wavelength link set (p^B), and an unintended wasted wavelength link set (p^W) are associated with the intended path p of each connection. In conclusion, when optimizing the use of the transmission capacity it should be taken into account that slot merging and copying increase the slot utilization, but also can reduce the available transmission capacity.

2.4. Synchronization Aspects

For proper operation of the slotted TDM protocol it is necessary to ensure that all the input ports to the switching nodes are synchronized (i.e., frames arriving at different input ports of a switching node have to be aligned in time). One scheme to achieve frame synchronization is to use delay lines with delays equal to a fractional part of a frame, which are connected in series. An electronic synchronization circuit compares the arrival time of the frame with the phase of a local frame clock and generates the appropriate setting of the delay lines, so that the input stream to the switch is aligned with the local frame clock.¹² This scheme requires synchronizers at all input links of every switching node. This increases the network latency, since every synchronizer introduces a delay, which can be as much as the frame transmission time. Possible architectures of synchronizers are described in Ref. 8. In this paper it is, however, assumed that frame synchronization can be achieved with 100% transmission capacity, without having knowledge of the implementation details.

Due to dispersion of the optical fiber, the different wavelengths in a photonic slot travel with different speeds. This results in slot skewing, the amount of which depends on the traveled distance. One way to cope with dispersion, is to use the maximum slot skewing as a guard-time, to ensure that the duration of a photonic slot does not exceed its allocated duration after propagating the maximum distance. Another way to deal with dispersion is to create a zero total dispersion transmission system, which uses non-zero dispersion fiber to cover the distances, in combination with dispersion-compensating fiber sections located at switching nodes and line amplifiers. Zero dispersion fiber is not an option, because multiple wavelength channels propagating with the same velocity are seriously degraded by non-linear effects.



3. TRANSMISSION FRAME CONSTRUCTION FOR THE E-PSR NETWORK

3.1. The Transmission Frame Construction Problem

Given an E-PSR network and the capacity demands, the problem is to construct TDM transmission frames for each link in the network, such that the constraints of the E-PSR architecture are satisfied, the capacity demands are allocated fairly, and the network throughput is maximized. Since the capacity demands, i.e., the number of connections between each pair of nodes, are assumed to be fixed and predetermined, the TDM transmission frame will be static over time. The problem is therefore referred to as the *Static E-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. In this paper it is assumed that routing of the connections is known beforehand and does not form part of the problem. By allowing a variable length of the TDM frame, it will always be possible to construct TDM frames that accommodate any finite number of connections. The length of the TDM frame generally increases with the number of connections to be established. The maximum network throughput is achieved by accommodating all connections in the shortest possible transmission frame. A formal definition of the Static PSR Frame Construction problem is given in the following.

INSTANCE: An *E-PSR network* is represented by a six-tuple $G(V, E, W, n, t, r)$. V is the set of vertices, representing the network nodes. $V = D \cup S$, with D the set of end nodes, and S the set of switching nodes. E is the set of edges, representing the unidirectional optical links. (It is assumed that $(u, v) \in E$ iff $(v, u) \in E \forall u, v \in V$.) A positive integer W specifies the number of wavelength channels available on each link. Function $n : S \mapsto \{PSR, IWS\}$ specifies the type of each switching node. Functions $t : D \mapsto N^+$, and $r : D \mapsto N^+$ specify the number of transmitters, and the number of receivers at each end node, respectively. The *capacity demand* is given by a set $L = \{(s, d, p)\}$ of connections. Each connection is specified by a source node s , a destination node d , and a path p .

