

Optimal Transceiver Scheduling in WDM/TDM Networks

Randall Berry, *Member, IEEE*, and Eytan Modiano, *Senior Member, IEEE*

Abstract—In this paper, we study the benefits of using tunable transceivers for reducing the required number of electronic ports in wavelength-division-multiplexing/time-division multiplexing optical networks. We show that such transceivers can be used to efficiently “groom” subwavelength traffic in the optical domain and so can significantly reduce the amount of terminal equipment needed compared with the fixed-tuned case. Formulations for this “tunable grooming” problem are provided, where the objective is to schedule transceivers so as to minimize the required number of ports needed for a given traffic demand. We establish a relationship between this problem and edge colorings of graphs which are determined by the offered traffic. Using this relationship, 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, each node is shown to only require 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, an analogous result is shown for both uniform and hub traffic in a ring. We then 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%. We also consider the case where traffic can dynamically change among an allowable set of traffic demands. Tunability is again shown to significantly reduce the port requirement for a nonblocking ring, both with and without rearrangements.

Index Terms—Graph coloring, integer linear programming (ILP), optical networks, traffic grooming, wavelength-division multiplexing (WDM).

I. INTRODUCTION

HIGH-CAPACITY optical networks are typically based on 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 Gb/s (OC-48). The offered traffic is typically at a finer granularity than 1 wavelength, e.g., a traffic demand of 155 Mb/s (OC-3) will only utilize 1/16 of a wavelength. To more efficiently utilize the network, this subwavelength traffic is 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/demultiplexing the subwavelength traffic. The amount of terminal equipment required is a predominate factor in the cost of such a network; reducing this cost is an important design consideration. Recently, there has been significant interest in reducing the required amount of terminal equipment 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 any terminal equipment.

Most work on grooming in WDM/TDM networks has focused on the case where optical transceivers are *fixed-tuned* so that a fixed subset of wavelengths are dropped at a each node; each dropped wavelength requires an electronic port (e.g., a SONET ADM). The basic traffic grooming problem as studied in [1]–[16] is to assign a given traffic requirement to wavelengths so that the total number of needed ports are minimized.¹ 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 bidirectional rings [2], [5]–[7], and unidirectional rings [1]. Heuristic algorithms for general (nonuniform) traffic have also been presented in [3], [9]–[11], [13], [14], and [16]. 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 nonnegligible 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, subwavelength 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

Manuscript received December 12, 2003; revised May 11, 2005. This paper was presented in part at INFOCOM 2004, Hong Kong, China, March 7–11, 2004 and in part at OFC 2004, Los Angeles, CA, February 22–27, 2004.

R. Berry is with the Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL 60208 USA (e-mail: rberry@ece.northwestern.edu).

E. Modiano is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: modiano@mit.edu).

Digital Object Identifier 10.1109/JSAC.2005.852240

¹Some authors, such as [10] and [11], actually consider the “dual” formulation of maximizing the throughput given a fixed number of ports per node.

times less than a time-slot, which may be on the order of μs . 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. 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 [17], [18]. 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 [19], 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 Section VIII, we extend this model to allow the traffic to dynamically change, but on a slow enough time-scale so that a new TDM schedule may be found.

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. We use the terms “tunable ports” and “tunable transceivers” interchangeably to refer 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. Our analysis of this problem is based on relating feasible schedules to edge colorings of a corresponding “traffic graph.” In Section IV, we discuss this correspondence and review several useful facts about edge colorings in graphs. Using this correspondence with edge coloring, we prove that with a limited number of wavelengths, the minimum tunable port problem is NP-complete. However, we show that this problem is efficiently solvable in a number of important cases including uniform traffic between an even number of nodes and when all traffic is to and from a single hub 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. Finally, we show that tunability is also beneficial when the traffic can dynamically change among a set of allowable traffic matrices, as in [3]. Both rearrangeable and nonrearrangeable rings are considered. In the rearrangeable

case, tunability can again reduce the needed number of ports by more than 60%. Without rearrangements, it is shown that the port savings can be arbitrarily large.

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.² 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 duplex. 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. In 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 the traffic is not symmetric, the number of transmitters and receivers at each node may be different. 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. For a unidirectional ring, we assume that the nodes are numbered sequentially around the ring in the direction that data is transmitted. 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} may be 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 terms “port” and “transceiver” 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

²We refer to each unidirectional connection as a “circuit”; a duplex connection contains two such circuits.

