Available online at www.sciencedirect.com
SCIENCE DIRECT®

Optical Switching and Networking 2 (2005) 100–112

Optical Switching and Networking

www.elsevier.com/locate/osn

Blocking performance of time switching in TDM wavelength routing networks[☆]

Mahesh Sivakumar[a,*], Suresh Subramaniam[b]

[a]Department of CSEE, University of Maryland, Baltimore County, Baltimore, MD 21250, United States
[b]Department of ECE, The George Washington University, Washington, DC 20052, United States

Received 16 February 2005; received in revised form 31 May 2005; accepted 6 June 2005
Available online 24 June 2005

Abstract

Advances in optical WDM technology have paved the way for high-capacity wavelength channels capable of carrying information at Gb/s rates. However, with current traffic streams requiring only a fraction of a wavelength's bandwidth, it becomes necessary to groom these independent low rate traffic streams on to higher capacity wavelength channels. An all-optical approach to grooming is to allow many connections to time-share a wavelength. Accordingly, in a TDM wavelength routing network, the establishment of a connection requires the assignment of time slots in addition to routing and wavelength assignment. One of the primary challenges in such networks is the need for quick reconfiguration at the routing nodes. In this paper, we investigate the effects of switch reconfigurability, wavelength conversion and time slot interchangers (TSIs) on the blocking performance of connections with multiple rates. Heuristics for time slot assignment that consider constraints imposed by six different node architectures are proposed, and the blocking performance of the TDM wavelength routing network is evaluated through simulations. Results indicate that limited reconfigurability at the nodes is sufficient to attain the performance obtained with full reconfigurability, especially when connections occupy only a small fraction of the wavelength capacity. Furthermore, the blocking performance is not seen to benefit significantly with the introduction of wavelength converters and TSIs, thus signifying that the improvement in blocking is largely dependent on the switch reconfigurability at the nodes.

Keywords: Optical networks; WDM; TDM wavelength routing; Reconfigurability; Wavelength converters; Time-slot interchangers; Architectures; Heuristics

1. Introduction

The widespread expansion in Internet services coupled with the ever-increasing demand for increased bandwidth capacity is beginning to apply brakes on the limits of conventional networking technology.

☆ This work was supported in part by NSF Grants ANI-9973111 and CNS-0434956.
* Corresponding author.
E-mail addresses: masiv1@cs.umbc.edu (M. Sivakumar), suresh@gwu.edu (S. Subramaniam).

1573-4277/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.osn.2005.06.001

In an effort to resolve the bandwidth crisis, network providers are looking to leverage the services of optical networks, networks based on the emergence of optical layer in transport networks, for reduced costs and higher capacity. Recent advances in optical switching and in particular *wavelength division multiplexing* (WDM) have enabled next generation networks to be able to operate at several terabits per second. The shortcomings of electronic or optoelectronic networks can be overcome by using all-optical or wavelength routed networks, which maintain the signal in optical form throughout the transmission.

Wavelength-routing optical networks consist of optical switching nodes interconnected by one or more fiber links. Data is transported in the network over *lightpaths*, which are all-optical communication paths. The procedure of setting up a lightpath between any source-destination pair involves choosing an appropriate route, and then reserving a wavelength on each link of the selected route. This is referred to as *Routing and Wavelength Assignment* (RWA). There are certain drawbacks with such a lightpath-based approach. Firstly, if there is no direct lightpath between two nodes, there needs to be intermediate store and forward node(s). However, these nodes limit the maximum throughput obtainable, due to the electronic processing and the buffer management involved. Secondly, most of the traffic streams supported by the network are expected to require only a fraction of the per-wavelength capacity. It is, hence, imperative that these low-rate traffic streams be multiplexed on to high-capacity wavelength channels and provisioned as a single lightpath, to improve resource utilization. The second issue gives rise to the concept of *traffic grooming*, which encapsulates techniques designed to efficiently groom low-rate traffic streams on to available wavelengths so as to increase the network throughput and utilize resources efficiently.

An all-optical approach to traffic grooming is to use TDM to divide time into frames of slots on each wavelength and multiplex several low-rate streams on to the high-capacity wavelength channel. The routing nodes here must be capable of routing wavelengths as well as time slots, i.e., the routing patterns of nodes must be configurable to switch on a time slot basis [1,2]. In addition, the complicated optical buffering

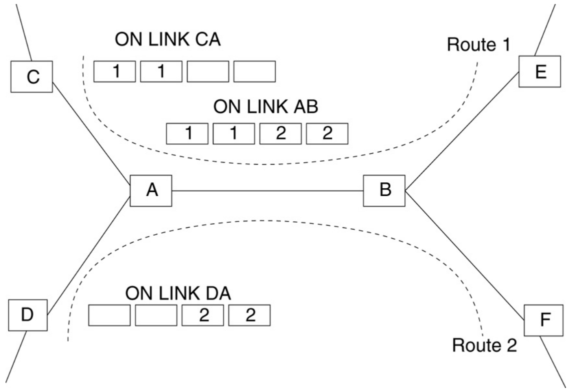

Fig. 1. Performance improvement with TDM wavelength routing networks.

and optical contention resolution schemes are not needed at the nodes [2]. All-optical routing avoids any O/E/O conversion and offers transparency within the time slots assigned to a connection. This scheme thus overcomes the drawbacks of traditional wavelength routing networks and avoids the difficulty involved in implementing optical packet switches. Routing is now done on the basis of the time slot a signal arrives on, along with the input port and wavelength. Hence, in addition to RWA (as done with traditional wavelength-routed networks), time slot assignment for each connection should also be performed. This problem has been referred to as Routing, Wavelength and Time slot Assignment (RWTA) [3].

1.1. Motivation and contributions

We consider a circuit-switched *time division multiplexed (TDM) wavelength routing network* [1] in which time is divided into frames of slots on wavelengths and the number of time slots per frame, T, is the same on all the wavelengths. Fig. 1 gives an illustration of the performance improvement possible with such networks. Each fiber link is assumed to have two wavelengths and time on each wavelength is assumed to be divided into frames of four slots each. The figure shows two lightpaths that use a link AB. Each of the lightpaths is assumed to request two slots. With traditional wavelength routing, this would require two wavelengths on the link AB. However, by using TDM wavelength routing to route these lightpaths, only one wavelength would be required, as shown in the figure. It is quite obvious that the latter is more likely to utilize the resources efficiently. However, the improved performance is obtained at the expense of increased network cost caused by the added flexibility and complexity at the routing nodes. Our goal is to determine if the improvement in performance is worth the additional cost.

The main contributor to the added complexity, and therefore the additional cost, relates to the design of the routing nodes that would allow quick reconfiguration of the switches at the nodes. In the worst case, the switches may be required to reconfigure after every time slot. Related work on RWTA [1–4] has assumed the switching node to be able to reconfigure after every slot in the wavelength frame. Such a switch will contribute more to the increase in cost and complexity of the routing nodes than a switch that can, say, reconfigure only once or twice within the frame. We define *switch reconfigurability* (R) as the number of times the switch at the node needs to reconfigure within a single wavelength frame. It not only determines the speed at which the optical switch at each node needs to operate but is also an indicator of the increase in switch cost and complexity required to improve the blocking performance. $R = 0$ implies that the switches are reconfigured only at connection setup time. On the other extreme, $R = T - 1$ corresponds to the most flexible case in which the switches may be reconfigured at the end of each time slot. Thus, the motivation behind this paper is to identify the amount of reconfigurability required at the routing nodes that would improve the blocking performance with a modest increase in cost. In short, we investigate the trade-off between the performance offered by reconfigurability and the cost involved in building such a routing node architecture.

The performance improvement with the available switch reconfigurability depends on the RWTA algorithm, and in particular the slot assignment algorithm used. With much work done in routing and wavelength assignment, we focus on time slot assignment in TDM wavelength routing networks to study the effect of reconfigurability provided by the routing nodes on the blocking performance of the network. Besides reconfigurability, we also study the effect of (i) *the constraint on number of wavelengths used per connection*, (ii) *wavelength conversion* and (iii) *time slot interchangers* (TSIs) on the blocking performance. We only consider *permutation switches*, i.e., switches that can be configured to route one input at a time to each output. Towards this goal, we propose six different network models for time slot assignment and use discrete event simulation to analyze their blocking performance. The models differ in the amount of flexibility offered to time slot assignment. Simulation results indicate that the time slot based approach would give a marked improvement in network throughput through decreased blocking probabilities. More importantly, a limited amount of switch reconfigurability is seen to be sufficient to obtain significantly improved blocking performance, indicating that full reconfigurability is not necessary at each node to obtain the best performance. The blocking improvement with the introduction of time slot interchangers is seen to be more pronounced when the average number of slots required by the connections is higher.

The rest of the paper is organized as follows. Section 2 presents the architecture of a TDM wavelength routing network with a summary of related work. The proposed models for time slot assignment and their performance benefits are presented in Section 3. Section 4 contains the performance evaluation of the proposed schemes and Section 5 concludes the paper.

2. Network architecture and related work

This section describes the architecture of a TDM wavelength-routed network followed by a summary of previous research efforts in this area.

2.1. Network architecture

The network architecture consists of TDM wavelength routers interconnected by bidirectional links. The performance improvement with TDM wavelength routing networks is made possible by equipping the crossconnects with switches that are capable of quick response to inputs, i.e., switches that can reconfigure at the granularity of a time slot. As far as setting up lightpaths is concerned, this network performs similarly to a circuit switched wavelength routing network in that the lightpaths are set up prior to transmission. However, each connection request now contains the

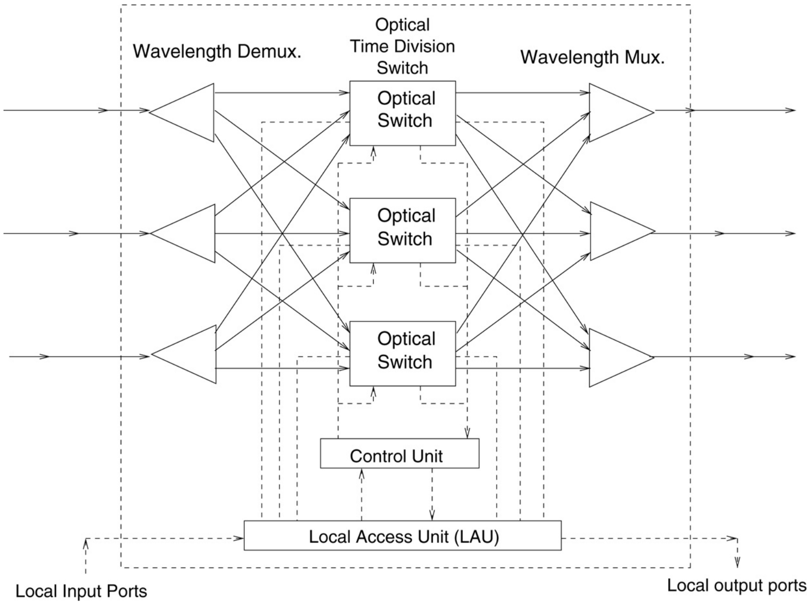

Fig. 2. Network architecture of a TDM wavelength routing node [2].

required bandwidth in terms of the number of time slot along with the source and destination nodes. The architecture of an $N \times N$ TDM wavelength router is shown in Fig. 2 [2].

The input ports are provided with wavelength demultiplexers that separate the wavelengths. These wavelengths are sequentially concentrated at an $N \times N$ optical time division switch. The N outputs of the switch are then connected to N $N \times 1$ multiplexers, respectively, before finally being combined on to an output port. There is a control unit that sets up the connectivity between the input and output ports and is capable of doing so on a time slot basis. The exact number of slots on the wavelength frame on which different connections can be multiplexed depends on the reconfigurability of the switch at the router. The control unit is also responsible for tracking and calibrating the local signal for synchronization. We assume that the total synchronization and propagation delay experienced by the slot traffic is an integral multiple of the frame length and is taken care of by the synchronizers. As a result, we ignore the delay in this study. Furthermore, due to the connections being circuit switched, the possibility of two slots needing to be switched to the same output slot at a node does not arise.

In such a network, we can have the flexibility of establishing a connection request on different slots of a wavelength(s) on each link if the node has a time slot interchanger (TSI). This would have an obvious improvement in the network throughput at the expense of the TSI.

2.2. Related work

A detailed survey of work on the electronic traffic grooming problem that has gained significant attention in recent times can be found in [5]. In the TDM wavelength routing network we consider, along with the route and the wavelength, time slots should also be assigned to the connection. Due to time being slotted on each wavelength, the space switching at one input port of the node should be synchronized with the others, i.e., the incoming time slots must be aligned before space switching. The possibility of such a synchronizer was described in [6], where the authors give an input synchronizer to align incoming slot data.

There has also been considerable work done on routing node architectures required to achieve reconfigurability. It is said that allowing guard times between time slots and increasing the duration of slots can be used to mitigate the effects of slow reconfiguration times [7]. In [8], the authors demonstrated the feasibility of the TDM wavelength routing network with tunable transceivers and a non-reconfigurable wavelength router. Addressing is achieved by launching the data on an appropriate wavelength where each wavelength frame was assumed to be 250 μs long with capability to accommodate 128 slots. In [2], the authors came up with a new architecture for optical transport networks based on their proposed Time–Wavelength–Space Routers (TWSRs) which are the same as the TDM-based wavelength routers we consider. They also proposed a heuristic for the problem of establishing the set of efficient time slot based lightpaths for a given set of connection requests and studied the effectiveness of such a network in terms of the throughput obtained.

The blocking performance of dual-rate connections (one and two slot requests) has been analytically modeled in [9,10]. The authors presented several architectures for channel switching, and developed analytical models for computing blocking probabilities for the various architectures. In [4], blocking probabilities are calculated for single-rate connections (i.e., all connections require one slot in a frame) under dynamic traffic, with random assignment of time slots and wavelengths. They also studied the effects of wavelength converters and TSIs on blocking probability. In [11], a fast, greedy, slot scheduling heuristic for transmitting packets is presented.

In [1], the authors considered the problem of scheduling multirate connections in TDM wavelength-routing networks. In particular, they addressed the off-line multirate connection scheduling problem, i.e., the problem of assigning time slots and wavelengths to a given static set of multirate connections, in ring topologies. They showed that the off-line single-rate connection scheduling problem is similar to the off-line wavelength assignment problem and find bounds on the frame length to maximize the network throughput. The problem of routing, wavelength and time slot assignment (RWTA) in TDM wavelength routing networks was studied in [3] with a goal of maximizing the overall throughput in the network. They proposed two algorithms for time slot assignment, namely, first-fit and least loaded time slot. While first-fit selects the first available free slot on each link of the lightpath, the least loaded time slot algorithm selects the least loaded time slots over all the links of the path. Discrete event simulations showed the superiority of the least loaded RWTA algorithm in terms of the amount of traffic supported for a given target blocking probability. Furthermore, the gains were seen to improve with increase in the number of fibers per link.

To the best of our knowledge, all earlier work on RWTA has assumed that each switching node is capable of reconfiguring after every time slot (fully reconfigurable nodes). As mentioned earlier, such a node architecture would be expensive in terms of the cost and complexity involved, and the performance improvement obtained may not be worth the additional cost. In this paper, we study the blocking performance of TDM wavelength routing networks with limited node reconfigurability with a view to study the cost-performance trade off. We also study time slot assignment and the effect of wavelength conversion and TSIs on the performance of such a network. As will be seen, the node reconfigurability tends to dominate the blocking performance and, more importantly, limited reconfigurability at each node is sufficient to obtain similar performance benefits to that provided by fully reconfigurable nodes.

We next present the various network models for time slot assignment and the wavelength and slot assignment algorithms evaluating the performance of a TDM wavelength routing network.

3. Network models and slot assignment algorithms

In this section, we describe the various network models we consider for wavelength and slot assignment (WSA) in studying the effects of factors such as node reconfigurability, wavelength conversion, and time slot interchanging on the blocking performance of TDM wavelength routing networks.

3.1. Assumptions and definitions

We consider circuit-switched networks under dynamic traffic where the connection requests arrive at a node according to a Poisson process and exist in the network for a random amount of time (exponentially distributed). During this time, an all-optical circuit is set up between the source and the destination of the lightpath request. For any two successive links of a lightpath, the same wavelength must be selected if the intermediate node does not have a converter and a different wavelength may be selected if the intermediate node has a wavelength converter. There are N nodes and L links in the network and each link has the same number of wavelengths W (numbered λ_1 through λ_W). The path for a lightpath request between a given pair of nodes is fixed and assumed to be a shortest path between the nodes. All the wavelengths have a fixed number of time slots (T). The number of time slots required by the lightpath requests (t_{req}) is assumed to be *uniformly distributed* with mean t and defined as follows.

- For $1 \leq t \leq T/2$, we have $t_{\text{req}} \sim \text{unif}(1, 2t - 1)$.
- For $T/2 \leq t \leq T$, we have $t_{\text{req}} \sim \text{unif}(2t - T, T)$.

We combine wavelength and slot assignment and use the first-fit algorithm explained later in this section. The slot assignment strategy depends on the network model used and the amount of reconfigurability available at the nodes and is designed to make use of the flexibility provided by each model. The following definitions will be used in the illustration of the time slot assignment algorithm.

- **Frame Length** (T): The number of slots in a single wavelength frame (numbered from 0 to $T - 1$).
- **NumSlotsReq** (t_{req}): The number of slots needed for a given randomly generated connection.
- **Switch Reconfigurability** (R): The number of times the switch at a node can reconfigure within a single wavelength frame. R takes values between 0 and $T - 1$ and is provided as input.
- **Block** (b): Set of contiguous slots between two points in a wavelength frame where the switch can reconfigure. For instance, if for a wavelength $\lambda_1, T = 10$ and the switch is designed to reconfigure after slot 3 and slot 6, [0–2], [3–5], and [6–9] would constitute its blocks. The cardinality of a block b $(|b|)$ is the number of slots within b.
- **BlockSet** (B): Set of blocks in a wavelength. All the blocks within a wavelength are chosen to contain approximately the same number of slots. This value depends on the frame length (T) and switch reconfigurability (R). For instance, if $T = 20$ and $R = 1$, the switch will be designed to reconfigure after slot 9 and B will contain the blocks [0–9] and [10–19]. The blocksets for the wavelengths are calculated offline.
- **Free Block** (FB): A block is said to be free if none of the slots of the block is allocated to any connection. There is an exception to this condition. Assume that a portion of a block is allocated to a connection between nodes x and y. Now, if there is a new connection request between the same *s–d* pair (x, y) and the block has a sufficient number of free slots to accommodate it, then the block is said to be free. This is because connections between the same *s–d* pair can share a block on a wavelength since they do not require the switch to reconfigure within the block.
- $FB_l(i)$: The set of free blocks for wavelength i on link l.
- $NS_l(i)$: The number of slots that are free for wavelength i on link l. If there are no TSIs in the network,

  $$NS_l(i) = \sum_{b \in FB_l(i)} |b|,$$

  and if every node has TSIs, this is just the number of free slots[1] for wavelength i on link l.
- $FB_p(i)$: The set of free blocks for wavelength i on path p.

  $$FB_p(i) = \bigcap_{l \in p} (FB_l(i)).$$

- $NS_p(i)$: The number of slots that are free for wavelength i on path p. If there are no TSIs in the network,

  $$NS_p(i) = \sum_{b \in FB_p(i)} |b|,$$

  and if every node has TSIs,

  $$NS_p(i) = \min_{l \in p} (NS_l(i)).$$

We first present the proposed network models with the corresponding wavelength and slot assignment strategies and then illustrate with examples the potential benefits from these architectural solutions.

3.2. Proposed network models

We consider six different models of a TDM wavelength routing network and examine their effect on the blocking probability of connection requests.

[1] A slot is free if it is not allocated to any connection.

The models differ in the flexibility they provide to wavelength and slot assignment. Note that the models considered are in increasing order of network cost and complexity. The switch reconfigurability, R, is a variable parameter in the first four models.

1. Single wavelength with no wavelength conversion (SWNC):

Under this scenario, the lightpath requests are restricted to a single wavelength in each link of their path, i.e., all the slots required by a lightpath request are to be allotted on a single wavelength. Moreover, due to lack of wavelength converters, every lightpath needs to be established on the same wavelength on all the links of its path. This type of restriction simplifies the node architecture but, on the other hand, is bound to increase the blocking probability due to the wavelength continuity constraint, and, more importantly, the restriction imposed by the need to find all the requested slots on a single wavelength.

For *wavelength and slot assignment* (WSA), we find the first wavelength[2] that satisfies the connection request in terms of the number of slots required. i.e., we find the first wavelength i such that $NS_p(i) \geq t_{\text{req}}$. In particular, when a connection comes in to the network, beginning with the wavelength that has the lowest index, the following are done:

- Check if each block in the wavelength frame beginning with first block is available on all the links of the path.
- Stop when either the sum of the cardinalities of free blocks for a wavelength i on path p ($NS_p(i)$) is greater than the number of slots required by the connection or if we reach the end of the wavelength frame (which indicates that the wavelength does not have enough free slots to fit the connection).
- If the wavelength does have enough free slots, select the wavelength and stop. Otherwise continue with the next wavelength.

If a wavelength with sufficient free capacity is found, the appropriate blocks (beginning with the first free block within the frame and ending when all the slots required have been allocated) on the wavelength are set to be busy on all the links of the path. If no such wavelength is found, the connection is blocked.

2. Single wavelength with full wavelength conversion (SWFC):

This is similar to SWNC except that every node in the network is assumed to have unlimited full wavelength converters. Full wavelength conversion is likely to improve the blocking performance by relaxing the wavelength continuity constraint, however, at the expense of increased network cost due to wavelength converters. The motivation behind such a model is to see whether the improvement in blocking performance is worth the additional cost of wavelength conversion. Furthermore, it allows us to study the performance improvement in blocking that could be obtained by imposing only the single wavelength restriction, as connections will not be blocked due to lack of wavelength converters.

WSA here is done as follows:

- As with SWNC, check if each block (beginning with the first block) is available on all the links of the path. The difference here is that the blocks can be free on any wavelength on each link of the path due to the availability of wavelength conversion at all nodes. Note that on each link, we begin searching for the requested block beginning at the first wavelength.
- Stop when either the sum of the cardinality of free blocks is at least the number of slots required by the connection or if we reach the end of the wavelength frame (which now indicates that the connection cannot be established).
- If a wavelength with sufficient free capacity is found on each link of the path, the appropriate blocks (beginning with the first free block within the frame and ending when all the slots required have been allocated) on the wavelength are set to be busy on all the links of the path.

If B_s represents the set of blocks selected (such that $\sum_{b \in B_S} |b| \geq t_{\text{req}}$), the algorithm would find the first wavelength (wavelength with the lowest index) on each link of the path that has these blocks free.

3. Multiple wavelengths with no wavelength conversion (MWNC):

Here, the slots requested by an arriving lightpath may be allotted on more than one wavelength within a single link. Thus, this case gives more flexibility in

[2] In the rest of this section, first wavelength refers to the wavelength with the lowest index when wavelengths are ordered in a fixed way. Similarly, first block refers to a block containing slots of the least index, i.e., 0, 1, ...

slot assignment and is likely to improve the blocking performance as compared to SWNC. Here again, no conversion is allowed and hence the wavelength continuity constraint may play a significant role in determining the blocking performance. Note that due to the single wavelength restriction being removed, the source node may need to split the traffic into multiple streams to allow it to be carried on multiple wavelengths and the destination node would have to combine them back to obtain the original traffic stream. This is likely to increase the complexity of the nodes that originate and terminate traffic, and therefore the cost of the network. The previous two cases (SWNC and SWFC) avoid this additional complexity and cost.

In this model, WSA is done as follows. As with SWNC, we find the set of free blocks for each wavelength i on the path p ($FB_p(i)$) beginning with the first wavelength. However, we cumulatively add $NS_p(i)$, the number of free slots for each wavelength on the path. We stop when the number of free slots available on the wavelengths is cumulatively greater than t_{req}. In other words, for each link l of the path, we find a set of one or more wavelengths (W_l) (as compared to a single wavelength in SWNC) such that $\sum_{i \in W_l}(NS_p(i)) \geq t_{\text{req}}$. It is important to note that the same set of wavelengths will be used on each link of the path due to absence of wavelength conversion.

4. Multiple wavelengths with full wavelength conversion (MWFC):

This is an extension to MWNC done with a view to improve the blocking performance by removing the wavelength continuity constraint. Compared to MWNC, this model would increase the network cost due to every node being provided with wavelength converters.

WSA in this model is similar to that in SWFC but the process may be repeated up to W times, where W is the number of wavelengths on a link. WSA is done as follows. During the first round, as with SWFC, we check if each block (beginning with the first block) is available on all the links of the path. The blocks can be free on *any* wavelength on each link of the path. However, unlike SWFC, when we reach the end of the frame and the requested number of slots is not available, we start the process again with the first block, cumulatively adding the number of slots found free in each round. During subsequent rounds, we avoid choosing the wavelengths that were selected in the previous round(s). This process continues until the requested bandwidth is allotted or we go through W rounds and we do not find enough capacity, in which case the connection will be blocked.

5. TSI with no wavelength conversion (SWNC-tsi and MWNC-tsi):

The previous four models had one common restriction. The slots occupied on a wavelength (or on multiple wavelengths) were required to be the same throughout the path, i.e., the same set of slots had to be assigned to a lightpath on all the links of the path. We term this limitation the *slot continuity constraint*. This limitation can however be relaxed by the use of a TSI. As mentioned earlier, the TSI allows different sets of slots to be assigned on different links of a lightpath. In this model, we assume TSIs to be present at all nodes and examine the performance when connections can be established on single and multiple wavelengths, respectively. Furthermore, we assume no wavelength conversion at all nodes in the network. TSIs do not make sense unless the switches are reconfigurable after every slot since any slot may require time slot interchanging, and hence we assume that $R = T - 1$. This model is likely to provide better blocking performance than the previous ones at the expense of increased network cost (due to TSIs).

WSA for this model differs from the previous ones in two aspects: (i) Since we assume that $R = T - 1$, each wavelength contains T blocks and each block in a wavelength consists of one slot, and (ii) the same set of slots need not be free on each link of the path for the chosen wavelength(s). For SWNC, we find the first wavelength that satisfies the connection request in terms of the number of slots required, i.e., we find the first wavelength i such that $NS_p(i) \geq t_{\text{req}}$. In particular, when a connection comes in to the network, we do the following for each wavelength i, starting with the first wavelength.

- Find the number of free slots on each link l of the path ($NS_l(i)$). Note that the slots need not be the same on each link.
- The number of slots for the wavelength that are free on path p is then calculated as $NS_p i = \min_{l \in p}(NS_l(i))$.
- If $NS_p(i) \geq t_{\text{req}}$, we select wavelength i and stop.
- Otherwise, we repeat with the next wavelength.

- If no such wavelength is found, the connection is blocked.

For MWNC-tsi, we calculate the number of slots for each wavelength i that are free on path p $(NS_p(i) = \min_{l \in p} NS_l(i))$. However, if the requested bandwidth cannot be found on a single wavelength, we find a set of one or more wavelengths (W_l) on each link of the path (searching the wavelengths in increasing order), such that $\sum_{i \in W_l}(NS_p(i)) \geq t_{\text{req}}$. It is important to note that the same set of wavelengths will be used on each link of the path due to absence of wavelength conversion.

6. TSI with full wavelength conversion (SWFC-tsi and MWFC-tsi):

This is similar to the previous model with the exception that all nodes are assumed to provide unlimited wavelength conversion. This model is expected to give the best blocking performance and contribute the most to network cost and complexity since it offers the maximum flexibility for WSA. In case of SWFC-tsi, we find the first wavelength i on each link l of the path (the wavelength with the lowest index) such that $NS_l(i) \geq t_{\text{req}}$. For MWFC-tsi, we find a set of wavelengths W_l on each link l such that $\sum_{i \in W_l}(NS_l(i)) \geq t_{\text{req}}$.

We next present the simulation results and discuss the implications.

4. Performance evaluation

In this section, we present the results of our performance evaluation.

4.1. Simulation model

We consider the 14-node NSFNet topology for our performance evaluation. We also obtained some results for the European Optical Network (EON) but since they showed a similar trend in results to what we obtained using the NSFNet, we have not included them in the paper. Each link is bidirectional and is implemented as two oppositely directed fibers. Each simulation is run for 1 million connection requests. Connection arrival rates between node pairs are in proportion to the packet traffic intensities recorded and reported in [12]. The actual arrival rates on the routes are obtained by using a scaling factor γ (10^{-6}) with the values used in [12]. The mean holding times of connections are set to 1 unit. In our results, we choose the scaling factor in such a way as to keep the average blocking probability approximately in the neighborhood of 10^{-3}. Converters are assumed to have full conversion capability. Each wavelength frame is assumed to consist of 20 slots, i.e., $T = 20$. Note that the case where every connection requests T slots simulates wavelength routing. We measure the connection blocking probability which is defined as the ratio of the total number of connections accommodated to the total number of incoming connections.

4.2. Single wavelength time slot assignment

This section illustrates the results obtained for SWNC, where connection requests are allotted slots only on one wavelength for each link of its path. We first plot the blocking probability against the load (defined as the average number of slots required (t) by a connection request), in Fig. 3(a) for various values of R. Note that $R = 0$ does not imply unslotted wavelength routing since multiple connections between the same node pairs could still be multiplexed on the same wavelength unlike in wavelength routing, although the performance difference could be very small.

It can be seen from the graph that there is a marked improvement in blocking probability when connections require very few slots. This can be noticed for all values of R, the reconfigurability factor. As expected, with increase in the load, the blocking probability also increases. A sharp increase in blocking can be noticed for $R = 3$, in the neighborhood of $t = 6$. This indicates that the TDM wavelength routing technique is likely to accommodate more sessions, and the improvement in blocking is likely to be significant especially in cases where connections occupy a small fraction of the wavelength capacity.

It is interesting to note the improvement in blocking obtained when switches are reconfigurable. With a small amount of reconfigurability (R), there is a marked improvement in blocking performance. For example, we can see that the blocking probability drops from 10^{-1} to 10^{-6} when R is increased to 3 from 0, for $t \leq 6$. However, as R is increased further,

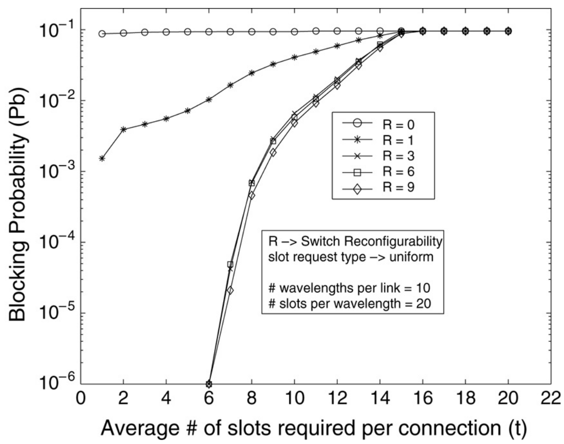

(a) Blocking performance (P_b) versus average number of slots required (t).

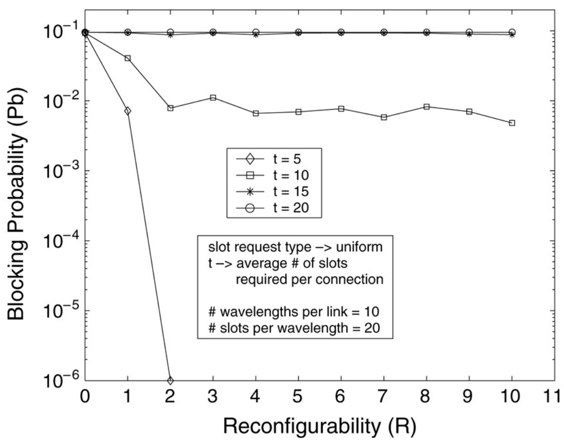

(b) Blocking probability (P_b) versus reconfigurability (R).

Fig. 3. Blocking performance for single wavelength with no conversion (SWNC).

the additional improvement is only marginal. It is further interesting to note the increase in load obtained at fixed P_b. For instance, in Fig. 3(a), it can be noticed that for a blocking probability of approximately 10^{-3}, $R = 3$ is sufficient to support an increase in load that is seven times the load obtained with $R = 1$. These factors signify the fact that we may not need maximum reconfigurability (where switches need to be reconfigurable after every slot) at the nodes to obtain the best blocking performance. This can also be clearly noticed in Fig. 3(b) where the blocking performance is plotted against the reconfigurability provided at the node (R) for varying loads. The curve for $t = 20$ represents the wavelength-routing case where connections are given an entire wavelength capacity. This case does not benefit at all from increase in reconfigurability, as can be seen. The slight increase in blocking for certain cases (e.g., curve for $t = 10$) with increased reconfigurability is due to the following: when the slots required per connection request exactly fit the number of slots after which the switch may be reconfigured, blocking is reduced as all the connection requests tend to fit perfectly in to the wavelength.

4.3. Multiple wavelength time slot assignment

This section presents the results obtained for MWNC, where the slots of a connection could be allotted on more than one wavelength in each link of its path. The additional flexibility provided to connection requests is likely to translate into better blocking performance than what could be obtained with SWNC. We first analyze the blocking performance for varying loads in Fig. 4(a) for the case when all nodes lack wavelength converters.

A similar trend to that seen in the single wavelength case (SWNC) is seen here too. Here again, the performance improvement is more pronounced when the number of slots required by the connection requests are close to or below one third of a wavelength capacity. The improvement with reconfigurability is similar to that seen in SWNC and is plotted in Fig. 4(b) for the case when no wavelength conversion is present in the network. However, compared to SWNC, the overall blocking performance with MWNC is better at all times, as expected. This is illustrated in Fig. 4(c), where the blocking performances of both cases (MWNC and SWNC) are compared for the case when $R = 6$ (equivalent to 30% reconfigurability). We also note that the performance improvement obtained with the use of multiple wavelengths is much less than what is obtained due to reconfigurability.

4.4. Improvement with wavelength conversion

This section illustrates the blocking performance of the TDM wavelength routing network for MWFC, where all the nodes in the network are assumed to provide unlimited wavelength conversion. A similar trend was seen (not shown here) when comparing

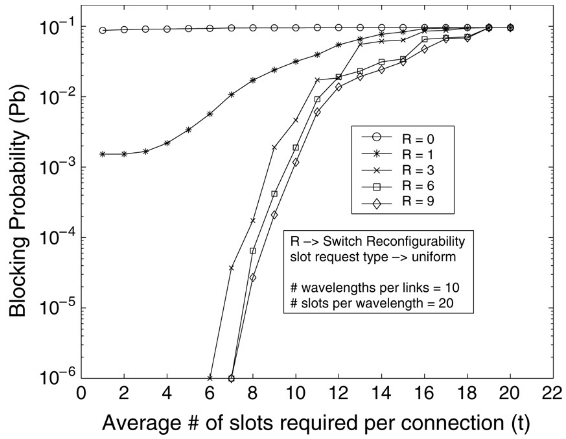

(a) Blocking performance (P_b) versus average number of slots required (t).

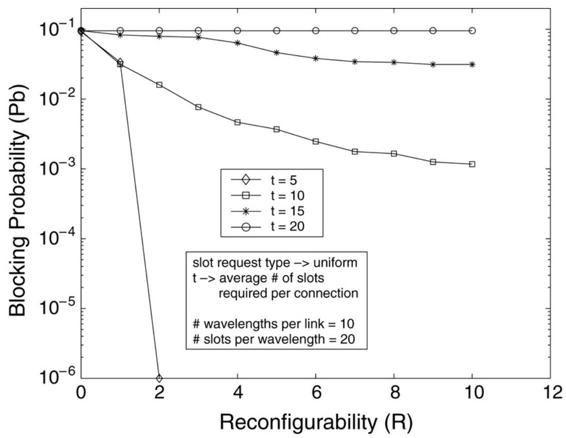

(b) Blocking probability (P_b) versus reconfigurability (R).

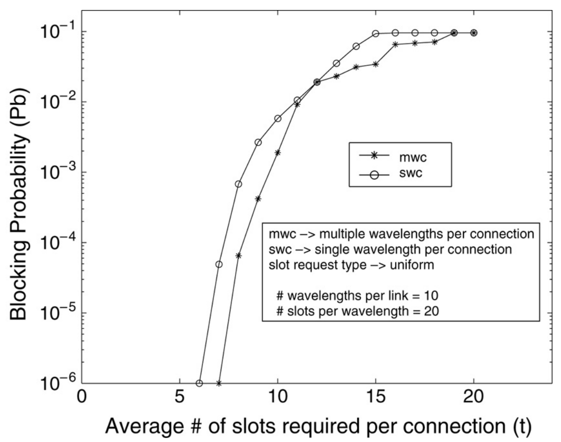

(c) Blocking performance comparison: MWNC versus SWNC.

Fig. 4. Blocking performance for multiple wavelengths with no conversion (MWNC).

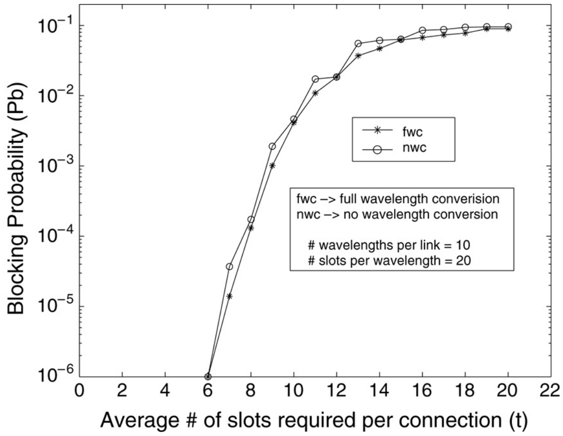

(a) Blocking performance (P_b) versus average number of slots required (t) with 30% reconfigurability.

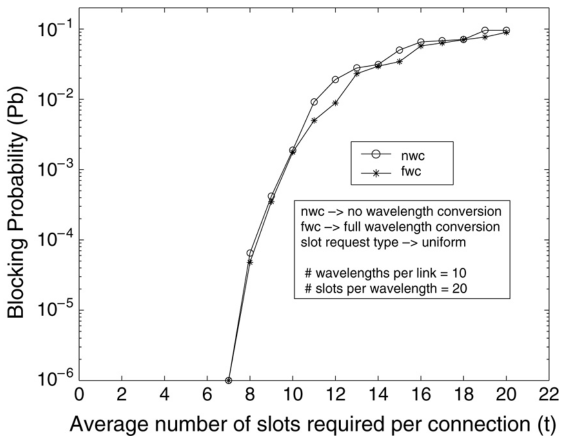

(b) Blocking probability (P_b) versus average number of slots required (t) with 15% reconfigurability.

Fig. 5. MWFC versus MWNC — Blocking performance.

the performances of SWNC and SWFC. We compare the blocking performance for MWFC with MWNC for the case when $R = 6$ (equivalent to 30% reconfigurability). This is illustrated in Fig. 5(a).

Improved blocking performance can be noticed when unlimited wavelength conversion is present at all nodes. However, the improvement is very marginal. This suggests that the blocking performance is dominated more by the flexibility (amount of reconfigurability at the nodes) provided to time slot assignment. The improvement in blocking probability from wavelength conversion is dependent only on the performance of the wavelength assignment algorithm used and not on the time slot assignment strategy used. Since the same wavelength assignment (first-fit) is used in both MWNC and MWFC, the improvement is very subtle with similar results obtained most of the time. The same can be witnessed in Fig. 5(b) for the case when $R = 3$ (equivalent to 15% reconfigurability).

4.5. Improvement with time slot interchangers

This section analyzes the performance improvement with time slot interchangers for both single and multiple wavelength models. We now consider the case when time slot interchangers are present at all nodes and compare its blocking performance with the case where none of the network nodes have time slot interchangers. This is illustrated for the SWFC case in Fig. 6(a), where nodes are assumed to provide unlimited wavelength conversion.

The results show a similar trend to what was seen in the previous cases. An interesting point to note though is that TSIs do not provide much improvement compared to the case without them when the average number of slots required is small. The improvement is more pronounced as the average number of slots required increases, but is still not significant. Similar results are also obtained in the case with time slot interchangers at all nodes for the multiple wavelengths case (MWFC) as shown in Fig. 6(b). The extent to which the performance of TSIs depends on load needs further investigation.

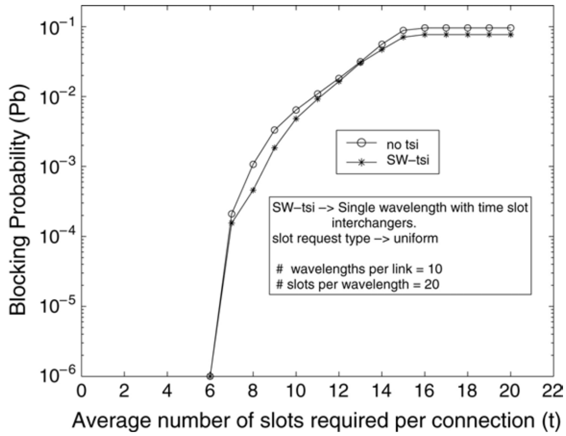

(a) Blocking performance (P_b) versus average number of slots required (t) for the single wavelength case.

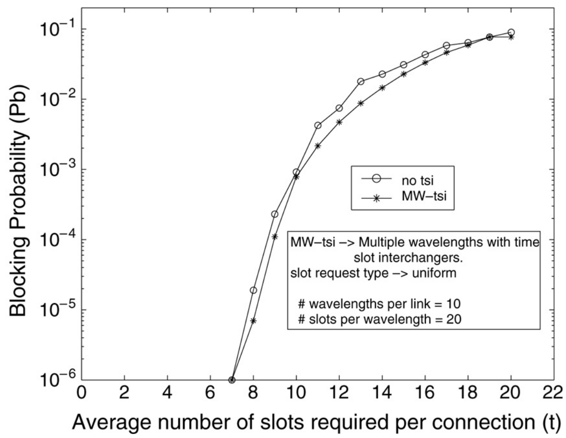

(b) Blocking probability (P_b) versus average number of slots required (t) for the multiple wavelengths case.

Fig. 6. Blocking performance with time slot interchangers (TSI).

5. Conclusions

We studied the blocking performance of TDM wavelength routing networks, via simulation, for varying degrees of reconfigurability at the nodes. We investigated six different scenarios based on the switch reconfigurability, constraints on the number of wavelengths used by a connection on each link, and the effect of time slot interchangers, and proposed heuristics for time slot assignment. Simulation results indicate that TDM wavelength routing is advantageous as compared to wavelength routing, especially when connections occupy a small fraction of the wavelength capacity. Interestingly, limited reconfigurability was seen to be sufficient to achieve the performance improvement obtained with full reconfigurability, suggesting that the overall network cost to obtain the benefits with TDM wavelength routing does not increase significantly. Furthermore, the improvement with time slot interchanging, wavelength conversion and the use of multiple wavelengths on a link per connection was not very significant, indicating that the blocking performance is dominated by the switch reconfigurability at the nodes.

References

[1] S. Subramaniam, E.J. Harder, H.-A. Choi, Scheduling multi-rate sessions in TDM wavelength-routing networks, in: Proc. GLOBECOM, 1999, pp. 1466–1472.

[2] N.F. Huang, G.H. Liaw, C.P. Wang, A novel all-optical transport network with time-shared wavelength channels, IEEE J. Sel. Areas Commun. 18 (10) (2000) 1863–1875.

[3] B. Wen, K. Sivalingam, Routing, wavelength and time-slot assignment in time division multiplexed wavelength-routed networks, in: Proc. INFOCOM, New York, NY, 2002.

[4] J. Yates, J. Lacey, D. Everitt, Blocking in multiwavelength TDM networks, in: Proc. 4th Intl. Conf. Telecommunication Systems, Modeling, and Analysis, 1996, pp. 535–541.

[5] K. Zhu, B. Mukherjee, A review of traffic grooming in WDM optical networks: Architectures and challenges, Opt. Netw. Mag. 1 (2000) 47–60.

[6] C. Guillemot et al., Transparent optical packet switching: Network, architecture and demonstrators in the KEOPS project, IEEE J. Sel. Areas Commun. 16 (1998) 1245–1259.

[7] H.S. Hinton, An Introduction to Photonic Switching Fabrics, Plenum Press, 1993.

[8] I.P. Kaminow et al., A wideband all-optical WDM network, IEEE J. Sel. Areas Commun. 14 (5) (1996) 780–799.

[9] V. Tamilraj, S. Subramaniam, An analytical blocking model for dual-rate sessions in multichannel optical networks, in: Proc. GLOBECOM, San Antonio, TX, 2001.

[10] V. Tamilraj, S. Subramaniam, Blocking of multirate circuits in multichannel optical networks, in: SPIE Opticomm, Boston, MA, 2002, pp. 89–100.

[11] A.C. Kam, K.-Y. Siu, R.A. Barry, E.A. Swanson, Toward best-effort services over WDM networks with fair access and minimum bandwidth guarantees, IEEE J. Sel. Areas Commun. 16 (7) (1998) 1024–1039.

[12] B. Mukherjee, D. Banerjee, S. Ramamurthy, A. Mukherjee, Some principles for designing a wide-area WDM optical network, IEEE/ACM Trans. Netw. 4 (5) (1996) 684–696.