QUESTION: Find a *wavelength assignment* $w : L \rightarrow \{1, \dots, W\}$, and a *time slot assignment* $\kappa : L \rightarrow \{1, \dots, K\}$, for the minimum value of K (the frame length), subject to the following constraints:

1. Wavelength continuity, i.e., a connection must be assigned the same wavelength channel on all the links that it traverses. This constraint is implicitly incorporated in the problem formulation by the assignment of a single wavelength to each connection.
2. Time slot continuity. Just as in the wavelength domain, the continuity constraint also applies to the time slot domain.
3. Limited transmitters, i.e., an end node cannot transmit in each time slot on more connections than it has transmitters:

$$|\{c : c = (s, d, p) \in L, s = m, \kappa(c) = k\}| \leq t(m), \quad \forall m \in D, k \in \{1, \dots, K\}. \quad (1)$$

4. Limited receivers, i.e., an end node cannot receive in each time slot on more connections than it has receivers:

$$|\{c : c = (s, d, p) \in L, d = m, \kappa(c) = k\}| \leq r(m), \quad \forall m \in D, k \in \{1, \dots, K\}. \quad (2)$$

5. Two connections sharing a link, cannot share the same resource (time slot and wavelength channel combination). This constraint can be formulated by using $C_{(\omega, \tau)} = \{c : c \in L, \kappa(c) = \tau, w(c) = \omega\}$, the set of connections that are assigned to the same wavelength channel ω and the same time slot τ :

$$p_i \cap p_j = \emptyset, \quad \forall c_i = (s_i, d_i, p_i), c_j = (s_j, d_j, p_j) \in C_{(\omega, \tau)}, \quad i \neq j, \quad \forall \omega, \tau. \quad (3)$$

6. Wasted and blocked wavelength link sets cannot be used by connections. Let $p^W(c)$ be the set of links on which the wavelength $w(c)$ is wasted by connection c , and $p^B(c)$ the set of links on which the wavelength $w(c)$ is blocked by connection c . The constraint can now be expressed as:

$$p_i \cap p^W(c_j) = \emptyset, \quad \forall c_i = (s_i, d_i, p_i), c_j = (s_j, d_j, p_j) \in C_{(\omega, \tau)}, \quad \forall \omega, \tau. \quad (4)$$

Constraint 4 also applies for the case $i = j$ in order to exclude loops in which a connection collides with a copy of itself. Note that it is not necessary to include the constraint $p_i \cap p^B(c_j) = \emptyset$, since violation of this constraint will always result in a violation of constraint 4 as well. Also note that blocked wavelength link sets are allowed to share links. The same applies to wasted wavelength link sets.



```

function WTA( E-PSR-network  $G$ , capacity-demand  $L$ );
  set  $S, T, A$ ; function  $w, \kappa$ ; connection  $c$ ;
  integer  $K$ ; boolean  $this-time-slot-done$ ;
   $S \leftarrow L$ ; { set of connections to do }
   $K \leftarrow 1$ ; { first time slot }
  do  $S \neq \emptyset \rightarrow$ 
     $this-time-slot-done \leftarrow$  false;
     $T \leftarrow \emptyset$ ; { set of connections in this time slot }
     $A \leftarrow S$ ; { set of feasible connections in this time slot }
    do  $\neg this-time-slot-done \rightarrow$ 
       $A \leftarrow feasible-connections(G, T, w, A)$ ;
      if  $A \neq \emptyset \rightarrow$ 
         $c \leftarrow most-eligible-connection(G, T, w, A)$ ;
         $w(c) \leftarrow least-used-wavelength(G, T, w, c)$ ;
         $\kappa(c) \leftarrow K$ ;
         $T \leftarrow T + \{c\}$ ;
         $A \leftarrow A - \{c\}$ ;
         $S \leftarrow S - \{c\}$ ;
      else
         $this-time-slot-done \leftarrow$  true;
         $K \leftarrow K + 1$ ;
      fi;
    od;
  od;
  return  $(\kappa, w, K)$ ;
end;

```

Table 1. Algorithm in pseudo-code for function WTA (Wavelength channel and Time slot Assignment)

3.2. A Transmission Frame Construction Algorithm

The Static E-PSR Frame Construction problem is an NP-hard problem⁹ *. It is therefore necessary to search for heuristic solutions. A greedy algorithm is proposed for assigning a time slot and a wavelength channel to each connection in the capacity demand. The algorithm builds the time slots of the transmission frames sequentially. The construction loop starts by determining the set of connections that have not been assigned a time slot and a wavelength channel. If this set is empty, the algorithm terminates. Otherwise, the set of feasible connections is determined, each of which can be incorporated into the current time slot without violating any of the E-PSR constraints. If the result is an empty set, the construction of a new time slot is started. Otherwise, the connection that is most eligible (the most difficult one) is selected from the set of feasible connections. Next, a wavelength channel is assigned to this connection and it is incorporated in the current time slot. Finally, the connection is removed from the set of connections that are not assigned a time slot and a wavelength channel, and the construction loop is started again.

Table 1 specifies the algorithm in pseudo-code. Function *feasible-connections* determines the set of connections that are compatible with the current time slot by evaluating the constraints specified in the question part of the problem definition. This evaluation involves, among other things, determination of the wasted and blocked paths of a connection under consideration. Blocked and wasted paths always terminate at IWS nodes, whereas they propagate through PSR nodes. Hence, the result of constraints evaluation is dependent on the specific choice of the node type (PSR or IWS) at each switching node location. Function *most-eligible-connection* selects the connection from the set of feasible connections that is the most eligible for inclusion in the current time slot, according to the following criteria (in order of priority):

*The E-PSR network is a generalization of the PSR network. The Static PSR Frame Construction problem was proven to be NP-hard in Ref. 5.



- Longest path length, i.e., the number of links traversed by a feasible connection. Intuitively a connection with a longer path is harder to establish, since an identical free wavelength channel must be found on more links.
- Smallest incremental used capacity, i.e., the additional network capacity that is consumed if the connection is incorporated in the current time slot, expressed in wavelength links. Apart from the wavelengths occupied by the connection on the links that it traverses, this includes also the wavelengths on the links that are wasted or blocked by the connection. Furthermore, it includes wavelength links that are wasted or blocked by already existing connections, due to the incorporation of the selected connection. By incorporating the connection that minimizes the incremental used capacity, the capacity left over for additional connections is maximized.
- Largest value of the maximum of the number of connections still to be transmitted by the source of the connection considered, and the number of connections still to be received by the destination of the connection considered. This is motivated by the fact that under certain conditions the network throughput is limited by the maximum number of connections originating from a node or the maximum number of connections destined to a node.
- In case there is more than one connection in the set of most eligible connections, the first one is selected.

Function *least-used-wavelength* determines a wavelength channel for the selected feasible connection. First, it determines the subset of wavelength channels that are free on all the links traversed by the connection. Then, the wavelength channel that has been assigned the least number of times during the current time slot is selected. This increases the chances that two connections with disjoint paths are assigned a different wavelength channel, and hence it reduces the chances of a collision when these paths are merged (later in the process) by the path of another connection.

The complexity of the algorithm can be determined by considering the worst case scenario, where at each iteration, none of the remaining connections conflicts with the current time slot. The complexity of the algorithm is thus $O(|L|^2 \cdot |E| \cdot W)$.

4. PERFORMANCE ANALYSIS

This section presents the results of an assessment of the performance that can be achieved by the proposed E-PSR network architecture, and using the WTA algorithm to construct the TDM link transmission frames. Dijkstra's shortest path algorithm⁶ is used to determine a priori the path of each connection. All links in the network are hereby assumed to have unit cost. The network performance is studied for two cases of randomly generated capacity demands: 1. a uniform capacity demand in which the network has to provide the same average capacity between each node pair, and 2. a non-uniform capacity demand in which the network has to provide a higher average capacity between a number of nodes designated as servers, and a lower average capacity between the remaining nodes (clients) and the servers. For both cases, fifty simulation runs are sufficient to obtain the average TDM transmission frame length with a confidence of 99% that the error is not more than 4%. The topology of the NSFNET backbone network,¹⁴ shown in Figure 4, is used as a benchmark network. Figure 4 depicts the switching nodes only. One end node with two tunable transmitters and two tunable receivers is connected at each switching node. All these end nodes have the same number of transmitters and receivers, in order to deal with uniform capacity demands. Switching nodes CA1, PA, and CO have a second end node with six transmitters and six receivers. These end nodes are the servers that have to deal with higher capacity demands in the non-uniform capacity demand case. Totally the network has 31 nodes (14 switching nodes plus 17 end nodes).