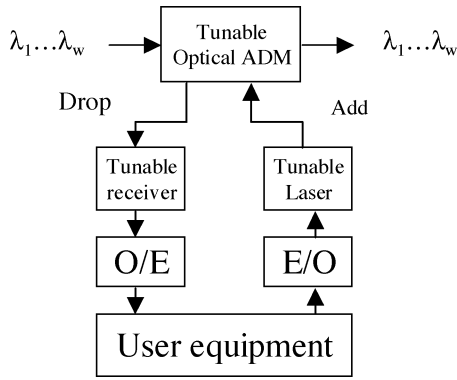


Fig. 1. Example node with tunable transceiver.

TABLE I
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)

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)

be assigned to these two wavelengths. With $g = 3$, as many as six circuits can be assigned to each wavelength; this can be accomplished by assigning both circuits for each duplex connection to same time-slot on the same wavelength. 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 W_{\min} wavelengths. Without the possibility of tunable transceivers, the assignment of circuits to wavelengths corresponds to the standard traffic grooming problem considered in [1]–[16]. 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 . In this assignment, each node must transmit and receive on both wavelengths. As a result, two transceivers are needed per node for a total of eight transceivers. A slightly more clever assignment, shown in Table II, only requires seven transceivers, as node 1 is only assigned to transmit and receive on λ_1 .

In this simple example, the above assignment is the best possible without tunability and results in a savings of one 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

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)

two and a total of six 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 seven to four 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 nontrivial combinatorial problem.

In the above example, the number of transceivers was reduced by 42% by using tunable components. This means that for tunable components to be cost-competitive for this example they need to cost less than 2.3 times their fixed-tuned counterparts. Our goal in this paper is to characterize this cost tradeoff for a range of different networks and traffic demands. This requires developing algorithms for assigning a given traffic requirement to time-slots/wavelengths so that the minimum number of tunable ports is needed. Here, we consider the case where all the ports in the network are tunable. Given that fixed-tuned transceivers will likely continue to be cheaper than tunable ones, further cost savings might be gained by using a hybrid architecture with some tunable and some fixed-tuned transceivers. For example, in Table III, node 1 is always transmitting on wavelength λ_1 and so does not require a tunable transceiver. Thus, this same assignment could be realized with three tunable and one fixed-tuned transceiver. This type of hybrid approach is outside of the scope of this paper, though some of the techniques used here may be useful in that context as well. The optimal hybrid architecture will depend strongly on the relative costs of the two types of transceivers. Also, practically, a hybrid approach may not be desirable as it could entail an increase in the operational cost of the network by requiring carriers to store on inventory both types of transceivers.³

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 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

³Another benefit of tunable transceivers is that a single transceiver can be used for all wavelengths eliminating the need for an inventory of different transceivers.

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 μ 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 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 μ s; these devices must be able to tune in sub- μ s 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., [20] and [21]). Fast-tunable receivers (i.e., filters) and associated issues such as fast clock recovery [20] 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 nonsymmetric 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)-constant 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 constants are needed only for simplicity of the presentation. The desired optimization problem is then

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

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

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

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

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

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

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.⁴ Note that the number of available wavelengths W does not appear as an explicit constraint in this problem; rather it enters the problem by determining the number of variables used, i.e., the variables $T_{i,j,m,k}$.

We also consider a restricted version of the problem with the assumption that every bidirectional pair must be assigned to the same wavelength/time-slot, as in the examples from Section 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

⁴Note, we do not need to constrain X_i to be an integer because the minimization will ensure this.

TABLE IV
OPTIMAL SOLUTION TO THE MTP PROBLEM
WITH ASYMMETRIC ASSIGNMENTS

	λ_1	λ_2	λ_3	λ_4
Slot 1	(1→3),(3→1)	(1→3)	(2→1)	(3→2)
Slot 2	(1→3),(3→1)	(1→2),(2→3)	(3→1)	

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

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

and finally, the constants $E_{(i,j)}^l$ 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. The following example shows that it can be strictly less. Consider a ring with $N = 3$ nodes, a granularity of $g = 2$, and $W = 4$ available wavelengths. Suppose that the traffic requirement is $R_{1,3} = R_{3,1} = 3$, $R_{1,2} = R_{2,1} = 1$, and $R_{2,3} = R_{3,2} = 1$. In this case, $\lceil W_1 \rceil = \lceil W_3 \rceil = 2$ and $\lceil W_2 \rceil = 1$. It can be shown that for the MTPS problem one node must use $\lceil W_i \rceil + 1$ transceivers, and so the solution to the MTPS problem is six transceivers. On the other hand, Table IV shows an optimal solution for the MTP problem requiring only $\lceil W_i \rceil$ transceiver per node, for a total of five transceivers. Here, $(i \rightarrow j)$ indicates a unidirectional circuit from i to j . Note that in slot 2 on wavelength λ_2 both $(1 \rightarrow 2)$ and $(2 \rightarrow 3)$ can be assigned to this slot, since they do not overlap on the ring.⁵ Though in general these two problems may have different solutions, in the following, we will see that in several important cases they have the same solution.

IV. GRAPH THEORETIC FORMULATIONS

In this section, we relate solutions of the MTP and MTPS problems to *edge colorings* of a corresponding “traffic graph.” This relationship will be the basis of our analysis in the subsequent sections. We will also review several useful results from the graph theory literature about specific types of edge colorings.

To begin, we define two different representations of the traffic requirement R in terms of a corresponding multigraph.⁶ The first representation we refer to as a *duplex traffic graph*; this

⁵Also, note that in this example, the transmitter and receiver of a node may be tuned to different wavelengths during a given time-slot, e.g., in time-slot 1, node 2 transmits on λ_3 and receives on λ_4 .

⁶A multigraph is a graph that may have multiple parallel edges between some of the nodes.

is a (undirected) multigraph $G_d = (V, E)$, where each node $i = 1, \dots, N$ in the ring is represented by a vertex $v_i \in V$ with $R_{i,j}$ edges $(v_i, v_j) \in E$ corresponding to each duplex circuit required between nodes i and j . In other words, the traffic requirement R is the adjacency matrix for the duplex traffic graph G_d . The second representation is a *simplex traffic graph*. This is a bipartite multigraph $G_s = (V, E)$, where the vertex set V is the union of two disjoint sets $C = \{v_1, \dots, v_N\}$ and $D = \{w_1, \dots, w_N\}$ each corresponding to the set of nodes $\{1, \dots, N\}$. The edge set E is a multiset consisting of $R_{i,j}$ edges between each node $v_i \in C$ and $w_j \in D$. In the duplex traffic graph G_d , both circuits in each duplex connection are represented by a single edge; in the simplex traffic graph G_s , each unidirectional circuit is represented by a separate edge. Fig. 2 shows an example of these two graphs for the case of uniform traffic demand of one circuit between $N = 5$ nodes.

For a given multigraph $G = (V, E)$, an edge coloring with K colors $\{c_1, \dots, c_K\}$ is any assignment of the colors to the edges of G . Unless explicitly mentioned, we do not require that an edge coloring be *proper*, i.e., that no two edges incident at the same node have the same color. Given an edge coloring with K colors, for each vertex $v \in V$, let $c_i(v)$ denote the number of edges incident with v assigned color c_i . Likewise, for each color c_i , let $\mathcal{E}(c_i)$ denote the number of edges in G assigned color c_i .

Any feasible solution to either the MTPS or MTP problem can be related to an edge coloring of the simplex traffic graph, G_s with g colors by identifying each of the g time-slots with a unique color. A coloring of the traffic graph is then obtained by assigning each edge the color of the time-slot that the corresponding circuit is assigned to. Given this edge coloring, the number of transceivers needed by node i , X_i , is the maximum number of edges assigned the same color incident at either of the corresponding vertices $v_i \in C$ or $w_i \in D$, i.e.,

$$X_i = \max_{k=1, \dots, g} \{ \max \{ c_k(v_i), c_k(w_i) \} \}. \quad (7)$$

By the same construction, a feasible solution to the MTPS problem determines an edge coloring of the duplex traffic graph, G_d . This is because in the MTPS problem both circuits in a duplex connection are always assigned to the same wavelength/time-slot. In this case

$$X_i = \max_{k=1, \dots, g} c_k(v_i). \quad (8)$$

Note that a feasible solution to the MTP problem cannot be identified with an edge coloring of G_d , unless it is also a feasible solution to the MTPS problem.⁷

This edge coloring representation of a solution to the MTP or MTPS problem indicates the time-slot for each circuit but does not indicate the wavelength that the circuit is assigned to. In other words, an edge coloring of the traffic graph does not explicitly determine a solution (e.g., the values of $T_{i,j,m,k}$) to the MTP or MTPS problems. In the case of an edge coloring for G_d , this information is not critical. This is because each edge represents a bidirectional circuit that takes up one time-slot/wavelength around the ring; hence, each bidirectional circuit can

⁷By definition, any feasible solution to the MTPS problem with W wavelengths will also be a feasible solution to the MTP problem with W wavelengths.




Fig. 2. (a) Example of the duplex traffic graph G_d and (b) simplex traffic graph G_s corresponding to uniform traffic between $N = 5$ nodes.

be arbitrarily assigned to any available wavelength without affecting its feasibility. This will require $\max_{i=1,\dots,g} \mathcal{E}(c_i)$ wavelengths. It follows that any edge coloring of G_d with g colors gives a feasible solution to the MTP and MTPS problems with W available wavelengths if and only if

$$\max_{i=1,\dots,g} \mathcal{E}(c_i) \leq W. \quad (9)$$

Therefore, the MTPS problem can be viewed completely in terms of finding an edge coloring of G_d that satisfies (9) and minimizes (8). On the other hand, for the MTP problem we must, in general, consider edge colorings of G_s . In this case, the assignment of circuits to wavelengths can be critical for ensuring that the given solution is feasible. For example, in the assignment in Table IV, the circuits transmitted in each time-slot must also be correctly packed into the available wavelengths. Since an edge coloring of G_s does not specify this information, the MTP problem can not, in general, be viewed simply in terms of edge colorings.

In the following, it will be useful to consider several specific types of edge colorings. An edge coloring of a multigraph $G = (V, E)$ with K colors is defined to be *equitable* [23], [24], if for each node $v \in V$ and any two colors c_k and c_l

$$|c_k(v) - c_l(v)| \leq 1 \quad (10)$$

i.e., the number of edges incident at a node assigned any color differs from the number assigned any other color by at most one. Note that a proper edge coloring is always equitable since $c_k(v)$ will be either 0 or 1. The following gives a relation between an equitable edge coloring of a traffic graph and the number of tunable ports needed.

Lemma 1: If a solution to either the MTP or MTPS problem can be represented as an equitable edge coloring of G_d or G_s with g colors, then in this solution each node i needs exactly $\lceil W_i \rceil$ tunable ports.

Proof: Let $d(v)$ denote the degree of vertex v in a given multigraph G . An equivalent definition of an equitable edge coloring with K colors is that for each node v and each color c_k

$$\left\lfloor \frac{d(v)}{K} \right\rfloor \leq c_k(v) \leq \left\lceil \frac{d(v)}{K} \right\rceil.$$

In G_d and G_s , $d(v_i) = R_i$. Therefore, if these graphs have an equitable edge coloring with g colors, it follows that for all i and k , $c_k(v_i) \leq \lceil R_i/g \rceil = \lceil W_i \rceil$. In the case of G_s , by the same reasoning, we have $c_k(w_i) \leq \lceil W_i \rceil$. Therefore, from (7) and (8), it follows that $X_i \leq \lceil W_i \rceil$. This must be met with equality, since $\lceil W_i \rceil$ is a lower bound on the number of ports needed. ■

Conditions for when a graph has an equitable edge coloring can be found in [23] and [24]. In particular, we note the following result from [23].

Lemma 2: For any integer $K \geq 2$, a bipartite multigraph always has an equitable edge coloring with K colors.

Another type of edge coloring is referred to as an *equalized* edge coloring. This is defined as a coloring of a multigraph G such that for any two colors c_k and c_l

$$|\mathcal{E}(c_k) - \mathcal{E}(c_l)| \leq 1 \quad (11)$$

i.e., the number of edges in the whole graph assigned any color differs from the number of edges assigned any other color by at most one. The following lemma shows that equalized edge colorings are related to the number of wavelengths needed in the entire network.

Lemma 3: If a solution to either the MTP or MTPS problem can be represented as an equalized edge coloring of G_d with g colors, then this solution requires only $\lceil W_{\min} \rceil$ wavelengths.

Proof: In this case, (11) is equivalent to requiring for each color c_k

$$\left\lfloor \frac{|E|}{K} \right\rfloor \leq \mathcal{E}(c_k) \leq \left\lceil \frac{|E|}{K} \right\rceil$$

where $|E|$ is the number of edges in the graph and K is the number of colors. For G_d , $|E| = \sum_i R_i/2$. And so, if G_d has an equalized edge coloring with g colors, then $\mathcal{E}(c_k) \leq \lceil \sum_i R_i/(2g) \rceil = \lceil W_{\min} \rceil$. From (9), it follows that this solution requires only W_{\min} wavelengths. ■

In [22], the following relationship between equitable and equalized edge coloring is given.⁸

Theorem 1: If a multigraph $G = (V, E)$ has an equitable edge coloring with K colors, then there exists an edge coloring with K colors that is both equalized and equitable.

V. 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 simple graph,⁹ G . The EDGE COLORING problem is to find the smallest K such that G has a proper edge coloring with K colors. The solution is called the chromatic index of the graph. Vizing's theorem [31], states that a simple 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 [25]. Moreover, for any $\Delta \geq 3$, EDGE COLORING is still NP-complete when restricted to the set of graphs with maximum degree Δ [26].

Theorem 2: For all $g \geq 3$, the MTPS problem is NP-complete.

⁸The result in [22] is for a more general class of *I-regular* coloring, of which equitable coloring are a special case.

⁹A simple graph is a graph with no parallel edges between the same pair of nodes, i.e., it is not a multigraph.




Fig. 3. Complete graph corresponding to a uniform traffic requirement of one circuit between all nodes. The indicated edge coloring corresponds to the slot assignment given in Table III.

Proof: To prove this, we show that any instance of EDGE COLORING can be transformed into an instance of the MTPS problem. Specifically, we can view any simple graph G as the duplex traffic graph G_d for an instance of the MTPS problem. (Since G is simple, $R_{i,j} \leq 1$ for all i, j .) In the MTPS problem, 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 is met with equality. With $X_i = 1$, from (8), it follows that the corresponding edge coloring of G satisfies $c_k(v_i) \leq 1$ for all k and i , i.e., this must be a proper edge coloring of G with g colors. And so, from Vizing's theorem, the chromatic index must be g .

Next, suppose that G 's chromatic index is g . In this case, a proper edge coloring using g colors can be mapped back into a time-slot assignment for the MTPS problem using exactly one transceiver for each node. Also, for each color c_k , it must be that $\mathcal{E}(c_k) \leq \lfloor N/2 \rfloor$ (since each edge takes-up two 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. 3, where G is a 4 node complete graph. 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 Section II. The MTPS solution for this traffic, using three slots, was given in Table III; this corresponds to the proper edge coloring shown in Fig. 3, where slot j in the table is identified with color c_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 4: 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 a multigraph 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 bidirectional 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). Again identifying time-slots with colors, the given solution to the MTPS problem will correspond to a proper edge coloring of H using 2 colors. Since a proper edge coloring is equitable, from Theorem 1, it follows that H has a equalized and equitable edge coloring. In an equalized edge coloring, it must be that $W - 1$ edges are assigned each color. This can be mapped back into a new time-slot assignment using $W - 1$ wavelengths for these two time-slots and requiring no additional ports. This shows that the desired solution to the MTPS problem must exist.

Next, we give polynomial time algorithm for constructing this solution. Since the maximum degree of H is 2, 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). Since H has a proper edge coloring with two colors, each cycle must have an even length. Also, H will have an even number $(W - 2 + W)$ of edges; hence, the number of paths with odd lengths must be even. We construct an equitable and equalized edge coloring of this subgraph with two 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 the desired time-slot assignment with exactly $W - 1$ edges being assigned each color. 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 V; this assignment is for a ring with $\lceil W_{\min} \rceil = 2$ and $N = 3$ nodes. One possible version of the

TABLE V
EXAMPLE OF TIME-SLOT ASSIGNMENT

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




Fig. 4. Graph H corresponding to the time-slot assignment in Table V. Each edge is labeled with $x(y)$, where x is the original time-slot assignment and y is the new assignment.

TABLE VI
NEW TIME-SLOT ASSIGNMENT AFTER APPLYING
THE APPROACH IN LEMMA 4

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

corresponding graph H is shown in Fig. 4.¹⁰ Each node is represented by two nodes in the graph since two 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 three 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 new time-slot assignment is shown in Table VI and requires only two wavelengths as desired.

A direct corollary of Lemma 4 is that the MTPS problem remains NP-complete when the wavelength limit is set to the minimum value. The next lemma shows that when W_{\min} is an integer, the MTP problem and MTPS problem have the same solution.

Lemma 5: If W_{\min} is an integer,¹¹ then any feasible solution to the MTP problem using $W = \lceil W_{\min} \rceil = W_{\min}$ wavelengths must also be a feasible 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 3: 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 2, we again use the a correspondence with the EDGE COLORING problem.

VI. 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, most of the results of this section are not topology dependent and so can be applied to a general network topology. In this section, we again use the relationship between edge coloring and the MTP problem; however, here we consider this in terms of the simplex traffic graph G_s .

Theorem 4: The optimal solution to the MTP problem with no wavelength limitations is for each node to use $X_i = \lceil W_i \rceil$ tunable ports.

Proof: As noted in Section IV, a solution to the MTP problem can be represented as an edge coloring of the simplex traffic graph G_s . This graph is always a bipartite multigraph. From Lemma 2, this graph always has an equitable edge coloring with g colors. As discussed in Section IV, in general, this edge coloring might not correspond to a feasible solution to the MTP problem because it does not account for whether the transmitted circuits can be correctly packed onto the available wavelengths. However, if there is no wavelength limitation then this is not a concern, i.e., each unidirectional circuit can be assigned to its own wavelength/time-slot. Thus, with no wavelength limitation, any edge coloring of G_s , can be mapped back into a feasible solution to the MTP problem. Since G_s has an equitable edge coloring with g colors, it follows from Lemma 1 that this corresponds to a solution to the MTP problem, where each node needs exactly $\lceil W_i \rceil$ tunable ports. ■

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. Also notice, that from Lemma 4, this is not the case for the MTPS problem; i.e., increasing the available wavelengths does not simplify the problem if symmetric assignments are still required.

In preceding proof, we used that fact that G_s always has an equitable edge coloring with g colors. From Theorem 1, it follows that G_s has an equitable and equalized edge coloring with g colors. Consider a traffic assignment corresponding to this equalized edge coloring, where again each unidirectional circuit is assigned to its own wavelength/time-slot. This will again require that each node has $X_i = \lceil W_i \rceil$ tunable ports. Furthermore, in a unidirectional ring as in the proof of Lemma 3, the number of edges assigned to each color c_k

$$\mathcal{E}(c_k) \leq \left\lceil \frac{|E|}{g} \right\rceil$$

where $|E| = 2gW_{\min}$ is the number of edges in G_s . This will require at most $\lceil 2W_{\min} \rceil$ wavelengths. Therefore, we have the following corollary.

¹⁰In general, H is not unique but depends on the assignment of circuits to the available transceivers; any such assignment can be used.

¹¹Recall that W_{\min} was defined in units of fractional wavelengths.




Fig. 5. Example of the graph G'_s corresponding to G_s in Fig. 2 with $g = 3$. The two children of each node i in G_s are labeled by i and i' in G'_s .

Corollary 1: In a unidirectional ring with symmetric traffic, the optimal solution to the MTP problem with $W \geq \lceil 2W_{\min} \rceil$ is for each node to use $X_i = \lceil W_i \rceil$ tunable ports.

In other words, approximately twice the minimum number of wavelengths are required for this approach. This is because we are not packing more than one circuit into a wavelength/time-slot. Indeed, using this type of approach, the above bound is the best possible in the sense that if fewer wavelengths were available, some circuits must share a wavelength/time-slot.

The above results show that with sufficient wavelengths available the optimal solution to the MTP problem is always to have each node use $\lceil W_i \rceil$ transceivers. Next, we discuss an efficient (polynomial time) algorithm for finding this solution. Here, we assume that there is not a wavelength limit; the resulting solution can then be converted into a solution that satisfies a wavelength limit $W \geq \lceil 2W_{\min} \rceil$ by applying the same type of algorithm as in the proof of Lemma 4.

To begin, consider the case where $W_i \leq 1$ for all $i = 1, \dots, N$. In this case, the maximum degree of G_s will be less than or equal to g , and from Theorem 4, the optimal solution to the MTP problem will be $X_i = 1$. Therefore, any proper edge coloring of G_s with g colors will give an optimal time-slot assignment. A bipartite graph always has a chromatic index equal to its maximum degree [27], and polynomial time algorithms for finding a proper edge coloring using this many colors can be found in [28]–[30]. Applying one of these algorithms to G_s gives us the desired time-slot allocation.

Next, we extend this approach to the case where some nodes may generate more than one wavelength worth of traffic. When $W_i > 1$ for some node i , the corresponding node in G_s will have a degree greater than g . Hence, a proper edge coloring of G_s will require more than g colors.¹² In this case, we decompose G_s into a set subgraphs, each of which have maximum degree g . Specifically, we construct a new bipartite graph $G'_s = (C', D', E')$ as follows: for each node $v_i \in C$, we put $\lceil W_i \rceil$ “children” nodes in C' ; likewise, for each node $v_i \in D$, we put $\lceil W_i \rceil$ “children” nodes in D' . For each edge $(v_i, w_j) \in E$, we place an edge in G'_s between one of the children of v_i and one of the children of w_j such that no node in G'_s has a degree larger than g . The assignment of edges to children can be done in a “greedy” fashion, i.e., we arbitrarily order the children and assign each edge to the first available child that has degree less than g . This can be done because the total amount traffic to be assigned to the $\lceil W_i \rceil$ children of a node v_i must be less than $\lceil W_i \rceil g$ by definition. The

¹²At first one might think that a proper 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.

graph G'_s will again have a chromatic index less than or equal to g . A proper edge coloring with g colors corresponds to a solution to the MTP problem, where each node i will need one tunable port for each of its $\lceil W_i \rceil$ children in G'_s . Therefore, finding an optimal solution to the MTP problem requires constructing the bipartite graph G'_s and finding a proper 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 simplex traffic graph G_s is given in Fig. 2(b). This graph has a maximum degree of $N - 1 = 4$ and, hence, has a proper coloring with four 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 Theorem 4, each node will require $\lceil W_i \rceil = 2$ tunable ports. Using the above algorithm, each node is split into two children. A corresponding graph¹³ G'_s is shown in Fig. 5; this graph has maximum degree of 3 and so can be colored using $g = 3$ colors.

Notice that in the above approach, the assignments can be arranged so that each node (child) in the graph G'_s 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 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. Likewise, since we are always using the simplex traffic graph, these results apply when the traffic is nonsymmetric as well.

VII. LIMITED WAVELENGTHS

When wavelengths are limited, the time-slot allocation of the previous section will no longer be feasible and the circuits 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., $W = \lceil W_{\min} \rceil$. 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 5, we know that if W_{\min} is an integer, then there is no loss in performance by restricting ourselves to such an assignment. For noninteger 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

¹³There are many different ways to construct the graph G'_s ; Fig. 5 illustrates one possible construction.

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 a uniform traffic requirement of exactly r circuits between each pair of nodes, i.e., $R_{i,j} = r$, for all $i \neq j$. In this case, each node i generates $W_i = (N - 1)r/g$ wavelengths of traffic. 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 5: 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 time polynomial in N and r .

Proof: Since the traffic demand is uniform, the duplex traffic graph, G_d , will be a complete multigraph with $r = R_{i,j}$ edges between each pair of vertices. A complete multigraph with an even number of vertices has a chromatic index equal to its degree [32]. Thus, we can find a proper 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 time polynomial in N and r . 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

$$\lceil W_i \rceil \left(\frac{N}{2} \right) g \geq W_i \left(\frac{N}{2} \right) g = \frac{N(N-1)r}{2}$$

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} may be greater than $\lceil W_{\min} \rceil$ so that the allocation uses more than the required number of wavelengths. However, from Lemma 4, the allocation can be transformed in polynomial time into an allocation using only $\lceil W_{\min} \rceil$ wavelengths. ■

In Theorem 5, 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.

Example: Consider a ring with $N = 6$ nodes, a granularity of $g = 3$, and a uniform demand of $r = 1$ circuit between each pair of nodes. In this case, $W_i = 5/3$ and $W_{\min} = 5$. The corresponding graph G_d has a chromatic index of 5. Applying an edge coloring algorithm from [32], we get the coloring of G_d using five colors shown in Table VII. Since $\lceil W_i \rceil = 2$, each node needs two 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 VIII. