

On the Benefit of Tunability in Reducing Electronic Port Counts in WDM/TDM Networks

Randall Berry
Dept. of ECE
Northwestern Univ.
Evanston, IL 60208, USA
e-mail: rberry@ece.northwestern.edu

Eytan Modiano
Laboratory for Information
and Decision Systems, MIT
Cambridge, MA 02139, USA
e-mail: modiano@mit.edu

Abstract—In this paper, we study the benefits of using tunable transceivers for reducing the required number of electronic ports in WDM/TDM networks. We show that such transceivers can be used to efficiently “groom” sub-wavelength traffic in the optical domain and so can significantly reduce the number of electronic ports compared to the fixed tuned case. We provide a new formulation for this “tunable grooming” problem. We show that in general this problem is NP-complete, but we are able to efficiently solve it for many cases of interest. When the number of wavelengths in the network is not limited, we show that each node only needs the minimum number of transceivers (i.e., no more transceivers than the amount of traffic that it generates). This holds regardless of the network topology or traffic pattern. When the number of wavelengths is limited, we show an analogous result for both uniform and hub traffic in a ring. We also develop a heuristic algorithm for general traffic that uses nearly the minimum number of transceivers. In most cases, tunable transceivers are shown to reduce the number of ports per node by as much as 60%.

I. INTRODUCTION

High capacity optical networks typically use a combination of wavelength division multiplexing (WDM) and time division multiplexing (TDM) techniques. In such a WDM/TDM network each fiber link supports multiple wavelength channels operating at a given bit rate, e.g., 2.5 Gbps (OC-48). The offered traffic is typically at a finer granularity than 1 wavelength, e.g. a traffic demand of 155 Mbps (OC-3) will only utilize 1/16th of a wavelength. To more efficiently utilize the network, this sub-wavelength traffic can be time-division multiplexed onto a wavelength, for example, using the Synchronous Optical Network (SONET) multiplexing hierarchy. Each node in such a WDM/TDM network requires some amount of terminal equipment for sending/receiving data. This includes optical transceivers for accessing the wavelength channels and electronic equipment for carrying out optoelectronic conversion and multiplexing the sub-wavelength traffic. The amount of terminal equipment required is a predominate factor in the cost of such a network, and reducing this cost is a important design consideration. Recently, there has been significant interest in reducing the required amount of terminal

equipment or ports by efficiently *grooming* the low rate traffic so that only a subset of the available wavelengths must be electronically processed at any node. The wavelengths that do not need processing can optically bypass the node without requiring electronic processing.

Most work on grooming in WDM/TDM networks has focused on the case where optical transceivers are fixed-tuned and so a fixed subset of wavelengths are dropped at a each node; each dropped wavelength requiring an electronic port (e.g. a SONET ADM). The basic traffic grooming problem as studied in [1–15] is to assign a given traffic requirement to wavelengths so that the total number of needed ports are minimized. The general traffic grooming problem has been shown to be NP-complete [1], even in the special case where all traffic is sent to a single egress node. However, for several special cases, algorithms have been found that significantly reduce the required number of ports. For example, for uniform all-to-all traffic, algorithms have been found for both bi-directional rings [2], [5–7] and unidirectional rings [1]. Heuristic algorithms for general (non-uniform) traffic have also been presented in [3], [9–11], [13], [14]. The port requirement in a network can be further reduced by using electronic switches (e.g. SONET cross-connects) to more efficiently groom the offered traffic [1], [4]. However, these switches also add a non-negligible cost to the network.

In this paper, we consider an alternative approach to designing WDM/TDM networks. This approach is based on using tunable optical transceivers, where these transceivers can be tuned from TDM time-slot to time-slot. With such transceivers, sub-wavelength traffic can be time-division multiplexed onto a wavelength *optically*. By shifting some the multiplexing functionality from the electronic domain to the optical domain, we can in many cases significantly reduce the amount of terminal equipment required over an architecture with fixed-tuned transceivers. Also, by grooming the traffic optically, there is no need for using electronic cross-connects as in [4]. The cost for these savings is in the tunable transceivers. These must have tuning times less than a time-slot, which may be on the order of μ s. Presently such fast-tunable transceivers are becoming available but are much more costly than their fixed tuned counterparts. It is reasonable to expect that as demand for tunable components increases, their cost will drop.

The work of R. Berry is supported in part by NSF CAREER award CCR-0238382. The work of E. Modiano is sponsored by DARPA under contract number MDA972-02-1-0021 and by NSF under grants ANI-0073730 and ANI-0335217.

One goal of this work is to quantify the savings in terminal equipment due to using tunable devices. As we show in this paper, this savings can be significant, both in terms of electronic and optical hardware.

This work compliments work on *reconfigurable* WDM networks, where tunable components are used to change the virtual topology in response to traffic variations or for protection purposes [16], [17]. Reconfiguration is generally thought of as occurring on a much slower time-scale than that considered here. It is also related to work on optical burst or packet switching [18], where fast-tunable components are used for switching bursty traffic. In optical burst or packet switching, the emphasis is on protocols for resolving contention and reserving bandwidth for randomly arriving traffic demands. Here our focus is on the case where traffic demands are changing slowly and a fixed TDM schedule can be calculated for each node. For example, this can be appropriate in the metro area.

In the following, we first describe the network model to be considered and give some examples to illustrate the benefits of tunability. For a given traffic demand, our goal is to design networks that use the minimum number of tunable ports, where a tunable port refers to all the hardware necessary to send and receive on a single wavelength including a tunable optical transceiver and an electronic port. Solving this problem, requires scheduling the offered traffic to the available time-slots and wavelengths so that the minimum number of ports are required. We give an integer linear programming (ILP) formulation for this *minimum tunable port problem* in a ring network. With a limited number of wavelengths, this problem is shown to be NP-complete, but we show it is efficiently solvable in a number of important cases including uniform traffic between an even number of nodes and when all traffic is sent to a single egress node. Additionally, in these cases each node is shown to need no more ports than the amount of traffic it generates. On the other hand, with sufficient wavelengths available, we show the problem admits a simple solution for any traffic demand. Again this solution results in each node using the minimum possible number of tunable ports. Moreover, this solution holds regardless of the network topology. For the limited wavelength case and a general traffic requirement, we also give heuristic algorithms with bounded approximation ratios. Numerical results show that using these approaches can provide as much as a 60% reduction in equipment.

II. NETWORK MODEL

We consider a network with N nodes numbered $1, \dots, N$. On each wavelength in the network, up to g low-rate circuits can be time division multiplexed; g is referred to as the traffic granularity. A static traffic requirement for the network is given by an $N \times N$ matrix $R = [R_{i,j}]$, where $R_{i,j}$ indicates the number of circuits required from node i to node j (the diagonal entries of R will be zero). A traffic requirement is *symmetric* if $R_{i,j} = R_{j,i}$ for all i, j ; this represents the case where all connections are bi-directional. Each node i generates

$R_i = \sum_j R_{i,j}$ circuits of traffic or, equivalently $W_i = R_i/g$ (fractional) wavelengths of traffic. For symmetric traffic, these quantities are also equal to the amount of traffic terminated by the node. For any architecture that supports the entire traffic requirement, each node i must have at least $\lceil W_i \rceil$ optical transmitters and, assuming the traffic demand is symmetric, $\lceil W_i \rceil$ optical receivers. When traffic is not symmetric, the number of transmitters and receivers at each node may be different. However, for simplicity of exposition we focus in this paper on the case of symmetric traffic. As will be evident, our results are easily applicable to asymmetric traffic as well. Also for simplicity, we focus in this paper on the case of unidirectional rings, and leave more general topologies for future research; although, as we point out, many of our results are applicable to general network topologies. Let W_{min} denote the minimum number of (fractional) wavelengths needed to support the given traffic requirement. In a unidirectional ring with symmetric traffic, each symmetric traffic demand $R_{i,j} = R_{j,i}$ uses exactly $R_{i,j}$ circuits around the ring, and so

$$W_{min} = \sum_{i \neq j} \frac{R_{i,j}}{2g} = \sum_{i=1}^N \frac{W_i}{2}.$$

Notice that since W_{min} maybe fractional, the actual number of wavelengths required is $\lceil W_{min} \rceil$.