4.1. Performance Bounds

First some performance bounds are derived, which apply for fixed capacity demands. Let S_i , be the number of connections in the set L with end node i as source, and D_i the number of connections with destination i . Let L_k be the number of connections routed on link k . The following lower bounds for the frame length are easily established:

$$K_1^{(min)} = \max_{i \in D} \left\{ \left\lceil \frac{S_i}{l(i)} \right\rceil \right\}, \quad (5)$$

$$K_2^{(min)} = \max_{i \in D} \left\{ \left\lceil \frac{D_i}{r(i)} \right\rceil \right\}, \quad (6)$$




Figure 4. Topology of the NSFNET backbone network

$$K_3^{(min)} = \max_{k \in E} \left\{ \left\lfloor \frac{L_k}{W} \right\rfloor \right\}. \quad (7)$$

An overall lower bound on the frame length is:

$$K^{(min)} = \max \{ K_1^{(min)}, K_2^{(min)}, K_3^{(min)} \}. \quad (8)$$

4.2. Uniform Capacity Demand

First, the network performance is investigated under a uniform capacity demand, whereby the number of connections between each source and destination is drawn using a uniform distribution with an average of 1 connection. For identical source and destination the number of connections is zero. This gives an average total capacity demand of 272 ($17 \times 16 \times 1$) connections.

A sequence of experiments is executed, whereby in the first experiment all switching nodes are PSR nodes, and in the subsequent experiments PSR nodes are gradually replaced by IWS nodes, until the all-IWS network is reached in the last experiment. For choosing the switching nodes to be converted into IWS node, the following heuristic is applied. First a connection set is generated deterministically using the uniform capacity demand, but instead of drawing the number of connections from a uniform distribution, the number of connections is set equal to the average. After routing these connections, the subset of longest connections in terms of number of links is selected, and the number of times each switching node occurs as central node of a path is counted. A central node on the path of a connection has the property that the number of links on the path from source to central node and the number of links on the path from central node to destination differ by not more than 1. Next, the switching nodes are listed in decreasing order of central node count. When M nodes are to be chosen for upgrading to IWS node, the first M nodes of the ordered list are selected. This heuristic is motivated by the fact that long connections (which are more difficult to accommodate) pass more often through the same centrally located node than through some node that is more decentralized. Therefore, if that node is upgraded to IWS, more connections can benefit from its switching flexibility. For the NSFNET network and a uniform capacity demand, the ordered list of switching nodes is TX, IL, UT, WA, MI, CO, CA2, PA, CA1, NJ, MD, GA, NE, and NY.

Figure 5 shows the average frame length for the NSFNET network as a function of the number of wavelength channels, for increasing number of switching nodes converted from PSR to IWS. Since the capacity demand is generated randomly, the lower bound for the frame length, $K^{(min)}$, (Equation 8) is a random variable. Figure 5 shows also 50% and 90% worst case probability intervals for $K^{(min)}$. The graph shows that the all-PSR network achieves an average frame length that is not more than 16% longer than for the all-IWS network. It is striking to notice that by converting only three nodes into IWS the performance gap between the two extreme cases is already closed by 47%. When 9 out of 14 nodes are converted into IWS, the performance is almost the same (1.5% longer average frame length) as in the all-IWS case.




Figure 5. Average frame length for the NSFNET network as a function of the number of wavelength channels, for a uniform capacity demand with a total average of 272 connections

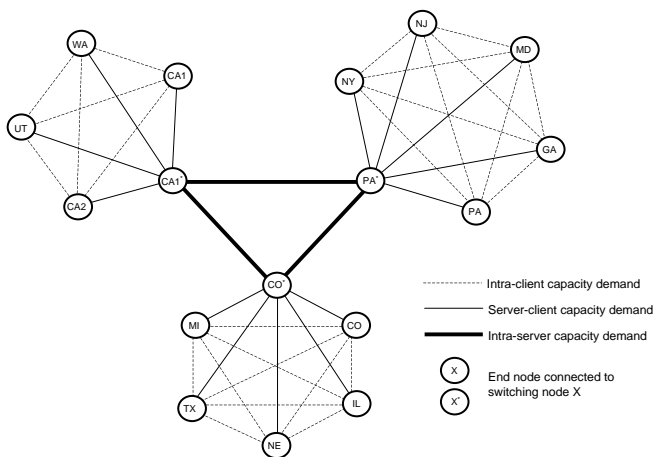


Figure 6. Server/client capacity demand in the NSFNET backbone

4.3. Non-Uniform Capacity Demand

To study the performance of the E-PSR network architecture under non-uniform capacity demands, a server/client capacity demand is used, that is represented in Figure 6. In this figure, nodes represent end nodes; switching nodes are not shown. The end nodes with six transmitters and six receivers, i.e., CA1*, PA*, and CO* act as servers. Each server communicates with a group of clients, whereby the capacity demand is asymmetric, i.e., the average capacity demand from server to client is 4 connections and the average capacity demand from client to server is 2 connections. In addition, there is an average demand of 2 connections (bidirectional) between each pair of clients associated with the same server. Also, the servers form a group with an average demand of 14 connections (bidirectional) between




Figure 7. Average frame length for the NSFNET network as a function of the number of wavelength channels, for a server/client capacity demand with a total average of 272 connections

each server pair. Actual numbers of connections are drawn using uniform distributions. The average total capacity demand is 272 connections, just as in the uniform capacity demand case.

The same heuristic as in Section 4.2 is used to determine the nodes that are upgraded from PSR to IWS. For the NSFNET network and the client/server capacity demand, the ordered list of switching nodes is IL, WA, TX, NE, GA, UT, CO, CA2, PA, NJ, CA1, MD, NY, and MI. Figure 7 shows the average frame length of the NSFNET network as a function of the number of wavelength channels, for increasing number of switching nodes converted from PSR to IWS.

The graph shows that the frame length does not decrease as fast with the number of wavelengths as in the uniform capacity demand case. This is because the lower bound on the frame length (Equation 7) is $\frac{48}{W}$, as opposed to $\frac{25}{W}$ in the uniform case. Furthermore, the performance gap between the all-PSR case and the all-IWS case is at most 8%, as opposed to 16% in the uniform traffic case. Explanation of this fact is that in the server/client case, a large number of the connections are between the server nodes that have six transmitters and receivers. Consequently, the server nodes are able to create full (or almost full) slots, resulting in reduced unusable transmission capacity due to wavelength wasting and blocking. Just as in the uniform traffic case, only 3 nodes need to be converted to IWS in order to close half of the performance gap, and 9 nodes need to be converted to close the performance gap almost completely.

5. CONCLUSION

This paper presented an architecture for all-optical WDM networks with arbitrary topologies, that offer optically transparent communication and a flexible allocation of the system capacity to the traffic demands. This architecture, called Enhanced Photonic Slot Routing (E-PSR), supports two types of switching nodes, i.e., Photonic Slot Routing (PSR) nodes, and Individual Wavelength Switching (IWS) nodes. The two node types provide different levels of functional flexibility on the one hand and cost/complexity on the other hand.

A heuristic scheduling algorithm has been presented that computes the composition, and the transmission time of the photonic slots for an E-PSR network employing TDM transmission control, given the network specification and the capacity demands, thereby maximizing the network throughput. Using the topology of the NSFNET backbone network, we showed that the performance of a network with only PSR nodes can be improved considerably by



upgrading a subset of nodes to IWS. The magnitude of the improvement depends strongly on the number and the choice of the specific nodes that are upgraded. Conversion of just one fifth of the nodes from PSR node into IWS node closes almost half of the performance gap between the all-PSR network and the all-IWS network. Conversion of two third of the PSR nodes virtually closes the performance gap completely.

The presented E-PSR network architecture is a flexible solution for accommodating growing capacity demands with minimal hardware upgrade. Only the nodes that are upgraded from PSR to IWS are effected and no other part of the network, needs to be changed. While exploring the network configurations derived by the various permutations of PSR and IWS nodes, a designer can use the scheduling algorithm to support him in the selection of the nodes that will be upgraded.

REFERENCES

1. E. Almström, C.P. Larsen, L. Gillner, W.H. van Berlo, M. Gustavsson, E. Berglind, "Experimental and Analytical Evaluation of Packaged 4x4 InGaAsP/InP Semiconductor Optical Amplifier Gate Switch Matrices for Optical Networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 14, no. 6, pp. 996-1004, June 1996.
2. I. Chlamtac, "Rational, Directions and Issues Surrounding High Speed Computer Networks," *IEEE Proceedings*, vol. 78, no. 1, pp. 94-120, January 1989.
3. 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, pp. 1171-1182, July 1992.
4. I. Chlamtac, A. Fumagalli, G. Wedzinga, "A Novel Approach for All-Optical Packet-Switching in Wide-Area Networks," *Proc. of SYBEN 98 Symposium on Broadband European Networks*, pp. 208-217, Zürich, Switzerland, 18-21 May 1998.
5. I. Chlamtac, A. Fumagalli, G. Wedzinga, "Slot Routing as a Solution for Optically Transparent Scalable WDM Wide-Area Networks," to appear in *Photonic Network Communications*.
6. E.W. Dijkstra, "A Note on Two Problems in Connexion with Graphs," *Numerische Mathematik*, vol. 1, pp. 269-271, October 1959.
7. P. Doussiere, "Recent Advances in Conventional and Gain Clamped Semiconductor Optical Amplifiers," *Proc. of Optical Amplifiers and their Applications, TOPS Vol. 5*, pp. 170-188, Monterey, CA, USA, 11-13 July 1996.
8. A. Franzen, D.K. Hunter, I. Andonovic, "A Low Loss Optical Packet Synchronization Architecture," *Proc. of SPIE, All-Optical networking: Architecture, Control, and Management Issues*, Vol. 3531, pp. 390-395, Boston, MA, USA, 3-5 November 1998.
9. 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.
10. P.E. Green, "Optical Networking Update," *IEEE Journal of Selected Areas in Communication*, vol. 14, no. 5, pp. 764-779, June 1996.
11. M. Gustavsson, B. Lagerstrom, L. Thylen, M. Janson, L. Lundgren, A.-C. Morner, M. Rask, B. Stolz, "Monolithically Integrated 4x4 InGaAsP/InP Laser Amplifier Gate Switch Arrays," *Electronic Letters*, vol. 28, pp. 2224-2225, November 1992.
12. Z. Haas, "The Staggering Switch: An Electronically Controlled Optical Packet Switch," *Journal of Lightwave Technology*, vol. 11, no. 5/6, pp. 925-936, May/June 1993.
13. R. Kannan, R. Bartoš, K.Y. Lee, F. Jordan, "STWnet: A High Bandwidth Space-Time-Wavelength Multiplexed Optical Switching Network," *Proc. of IEEE INFOCOM 97*, vol. 2, pp. 777-784, Kobe, Japan, 7-11 April 1997.
14. B. Mukherjee, S. Ramamurthy, D. Banerjee, A. Mukherjee, "Some Principles for Designing a Wide-Area WDM Optical Network," *Proc. of IEEE INFOCOM 94*, pp. 110-119, Toronto, Canada, 6-12 June 1994.
15. J.D. Walker, F.G. Patterson, S.P. Djajili, R.J. Deri, "A Gain-clamped, Crosstalk Free, Vertical Cavity Lasing Semiconductor Optical Amplifier for WDM Applications," *Proc. of Integrated Photonics*, pp. 474-477, Boston, MA, USA, 29 April - 2 May 1996.