Each node in the network is assumed to have a set of *tunable ports*, where each port includes a tunable optical transmitter and a tunable optical receiver. As noted above, we use the term “port” to refer to all of the equipment required to receive and transmit on one wavelength (see Fig. 1). From the above discussion, for symmetric traffic, each node requires at least $\lceil W_i \rceil$ tunable ports. Notice that our definition of a port allows a node to receive traffic on one wavelength and transmit on another simultaneously.

We begin by considering several simple examples to illustrate the potential advantages of tunability. Consider a unidirectional ring with $N = 4$ nodes, a granularity of $g = 3$, and assume that there is a uniform demand of one circuit between every pair of nodes, i.e. $R_{i,j} = 1$ for all $i \neq j$. In this case $W_{min} = 2$ and $W_i = 1$ for all i . Assume exactly W_{min} wavelengths are available. There are a total of $N(N - 1)$ circuits that need to be assigned to these two wavelengths. With $g = 3$, as many as 6 circuits can be assigned to each wavelength; this can be accomplished by assigning both circuits for each duplex connection to same time-slot. The traffic demand can then be supported by finding an assignment of each duplex connection to one of the g time-slots in the TDM frame, on one of the wavelengths. Without the possibility of tunable transceivers, the assignment of circuits to wavelengths corresponds to the standard traffic grooming problem considered in [1–14]. A simple approach would be to arbitrarily assign circuits to the wavelengths. For example, one such assignment is shown in Table I. Here $(i - j)$ indicates the duplex connection between nodes i and j . Notice that since transmissions from i to j and from j to i do not overlap on the ring, they can share the same time-

slot/wavelength. In this assignment, each node must transmit

TABLE I
AN ARBITRARY TRAFFIC ASSIGNMENT FOR $g = 3$.

| | λ_1 | λ_2 |
|--------|-------------|-------------|
| Slot 1 | (1-2) | (1-4) |
| Slot 2 | (1-3) | (2-3) |
| Slot 3 | (3-4) | (2-4) |

and receive on both wavelengths. As a result, two transceivers are needed per node for a total of 8 transceivers. A slightly more clever assignment, shown in Table II, only requires 7 transceivers, as node 1 is only assigned to transmit and receive on λ_1 .

TABLE II
OPTIMAL ASSIGNMENT FOR FIXED TUNED TRANSCEIVERS.

| | λ_1 | λ_2 |
|--------|-------------|-------------|
| Slot 1 | (1-2) | (2-3) |
| Slot 2 | (1-3) | (2-4) |
| Slot 3 | (1-4) | (3-4) |

In this simple example, the above assignment is the best possible without tunability and results in a savings of 1 transceiver. Many researchers have studied the static traffic grooming problem, and in general average transceiver savings that exceed 50% has been obtained for various traffic scenarios [1], [2], [10], [13]. All of the previous works have assumed that the transceivers are fixed tuned. However, as mentioned earlier, using tunable transceivers can help reduce the number of transceivers significantly. For example, consider the traffic assignment given in Table II. Notice that node 3 only transmits and receives on one wavelength at any given time (i.e., λ_2 in slot 1, λ_1 in slot 2 and λ_2 in slot 3). Hence, if node 3 were equipped with a tunable transceiver, it would only need one transceiver rather than 2 and a total of 6 transceivers would be required. In the assignment in Table II, nodes 2 and 4 must transmit on both wavelengths in the same slot and hence must each be equipped with two transceivers. Alternatively, a more clever assignment, shown in Table III, requires each node to transmit and receive only on one wavelength during each slot and so each node need only be equipped with a single tunable transceiver. Thus, the number of transceivers can be reduced from 7 to 4 by proper slot assignment. In this case, the optimal assignment can be found by inspection; however, in larger networks we will see that this can be a non-trivial combinatorial problem.

In what follows, we develop slot assignment algorithms, for certain cases, whereby each node only needs the minimum number of transceivers ($\lceil W_n \rceil$). In order to accomplish this, we must assign circuits to slots in such a way that each node is never assigned (to receive or transmit) on more than $\lceil W_n \rceil$ circuits during the same slot. Note that we allow the transmitter and receiver to be tuned to different wavelengths during a slot and hence it is possible for a node (even with just one

TABLE III
OPTIMAL ASSIGNMENT WITH TUNABLE TRANSCEIVERS.

| | λ_1 | λ_2 |
|--------|-------------|-------------|
| Slot 1 | (1-2) | (3-4) |
| Slot 2 | (1-3) | (2-4) |
| Slot 3 | (1-4) | (2-3) |

transceiver) to receive traffic from one wavelength and transmit to another during the same slot.

A. Synchronization issues

From the above discussion, it is clear that we require nodes to be synchronized at the slot level. Obtaining synchronization in a linear network (e.g., bus or ring) is rather straightforward, as all of the nodes can be synchronized to a single point of reference, as is commonly done with SONET rings. However, when the propagation delays around the ring are not negligible, a subtle problem arises in a slotted ring. Consider a slotted ring where slots are of duration t_s seconds and the propagation delay around the ring is t_p seconds. In order to maintain synchronization, when a slot propagates around the ring, it should return to its source on a slot boundary. Hence, as long as t_p is an integer multiple of t_s , synchronization is maintained. In practice, t_p may be arbitrary and not a multiple of t_s . In SONET rings, this problem is easily overcome by adding, using electronic buffers, a small delay at one of the nodes in order to make sure that the effective t_p is an integer multiple of t_s . In an optical ring adding delay is also possible using fiber delay lines. However, such an approach may be cumbersome. A simpler alternative is to use a framing whereby transmissions are synchronized along frame boundaries. Let g be the number of slots per frame, and suppose that $t_p > gt_s$, then, by starting a new frame every t_p seconds, synchronization can be achieved. Of course, this may result in the ring being idle for a duration $t_p - gt_s$. However, this idle time can be reduced by transmitting multiple frames every t_p seconds. Specifically, we can transmit $n = \lfloor t_p / (gt_s) \rfloor$ frames every t_p seconds and the amount of time during which the ring is idle would only be $t_p - ngt_s$. As with SONET, slotted rings require that the propagation delay around the ring is greater than the frame duration. For example, SONET frames are $125\mu s$ in duration; propagating at the speed of light this requires a ring of at least 25 miles. In practice, this minimum propagation delay around the ring can be artificially added.

B. Tunable transceivers

In this architecture, nodes are equipped with tunable transceivers. An example of such a node is shown in Fig. 1; however, many different implementations are possible. As shown in the figure, each tunable transceiver consists of a tunable optical ADM, a tunable receiver and a tunable laser. In addition, each node must also be equipped with optical-to-electrical (OE) and electrical-to-optical (EO) converters. There are a number of possible implementations of such tunable

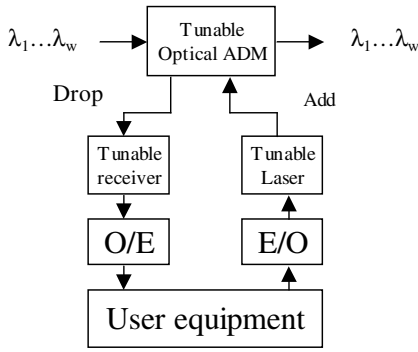


Fig. 1. An example node with tunable transceiver.

transceivers; the specifics of which are not of interest to this paper. However, the functionality of the tunable transceiver is to allow a node to remove and add data to a wavelength during a time slot. With time-slots durations on the order of μs ; these devices must be able to tune in sub- μs time.

As noted in the introduction, fast-tunable transceivers are relatively expensive. Over the past few years a number of manufacturers have begun to offer such devices in the market place. For example, fast-tunable lasers with switching times on the order of a nanosecond or less have been demonstrated (e.g., [19], [20]). Fast-tunable receivers (i.e. filters) and associated issues such as fast clock recovery [19] are also being actively researched. A goal of this paper is to examine the benefits of tunable components, so that the cost trade-off between tunable and fixed-tuned devices can be better understood. As we show in this paper, the use of tunable transceivers can reduce the number of transceivers considerably. It is important to point out that these savings are not only in terms of optical devices (i.e., saving on optical receivers and lasers), but also on costly OE and EO converters.

III. PROBLEM FORMULATION

We are interested in finding a time-slot assignment that minimizes the number of tunable ports needed for a network given a traffic requirement $R = [R_{i,j}]$, and W available wavelengths (clearly, to be feasible it must be that $W \geq W_{\min}$). We refer to this as the *minimum tunable port (MTP)* problem. Next, we give a precise ILP formulation for this problem. For simplicity, we restrict our attention to the symmetric traffic case and assume that the network is a unidirectional ring. However, the formulation can easily be extended to non-symmetric traffic or other network topologies.

Let X_i be an integer variable indicating the number of transceivers at node i . For $k = 1, \dots, g$ and $m = 1, \dots, W$, let $T_{i,j,m,k}$ be a $(0,1)$ -variable indicating that a transmission from node i to j is scheduled in time-slot k , on wavelength m . In a unidirectional ring, routing is fixed and a transmission from node i to j will use all of the links along the ring from node i to j . Number the links along the ring l_1, \dots, l_N , where

l_i is the link between node i and $i+1$. Let $E_{(i,j)}^l$ be a $(0,1)$ -variable indicating that the transmission from node i to j uses link l . Again, the values of $E_{(i,j)}^l$ are a deterministic function of i and j , and these variables are needed only for simplicity of the presentation. The desired optimization problem is then

$$\min \sum_{i=1}^N X_i$$

$$\text{subject to: } \sum_{m,j} T_{i,j,m,k} \leq X_i \text{ for all } i, k \quad (1)$$

$$\sum_{m,i} T_{i,j,m,k} \leq X_j \text{ for all } j, k \quad (2)$$

$$\sum_{m,k} T_{i,j,m,k} = R_{i,j} \text{ for all } i, j \quad (3)$$

$$\sum_{i \neq j} E_{(i,j)}^l T_{i,j,m,k} \leq 1 \text{ for all } l, k, m \quad (4)$$

$$T_{i,j,m,k} \in \{0,1\}, E_{(i,j)}^l \in \{0,1\} \quad (5)$$

Constraints (1) and (2) ensure that no node can transmit or receive on more wavelengths at any time than it has ports. Constraint (3) ensures that the traffic demand is satisfied, and (4) ensures that each time-slot on each wavelength is not used more than once on any link on the ring. Constraint (5) is the integer constraint.

We also consider a restricted version of the problem with the assumption that every bi-directional pair must be assigned to the same wavelength/time-slot, as in the examples from Sect. II. We refer to this as the *minimum tunable port with symmetric assignments (MTPS)* problem. With a unidirectional ring, symmetric assignments imply that each pair will take up one time-slot on one wavelength around the entire ring. This simplifies the optimization. In particular, $T_{i,j,m,k} = T_{j,i,m,k}$, and, due to the unidirectional ring routing, $E_{(i,j)}^l + E_{(j,i)}^l = 1$ (for all l). Specifically, the MTPS problem can be formulated identically to the MTP problem except for the following three changes. First in constraint (3), we only need to consider pairs $j > i$, due to the symmetry. Second, constraint (4) is replaced by

$$\sum_{i=1}^N \sum_{j=i+1}^N T_{i,j,m,k} \leq 1 \text{ for all } m, k, \quad (6)$$

and finally, the variables $E_{(i,j)}^l$ are not needed. Note that in (6) there are Wg constraints, while in (4) there are NWg constraints; this decrease reflects the reduced flexibility due to the symmetric assignment assumption. Since any solution to the MTPS problem is also a feasible solution for MTP, the solution to MTP will clearly be less than or equal to the solution to MTPS. Examples can be found where it is strictly less; however, in the following, we will see that in several important cases these two problems have the same solution.

IV. COMPLEXITY RESULTS

In this section we address the complexity of the MTP problem. We show that in general this problem is NP-complete.

We first show NP-completeness for the MTPS problem. Then we proceed to show that the problem without the symmetric assignment restriction is also NP-complete. Our proof is based on relating the MTP problem to the EDGE COLORING problem for an arbitrary graph. The EDGE COLORING problem is to find the fewest colors required to color the edges of a given graph G so that no two adjacent edges have the same color. The solution is called the chromatic index of the graph. Vizing's theorem [25], states that a graph's chromatic index is either Δ or $\Delta + 1$, where Δ is the maximum degree of any node in the graph. Determining which of these values is the chromatic index for an arbitrary graph is NP-complete [21].

Theorem 1: The MTPS problem is NP-complete.

Proof: To prove this, we show that any instance of EDGE COLORING can be transformed into an instance of the MTPS problem. Given a graph G , we identify it with a traffic matrix for the MTPS problem. Specifically, the nodes in the ring correspond to the nodes in G . The traffic demand is such that there is one bi-directional circuit ($R_{i,j}=R_{j,i}=1$) between each pair i, j for which there is an edge between the corresponding nodes in G , and $R_{i,j} = 0$ otherwise. Let the traffic granularity g equal the maximum degree in the graph G (hence, each node will generate at most one wavelength of traffic), and let the number of available wavelengths $W = \lfloor N/2 \rfloor$. We show that the solution to the MTPS problem is N transceivers (one per node) if and only if G has a chromatic index of g . From Vizing's theorem, it follows that any algorithm that can solve MTPS can be used to determine the chromatic index of G , and so MTPS must be NP-complete.

First, assume that the solution to MTPS uses N tunable transceivers. Notice that by construction each node in the ring must have at least one transceiver; if there are exactly N transceivers then this bound must be met with equality. With 1 transceiver per node, each node can transmit at most once in each time-slot, and by associating each time-slot with one color, it is clear that the corresponding solution yields an edge coloring of G using g colors. Hence, from Vizing's theorem, the chromatic index must be g . Notice that this must be an edge coloring because if two adjacent arcs had the same color, then the common node must be scheduled twice in the same time-slot.

Next, suppose that G 's chromatic index is g . In this case, an edge coloring using g colors can be mapped back into a time-slot assignment using exactly one transceiver for each node. Also, the number of edges labeled with a given color must be less than $\lfloor N/2 \rfloor$ (since each edge takes-up 2 of the N nodes), and so this assignment can be accommodated on the W available wavelengths. Therefore, we have shown that the chromatic index of G is g if and only if the solution to MTPS is N . ■

An example of the mapping between EDGE COLORING and MTPS is given in Fig. 2, where a 4 node complete graph is colored using 3 colors ($C1, C2, C3$). Since the graph is complete, each node has degree 3; hence, $g = 3$ and the corresponding traffic matrix is the uniform all-to-all traffic matrix (i.e. $R_{i,j} = 1$ for all $i \neq j$) used in the example of

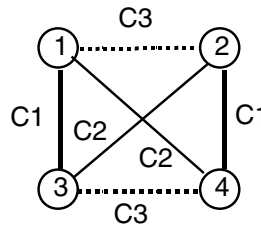


Fig. 2. A complete graph corresponding to the uniform traffic matrix. The coloring corresponds to the slot assignment given in Table III.

Sect. II. The MTPS solution for this traffic, using 3 slots, was given in Table III; this corresponds to the coloring in Fig. 2, where slot j in the table is identified with color j in the figure.

The next lemma implies that using more than the minimum number of wavelengths in the MTPS problem does not provide any benefit in terms of the number of ports. This result is useful in proving that the MTP problem is also NP-complete, and it will also be used several times when we discuss specific algorithms in the following sections.

Lemma 1: Any solution to the MTPS problem using $W > \lceil W_{\min} \rceil$ wavelengths can be converted in polynomial time into a solution using exactly $\lceil W_{\min} \rceil$ wavelengths without increasing the number of ports.

Proof: Given a solution to the MTPS problem (i.e. a time-slot assignment) using $W > \lceil W_{\min} \rceil$ wavelengths. We show that this can be converted into a time-slot assignment using $W - 1$ wavelengths without increasing the number of ports. The lemma then follows by iterating this argument.

Suppose that the n th time-slot has W circuits assigned to it. Then there must be some other time-slot, m , with $W - 2$ or fewer circuits. We can assume that time-slot m has exactly $W - 2$ circuits (if this is not the case, we can add extra "dummy" circuits to this time-slot). Consider the multi-graph¹ H constructed as follows. Identify a node in H with each transceiver of each node i assigned to the time-slots n or m . Thus if node i has X_i ports, there will be at most X_i nodes in H identified with this node, and possibly fewer if some of these ports are not used in these time-slots. Place an edge in H between the corresponding ports for each bi-directional circuit assigned to one of these time-slots. This graph will have a maximum degree of 2 since each port can be used at most once in each time-slot. Thus it will consist of one or more disjoint components where each component is either a cycle or a path (a sequence of edges with no repeated nodes). As in Theorem 1, the original time-slot assignment corresponds to a proper edge coloring of H using 2 colors. Therefore, each cycle must have an even length. If this were not the case then there is no way the corresponding edges could have been colored with 2 colors. Finally, note that this graph will have an even number ($W - 2 + W$) of edges; hence, the number of paths with odd lengths must be even. We construct a new

¹A multi-graph is a graph that may have multiple parallel edges between some of the nodes; H will have multiple edges if there is a circuit between the same pair of transceivers in both time-slots.

edge coloring of this subgraph using the same 2 colors in the following manner. Each cycle and each even length path can be colored using an equal number of each color. Since there are an even number of paths with odd lengths, these paths can also be colored so that each color is used an equal number of times. This results in exactly $W - 1$ edges being assigned each color; thus, the corresponding time-slot assignment will require no more than $W - 1$ wavelengths for these two time-slots and no additional ports as desired. Furthermore, the above reduction can be done in polynomial time and will be repeated at most $g(W - \lceil W_{\min} \rceil)$ times. ■

An example of the construction used in this proof is given next. Consider the time-slot assignment using $W = 3$ wavelengths shown in Table IV; this assignment is for a ring with $\lceil W_{\min} \rceil = 2$ and $N = 3$ nodes. One possible version of

TABLE IV
EXAMPLE TIME-SLOT ASSIGNMENT

| | λ_1 | λ_2 | λ_3 |
|--------|-------------|-------------|-------------|
| Slot 1 | (1-2) | (1-3) | (2-3) |
| Slot 2 | (2-3) | | |

the corresponding graph H is shown in Fig. 3.² Each node is represented by 2 nodes in the graph since 2 transceivers per node are required for the above assignment. The graph consists of two disjoint paths— one of length 3 and one of length 1. The edges are labeled with the original time-slot assignment as well as the new time-slot assignment given by the above lemma. The original assignment results in 3 edges labeled “1” and only one edge labeled “2.” Changing the label on the path of length 1, results in $\lceil W_{\min} \rceil = 2$ edges with each label. The

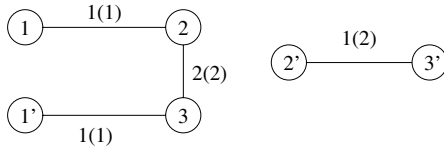


Fig. 3. The graph H corresponding to the time-slot assignment in Table IV. Each edge is labeled with $x(y)$, where x is the original time-slot assignment and y is the new assignment.

new time-slot assignment is shown in Table V and requires only 2 wavelengths as desired.

TABLE V
THE NEW TIME-SLOT ASSIGNMENT AFTER APPLYING THE APPROACH IN LEMMA 1.

| | λ_1 | λ_2 |
|--------|-------------|-------------|
| Slot 1 | (1-2) | (1-3) |
| Slot 2 | (2-3) | (2-3) |

A direct corollary of Lemma 1 is that the MTPS problem remains NP-complete when the wavelength limit is set to the

²In general, H is not unique but depends on the assignment of circuits to the available transceivers; any such assignment can be used.

minimum value. The next lemma shows that when W_{\min} is an integer, the MTP problem and MTPS problem have the same solution.

Lemma 2: If W_{\min} is an integer³, then any solution to the MTP problem using $W = \lceil W_{\min} \rceil = W_{\min}$ wavelengths must also be a solution to the MTPS problem.

The proof of this is given in Appendix I. The basic idea is that when W_{\min} is an integer, it can be shown that the only way to accommodate the entire traffic requirement on W_{\min} wavelengths is if a symmetric assignment is used. Using this lemma, it can be shown that the MTP problem is also NP-complete.

Theorem 2: The MTP problem with a wavelength limit of $\lceil W_{\min} \rceil$ is NP-complete.

The proof of this theorem is given in Appendix II. As in the proof of Theorem 1, we again use a correspondence with the EDGE COLORING problem.

V. MTP WITHOUT WAVELENGTH LIMITATION

We have shown that the MTP problem with a wavelength limit of $\lceil W_{\min} \rceil$ is NP-complete. In this section we show that without this wavelength limitation, the MTP problem can be efficiently solved and the solution requires exactly the minimum possible number of ports. The discussion of this section serves two purposes. First, it facilitates the discussion of the next section where we consider the problem with the restriction of using $\lceil W_{\min} \rceil$ wavelengths. Second, all of the results of this section are not topology dependent and so can be applied to a general network topology. We first state this result for the case where each node generates no more than one wavelength worth of traffic. As in Sect. IV, the proof relies on a relationship between EDGE COLORING and the MTP problem; however, here a different correspondence between these problems is used.

Theorem 3: If a network has no wavelength limitation and $W_i \leq 1$ for all $i = 1, \dots, N$, then each node requires only one tunable transceiver. Moreover, an optimal time-slot allocation can be found in polynomial time.

Proof: For this proof we represent the traffic requirement using a bipartite multi-graph, $B = (C, D, E)$, where C and D are two disjoint sets of nodes and E is a multi-set consisting of edges between nodes in C and nodes in D . Here C and D will each correspond to the set of nodes, $\{1, \dots, N\}$, and there will be $R_{i,j}$ edges between node $i \in C$ and $j \in D$. Note each edge in B represents a unidirectional circuit, while in Sect. IV an edge in the graph G represented a bi-directional pair of circuits. For a bi-partite multi-graph, it is known that the chromatic index is equal to the maximum degree [23]. For $W_i \leq 1$, the maximum degree of B must be less than or equal to g . Hence, an edge coloring of B can be found that uses at most g colors. In such a coloring, each color can again be identified with a particular time-slot during which the corresponding circuit will be transmitted. This is accomplished by having the transmitter and receiver

³Recall that W_{\min} was defined in units of fractional wavelengths.

both tune to an unused wavelength during the time-slot. Since the number of wavelengths is not limited, a free wavelength can always be found. Also, since no two adjacent edges have the same color, each node will not have to send or receive on more than one wavelength at any time, and so each requires at most one tunable transceiver. Therefore, finding an optimal solution to the MTP problem requires finding an optimal edge-coloring of B . In a bipartite graph, an optimal edge coloring can be found in polynomial time [22–24]. ■

In the proof of Theorem 3, an assignment of circuits to time-slots is identified with an edge coloring in the bipartite graph B . Equivalently, this can be thought of as a decomposition of B into a set of matchings, where each matching is assigned to one time-slot. Also notice that the solution given in this proof is not a “symmetric assignment” as required for the MTPS problem (i.e., traffic from i to j is not carried on the same wavelength/time-slot as traffic from j to i). In fact, only one unidirectional circuit is assigned to each wavelength/time-slot. This simplification results from not having a wavelength limitation and allows us to solve an otherwise NP-complete problem in polynomial time. Next, we extend this result to the case where some nodes may generate more than one wavelength worth of traffic.

Corollary 1: If a network has no wavelength limitations, then each node requires $\lceil W_i \rceil$ tunable ports. Moreover, a time-slot allocation that achieves this can be found in polynomial time.

Proof: If $W_i \leq 1$, then this follows from Theorem 3. When $W_i > 1$, for some i , the corresponding node in the bipartite graph B will have a degree greater than g . Hence an edge coloring of B will require more than g colors.⁴ In this case, we construct a new bipartite graph $B' = (C', D', E')$ as follows: for each node $i \in C$, we put $\lceil W_i \rceil$ “children” nodes in C' ; likewise, for each node $j \in D$, we put $\lceil W_j \rceil$ “children” nodes in D' . For each edge $(i, j) \in E$, we place an edge in B' between one of the children of i and one of the children of j such that no node in B' has a degree larger than g . This can be done because the total amount of traffic to be assigned to the $\lceil W_i \rceil$ children of a node i must be less than $\lceil W_i \rceil g$ by definition. The graph B' will again have a chromatic index less than or equal to g . Given a coloring of this graph using g colors, we again identify each color with a time-slot. Each node i will now need one tunable port for each of its $\lceil W_i \rceil$ children in the graph B' . In this case, finding an optimal solution to the MTP problem requires constructing the graph B' and finding an edge coloring in this graph, both of which require only polynomial complexity. ■

As an example, consider a ring with $N = 5$ nodes, and assume that there is a uniform traffic demand of one circuit between each pair of nodes. The bipartite graph B in Theorem 3 for this example is shown in Fig. 4. This graph has a

maximum degree of $N - 1 = 4$ and hence can be colored using 4 colors. If $g \geq 4$, this coloring can be used to provide a time-slot assignment using one tunable port per node. However, if $g = 3$, then $W_i = 4/3$ and from Corollary 1, each node will require $\lceil W_i \rceil = 2$ tunable ports. Following the proof of Corollary 1, each node is split into 2 children. A corresponding graph⁵ B' is shown in Fig. 5; this graph has maximum degree of 3 and so can be colored using $g = 3$ colors.

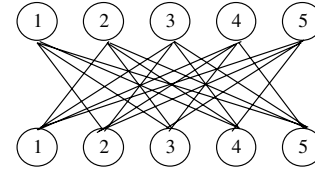


Fig. 4. Example of bipartite graph, B , corresponding to uniform traffic between $N = 5$ nodes.

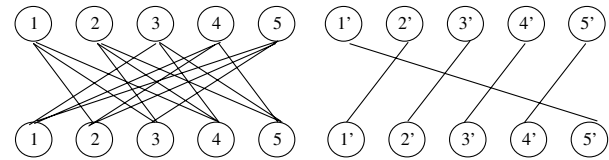


Fig. 5. Example of the graph B' used in Corollary 1. The two children of each node i in B are labeled by i and i' in B' .

In a unidirectional ring, the above approach requires $\sum_i \lceil W_i \rceil$ wavelengths. This follows because each child of a node as given in Corollary 1 will require one wavelength. Furthermore, $\sum_i \lceil W_i \rceil \geq 2W_{\min}$, with equality when W_i is an integer for all i . In other words, approximately twice the minimum number of wavelengths is required for this approach. This is because the algorithm only assigns one circuit to each wavelength during each time-slot. In the next section we consider algorithms that use wavelengths more efficiently by packing more than one circuit into a wavelength/time-slot pair.

Also, notice that in the above approach, the assignments can be arranged so that each node (child) in the graph B' always transmits on the same wavelength. This is because there is only one transmission on each wavelength/time-slot pair and so which wavelength this occurs on does not matter. Hence, this solution can be realized if each node has fixed tuned transmitters and only tunable receivers. Alternatively, it is also possible for each node to have fixed tuned receivers and only tunable transmitters. Finally, we note that the above solutions extend directly to the case where the network does not have a unidirectional ring architecture, or even a ring architecture for that matter.

VI. LIMITED WAVELENGTHS

When wavelengths are limited, the time-slot allocation of the previous section will no longer be feasible and the circuits

⁴At first one might think that an edge coloring of G could directly be used to find a time-slot assignment, by assigning no more than $\lceil W_i \rceil$ colors to each time-slot. The problem with this approach is that if the $\lceil W_i \rceil$'s are not equal for all i it is not straightforward to do this in way that ensures no node will need more than $\lceil W_i \rceil$ ports.

⁵There are many different ways to construct the graph B' ; Fig. 5 illustrates one possible construction.

must be more efficiently packed onto the available wavelengths. We consider this case in the following; in particular, we focus on the case with the tightest wavelength restriction, i.e., $W = W_{min}$. We assume that the traffic requirement is symmetric and consider slot assignment algorithms that use symmetric assignments, as required for the MTPS problem. From Lemma 2, we know that if W_{min} is an integer then there is no loss in performance by restricting ourselves to such an assignment. For non-integer W_{min} , symmetric assignments may not be optimal for the MTP problem, but this restriction simplifies the problem considerably. First, we show that in a number of special cases the optimal slot-assignment can be efficiently found. We then consider heuristics for the general case.

A. Uniform all-to-all traffic

Our first result applies to the case of uniform traffic, i.e., there are exactly r circuits between each pair of nodes, $r = R_{i,j}, i \neq j$. The following theorem states that in this important special case, if the number of nodes is even then each node need only be equipped with the minimum number of transceivers.

Theorem 4: In a ring with a uniform traffic requirement, N even and $\lceil W_{min} \rceil$ wavelengths, each node requires $\lceil W_i \rceil$ tunable transceivers. Moreover, an optimal time-slot allocation can be found in polynomial time.

Proof: Similar to the proof of Theorem 1, we again use a correspondence between the traffic requirement and a graph G . In this case, since there may be multiple circuits between a pair of nodes, G may be a multi-graph. Each node in G again corresponds to a node in the ring and there is an edge between each pair of nodes corresponding to each bi-directional circuit required between the nodes. Since the traffic demand is uniform, G will be a complete multi-graph, with $r = R_{i,j}$ edges between each pair of nodes. A complete multi-graph with an even number of nodes has a chromatic index equal to its degree [26]. Thus, we can find an edge coloring for this graph using $(N-1)r$ colors. Since the traffic is uniform, each node will have one edge incident to it with each of the $(N-1)r$ colors. Each edge is shared by two nodes, so there will be a total of $N/2$ edges of each color. Such a coloring can be found in polynomial time. We next show how to use this coloring to find the desired time-slot allocation.

First, assume that there are $\hat{W} = (\lceil W_i \rceil)(N/2)$ wavelengths available. Since there are $N/2$ edges of each color, it follows that using \hat{W} wavelengths, all of the traffic corresponding to $\lceil W_i \rceil$ distinct colors can be assigned to a single time-slot. Also, since $\lceil W_i \rceil (N/2)g > W_i(N/2)g = N(N-1)r/2$, we can assign all of the traffic in this way. This results in a time-slot allocation using $\lceil W_i \rceil$ transceivers per node. If W_i is an integer then, since N is even, $\hat{W} = \lceil W_{min} \rceil$ and we are done. If W_i is not an integer, then $\hat{W} > \lceil W_{min} \rceil$ and this allocation uses more than the required number of wavelengths. However, from Lemma 1, the allocation can be transformed in polynomial time into an allocation using only $\lceil W_{min} \rceil$ wavelengths. ■

In Theorem 4, each node meets the lower bound on the required number of ports. Therefore, we have found a solution to both the MTPS problem as well as the MTP problem, i.e., there is no loss from requiring all assignments to be symmetric. This theorem only applies when there is an even number of nodes in the ring. A complete graph with an odd number of nodes does not have a chromatic index equal to its degree; this case will be addressed in Section VI.C.

Example: Consider a ring with $N = 6$ nodes, a granularity of $g = 3$, and assume that there is a uniform demand of one circuit between each pair of nodes. In this case, $W_i = 5/3$ and $W_{min} = 5$. The corresponding graph G has a chromatic index of 5. Applying an edge coloring algorithm from [26], we get the coloring of G using 5 colors shown in Table VI.

TABLE VI
AN EDGE COLORING FOR UNIFORM TRAFFIC.

| Color | Node pairs |
|-------|---------------------|
| C_1 | (1-6), (5-2), (4-3) |
| C_2 | (2-6), (1-3), (5-4) |
| C_3 | (3-6), (2-4), (1-5) |
| C_4 | (4-6), (3-5), (2-1) |
| C_5 | (5-6), (4-1), (3-2) |

Since $\lceil W_i \rceil = 2$, each node needs 2 transceivers and so we can assign two colors per time-slot. With $\hat{W} = 6$ wavelengths, a valid time-slot assignment results by assigning colors C_1 and C_2 to time-slot 1, C_3 and C_4 to time-slot 2, and C_5 to time-slot 3. The resulting assignment is shown in Table VII. Using Lemma 1, this can be transformed into an assignment using $W_{min} = 5$ wavelengths, as shown in Table VIII. Notice that in both cases, each node appears at most twice in each time-slot and hence requires $\lceil W_i \rceil = 2$ tunable transceivers.

TABLE VII
TIME-SLOT ASSIGNMENT USING 6 WAVELENGTHS.

| | λ_1 | λ_2 | λ_3 | λ_4 | λ_5 | λ_6 |
|--------|-------------|-------------|-------------|-------------|-------------|-------------|
| Slot 1 | (1-6) | (5-2) | (4-3) | (2-6) | (1-3) | (5-4) |
| Slot 2 | (3-6) | (2-4) | (1-5) | (4-6) | (3-5) | (2-1) |
| Slot 3 | (5-6) | (4-1) | (3-2) | | | |

TABLE VIII
TIME-SLOT ASSIGNMENT USING $W_{min} = 5$ WAVELENGTHS

| | λ_1 | λ_2 | λ_3 | λ_4 | λ_5 |
|--------|-------------|-------------|-------------|-------------|-------------|
| Slot 1 | (1-6) | (5-2) | (4-3) | (2-6) | (1-3) |
| Slot 2 | (3-6) | (2-4) | (1-5) | (4-6) | (3-5) |
| Slot 3 | (5-6) | (4-1) | (3-2) | (5-4) | (2-1) |

B. Egress Traffic

Another important class of traffic for which the optimal time-slot assignment can be found is a ring with "egress" traffic. That is a traffic demand where all circuits are either to or from a single "hub" node in the ring, i.e. $R_{i,j} > 0$ only if

either $i = h$ or $j = h$, where node h is the hub node. This traffic demand is particularly relevant in metro area networks where most of the traffic on the ring goes to one or two central office hubs. We note that without tunable transceivers the traffic grooming problem for egress traffic is NP-complete [1]. However, with tunability, the MTP problem for egress traffic can be solved optimally in polynomial time.

Theorem 5: In a ring with egress traffic and $\lceil W_{\min} \rceil$ wavelengths, each node requires $\lceil W_i \rceil$ tunable ports and an optimal time-slot allocation can be found in polynomial time.

Proof: As in Theorem 4, let the traffic requirement be represented by a multi-graph G . With egress traffic, this graph will be a bipartite multi-graph (where one set of nodes corresponds to the hub node(s) and the other to the non-hub nodes). As noted in Sect. V, the chromatic index of a bipartite multi-graph is always the maximum degree and an edge coloring achieving this can be found in polynomial time. Once again this edge coloring can be used to find the desired time-slot assignment.⁶ ■

Theorem 5 does not require that the ring have a unidirectional topology and applies to bi-directional rings as well (with W_{\min} appropriately re-defined). It also extends directly to a traffic pattern with multiple egress nodes, where there is no traffic between the egress points, or to any other traffic requirement for which the corresponding graph is bi-partite (i.e., the nodes can be divided into two groups and traffic only flows between the groups but there is no traffic within a group). Finally, we note that with a single egress node, the hub node must be able to receive traffic on all wavelengths; so, tunable components are only needed at the non-egress nodes.

Example: Consider a ring with N nodes, a granularity of g , and assume that there is uniform egress traffic to a single hub, i.e. the traffic requirement is for r bi-directional circuits between each non-hub node and the hub. In this case, for each non-hub node $W_i = r/g$ and for the hub node, h , $W_h = (N - 1)r/g$. Hence from Theorem 5, the traffic can be supported using

$$\left\lceil \frac{(N-1)r}{g} \right\rceil + (N-1) \left\lceil \frac{r}{g} \right\rceil \quad (7)$$

tunable ports. With fixed tuned transceivers, the minimum number of transceivers for uniform egress traffic is found in [1]. Comparing (7) to the results in [1], it can be shown that tunable ports can reduce the number of transceivers required at the hub by up to 50%.

C. Heuristics Algorithms for arbitrary traffic demands:

We next present heuristic time-slot assignment algorithms for an arbitrary (symmetric) traffic demand. First, we consider a “0-1” traffic requirement, where $R_{i,j}$ is either 0 or 1 for all i and j ; i.e. at most one circuit is established between each pair of nodes. In this case the corresponding graph G will be a simple graph⁷. For simple graphs, from Vizing’s theorem, the

⁶In fact, with a single hub node the time-slot assignment can be found directly using a simple greedy algorithm.

⁷A simple graph is a graph with no parallel edges between the same pair of nodes, i.e. it is not a multi-graph

chromatic index is at most $\Delta + 1$, where Δ is the maximum degree of the graph. Moreover, polynomial time algorithms for coloring any simple graph with $\Delta + 1$ colors are known [25]. We use this fact to develop time-slot assignment algorithms for the ring. To begin we construct the graph G based on the traffic requirement. Following a similar approach to Corollary 1, we then construct a new graph G' by replacing each node i in G with $\lceil R_i/(g-1) \rceil$ “children nodes” in G' , and each edge in G with an edge between two of the corresponding children in G' so that no node in G' has a degree greater than $g-1$. In this case G' can be colored using at most g colors and this coloring can be used for a time-slot assignment, where each node in G' requires one tunable transceiver. Hence, using this time-slot assignment, each node will need $\lceil R_i/(g-1) \rceil$ tunable transceivers. With $\lceil R_i/(g-1) \rceil$ tunable transceivers, each node will transmit on at most $\lceil R_i/(g-1) \rceil$ wavelengths and the total number of wavelengths required will be at most $\frac{1}{2}(\sum_i \lceil R_i/(g-1) \rceil)$ wavelengths. This may be greater than W_{\min} , but, using Lemma 1, it can be transformed into an assignment using W_{\min} wavelengths without increasing the number of ports. To summarize, we have shown the following:

Theorem 6: For any 0-1 traffic requirement, a time-slot assignment using W_{\min} wavelengths can be found (in polynomial time) where each node i uses $\lceil R_i/(g-1) \rceil$ tunable transceivers.

The optimal time-slot assignment requires at least $\lceil W_i \rceil = \lceil R_i/g \rceil$ transceivers. Hence, the performance ratio of the above heuristic is bounded by

$$\frac{\lceil R_i/(g-1) \rceil}{\lceil R_i/g \rceil}$$

for each node i . This is at most 2 and in many cases will be equal to one.

Example: Consider a ring with $N=5$ nodes a granularity of $g=3$, and assume that the traffic is a uniform demand of one circuit between each pair of nodes. In this case, $W_i = 4/3$ and $W_{\min} = 10/3$. Since N is odd, Theorem 4 does not cover this case; hence, we apply the heuristic in Theorem 6. Following this procedure we split each node into $\lceil R_i/(g-1) \rceil = 2$ children and form the graph G' where each node has a degree no greater than $2(g-1)$, as shown in Fig. 6. A proper edge coloring of G' using at most $g=3$ colors is also indicated in the figure. This coloring translates into the time-slot assignment shown in Table IX, which uses 2 tunable transceivers per node and $\lceil W_{\min} \rceil = 4$ wavelengths. Notice that in this case $\lceil W_i \rceil = 2$, and therefore, this is in fact the optimal time-slot assignment.

TABLE IX
TIME-SLOT ASSIGNMENT CORRESPONDING TO FIG. 6

| | λ_1 | λ_2 | λ_3 | λ_4 |
|--------|-------------|-------------|-------------|-------------|
| Slot 1 | (1-2) | (1-4) | (2-4) | |
| Slot 2 | (2-3) | (4-5) | (3-4) | (2-5) |
| Slot 3 | (1-3) | (1-5) | (3-5) | |

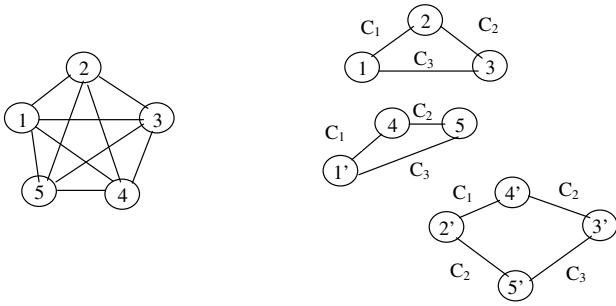


Fig. 6. Example graphs. On the left is the original graph G ; on the right is the derived graph G' which has a maximum degree of 2. The two “children” of each node i in G are labeled by i and i' in G' . An edge coloring of G' using three colors (C_1, C_2, C_3) is also shown.

For a general traffic requirement, the graph G will not be a simple graph. For a multi-graph, Vizing’s theorem does not apply, but two other upper bounds on the chromatic index are known. The first, also due to Vizing [25], states that the chromatic index is less than or equal to $\Delta + m$, where Δ is the maximum degree and m is the maximum number of parallel edges between any two nodes. The second upper bound, due to Shannon [27], is $(3/2)\Delta$. Cases can be found where either of these is the tighter. Next, we consider a heuristic for a general traffic matrix that uses Shannon’s bound (a similar approach can be developed using Vizing’s bound). Once again we begin the traffic graph G and construct a new graph G' . This time we replace each node i in G with $\lceil (3R_i)/(2g) \rceil$ “children” nodes in G' , and each edge in G with an edge between two of the corresponding children in G' so that no node in G' has a degree greater than $(2/3)g$. Hence, G' can be colored using no more than g colors and this can again be mapped into a time-slot assignment. Therefore, we have the following:

Theorem 7: For any symmetric traffic requirement, a (symmetric) time-slot assignment can be found (in polynomial time) where each node uses $\lceil (3R_i)/(2g) \rceil$ tunable transceivers and W_{min} wavelengths.

D. Numerical examples

We present some numerical examples that compare the number of tunable ports required to the number of fixed tuned ports needed. Figure 7 shows the number of ports in a ring with $g = 4$ and a uniform demand of $r = 1$ circuit for different values of N . Three curves are shown in the figure. The top curve is a lower bound on the number of ports required in a ring with fixed-tuned transceivers given in [1]; we note that in general this bound is not tight. The middle curve is the number of ports needed with tunable transceivers and W_{min} wavelengths. When N is even this is given by Theorem 4; when N is odd, the heuristic algorithms in Sect. VI.C are used. The bottom curve is the number of tunable ports needed without any wavelength restrictions; this is found using Theorem 3 and Corollary 1 in Sect. V. In this case with tunability, the number of ports can be reduced by over 40%. Also note that there is little difference between the case with

wavelength limitation and without. This is expected because when N is even we know the two cases should be equal. Fig. 8 shows an analogous set of curves for the case where $g = 16$; here, tunability reduces the number of ports by up to 60%.

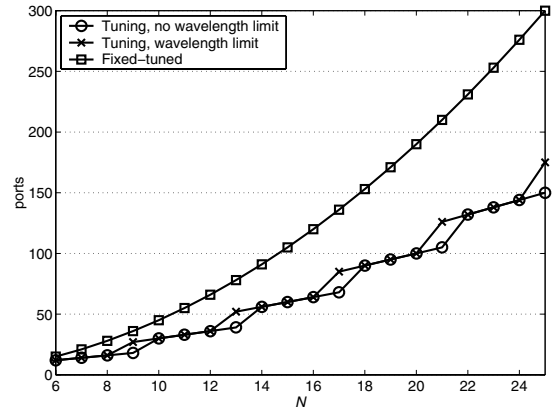


Fig. 7. Number of ports vs. N for a ring with uniform demand of $r = 1$ circuits and $g = 4$.

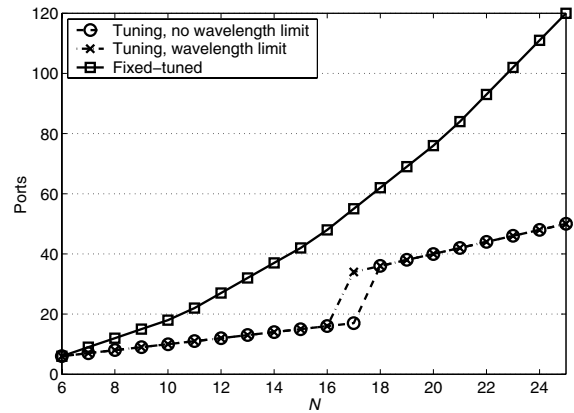


Fig. 8. Number of ports vs. N for a ring with uniform demand of $r = 1$ circuits and $g = 16$.

VII. CONCLUSIONS

In this paper we considered the problem of traffic grooming in a WDM/TDM network with tunable transceivers. While we show that the problem is generally NP-complete, we are able to solve it for many cases of interest. When the number of wavelengths in the network is not limited, we show that each node only needs as many transceivers as the number of wavelength worth of traffic that it generates. This results holds regardless of the network topology or the traffic pattern. When the number of wavelengths is limited, we show that the same holds for uniform and hub traffic in a ring network. We also provide heuristic algorithms for general traffic in a ring. In all cases we observe transceiver savings of up to 60% as compared to fixed-tuned transceivers.

One goal of this work is to quantify the benefits of tunable optical components. While presently tunable transceivers are

much more costly than their fixed tuned counterparts, the fact that they can significantly reduce the amount of hardware required in the network (both optical and electronic) may justify their use. Our work is preliminary, in that for the most part we focus on a unidirectional ring topology. However, the promising results that we observe open up many new avenues for future research. For example, we would like to generalize this work to topologies other than rings. One avenue of possible research is to examine the trade-off between the number of available wavelengths and the number of transceivers needed. Another possible avenue is to examine the effects of limited tunability (e.g., tunable transmitters and fixed tuned receivers or vice versa). In all cases our objective is to tradeoff additional complexity in optical hardware for a significant reduction in the electronic hardware. As the cost and capabilities of optical hardware improve, such a trade-off may become extremely beneficial.

APPENDIX I PROOF OF LEMMA 2

Proof: To prove this lemma, we simply need to show that any traffic assignment that uses the minimum number of wavelengths must be symmetric. Define the *length* of a circuit to be the number of links crossed by the circuit, (e.g. a circuit between nodes i and $i+1$ has a length of 1). We say a circuit of length L takes up L/N of a full wavelength/time-slot around the ring. Since each bi-directional pair takes up one full wavelength/time-slot around the ring, the number of full wavelength/time-slots occupied by a given traffic matrix is equal to

$$r = \sum_{i=1}^N \sum_{j=1}^N R_{i,j} / 2 = gW_{min}.$$

Therefore, when $W = W_{min}$ is an integer, there are exactly r full wavelength/time-slots available, all of which must be occupied. Consider a traffic assignment to these r wavelength/time-slots and let r_i be the number of unidirectional circuits assigned to wavelength/time-slot i . Since each wavelength/time-slot must be fully occupied, $r_i \geq 2$ for each i . Also let L_i^j be the length of the j th circuit assigned to the i th ring. Then the total length of calls around the ring is

$$L = \sum_{i=1}^r \sum_{j=1}^{r_i} L_i^j = rN.$$

Since the traffic is symmetric, this number must equal the length of adjacent circuits around the ring, where the circuit from j to i is defined to be adjacent to the circuit from i to j . Moreover, if a circuit from i to j has length l , then the adjacent circuit will have length $N - l$. Hence counting the

length of adjacent circuits we obtain,

$$\begin{aligned} L_{adj} &= \sum_{i=1}^r \sum_{j=1}^{r_i} (N - L_i^j) \\ &= \sum_{i=1}^r \sum_{j=1}^{r_i} N - \sum_{i=1}^r \sum_{j=1}^{r_i} L_i^j \\ &= \sum_{i=1}^r \sum_{j=1}^{r_i} N - rN \geq 2rN - rN = rN. \end{aligned}$$

The last inequality holds because $r_i \geq 2$ and equality is obtained if and only if $r_i = 2$, for all i . However, the only traffic assignment that satisfies $r_i = 2$ is a symmetric assignment. ■

APPENDIX II PROOF OF THEOREM 2

Proof: To prove this theorem, we again use a reduction from EDGE COLORING. Given a graph G with maximum degree Δ , we use the correspondence from the proof of Theorem 1 to map this graph into a traffic matrix for the MTP problem in a ring with $g = \Delta$. Hence, $W_{min} = \sum_i \frac{d_i}{2g}$, where d_i is the degree of node i .

First, consider the case where W_{min} is an integer and assume that there are W_{min} wavelengths. In this case from Lemma 2, the solution to the the MTP problem is N if and only if the solution to the MTPS problem with W_{min} wavelengths is N . Following the proof of Theorem 1, G has a chromatic index of g if and only if the solution MTPS problem with $W = \lfloor N/2 \rfloor$ wavelengths is N . Finally, from lemma 1, the solution to MTPS problem with $W = \lfloor N/2 \rfloor$ wavelengths is N if and only if the solution to the MTPS problem with W_{min} wavelengths is N . Combining these observations we have that the solution to the MTP problem is N if and only if G has a chromatic index of g .

Next, assume W_{min} is not an integer. In this case, we cannot directly use lemma 2. Instead, define a new graph G' as follows. Let G' be a graph with $2g$ disconnected components, where each component is isomorphic to G . Clearly the maximum degree of G' is also g and the chromatic index of G' is equal to the chromatic index of G . Consider the instance of the MTP problem with traffic corresponding to G' in a ring with granularity g and $W = W'_{min} = \sum_i d'_i / (2g)$ wavelengths, where d'_i is the degree of the i th node in G' . Therefore, $W'_{min} = W_{min}(2g)$ is an integer, and, by the same arguments as above, the solution to the MTP problem is $2Ng$ if and only if G has a chromatic index of g . ■

REFERENCES

- [1] E. Modiano and A. Chiu, "Traffic Grooming Algorithms for Minimizing Electronic Multiplexing Costs in Unidirectional SONET/WDM Ring Networks," *CISS 98*, Princeton, NJ, Feb., 1998. Extended version appeared in *IEEE Journal of Lightwave Technology*, Jan., 2000.
- [2] J. Simmons, A. Saleh, "Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic," *IEEE Journal of Lightwave Technology*, vol. 17, pp. 48-57, Jan. 1999.
- [3] R. Berry and E. Modiano, "Reducing Electronic Multiplexing Costs in SONET/WDM Rings with Dynamically Changing Traffic," *IEEE Journal of Selected Areas in Communication*, October 2000.

- [4] E. Modiano and R. Berry, "Using grooming cross-connects to reduce ADM costs in SONET/WDM ring networks," *OFC 2001*, Anaheim, CA, March 2001.
- [5] O. Gerstel and R. Ramaswami, "Cost effective Grooming in WDM Rings," *Infocom 98*, San Francisco, CA, April, 1998.
- [6] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength assignment in a WDM ring to minimize the cost of embedded SONET rings," *Proc. Infocom 98*, San Francisco, CA, Apr. 1998.
- [7] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET network design," *Proc. Infocom 99*, New York, Mar. 1999.
- [8] T. Y. Chow, P. J. Lin "The Ring Grooming Problem" *SIAM Journal of Discrete Mathematics*, 2001.
- [9] V. R. Konda, T. Y. Chow "Algorithm for traffic grooming in optical networks to minimize the number of transceivers," *IEEE 2001 Workshop on High Performance Switching and Routing*.
- [10] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network", *IEEE Journal of Selected Areas in Communication*, January, 2002.
- [11] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network", *Proc. ICC*, 2001.
- [12] E. Modiano and P. Lin, "Traffic grooming in WDM networks", *IEEE Communications Magazine*, July, 2001.
- [13] X. Zhang and C. Qiao, "An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings", *IEEE/ACM Trans. on Net.*, Oct., 2000.
- [14] H. Ghafouri-Shiraz, G. Zhu, and Y. Fei, "Effective Wavelength Assignment Algorithms for Optimizing Design Costs in SONET/WDM Rings," *IEEE Journal of Lightwave Technology*, Vol. 19, No. 10, Oct.2001.
- [15] R. Dutta and G. N. Rouskas, "On Optimal Traffic Grooming in WDM Rings," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 1, Jan. 2002, pp. 110-21.
- [16] G. N. Rouskas and M. H. Ammar, "Dynamic reconfiguration in multihop WDM networks," *Journal of High Speed Networks*, vol. 4, no. 3, pp. 221-238, June 1995.
- [17] A. Narula-Tam and E. Modiano, "Dynamic load balancing in WDM packet networks with and without wavelength constraints," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, pp. 1972-1979, Oct. 2000.
- [18] C. Qiao and M. Yoo, "Optical Burst Switching - A New Paradigm for an Optical Internet," *Journal of High Speed Networks*, vol. 8, No. 1, 1999.
- [19] K. Shrikhande, et al., "Performance Demonstration of a Fast-tunable Transmitter and Burst Mode Packet Receiver for HORNET," *Proc. of OFC 2001*.
- [20] M. Kaurer, et al., "16-Channel Digitally Tunable Packet Switching Transmitter with sub-nanosecond switching time," *Proc. ECOC*, 2002.
- [21] I. Holyar, "The NP-completeness of edge-coloring," *SIAM J. on Computing*, vol. 10, no. 4, 718-720, 1981.
- [22] H. N. Gabow, O. Kariv, "Algorithms for edge coloring bipartite graphs and multigraphs," *SIAM J. on Computing*. **11**(1982),117-129.
- [23] R. Cole, J. Hopcroft, "On edge coloring bipartite graphs," *SIAM J. on Computing* **11** (1982) 540-546.
- [24] A. Schrijver, "Bipartite edge-coloring in $O(\Delta m)$ time," *SIAM J. on Computing* **28**(1999),841-846.
- [25] V. G. Vizing, "On an estimate of the chromatic class of a p-graph," *Diskret. Analiz* **3**(1964),25-30.
- [26] F. Harary, *Graph Theory*, Addison Wesley, 1971.
- [27] C. E. Shannon, "A Theorem on coloring the lines in a network," *J. Math. Phys.*, vol. 28, 148-151, 1949.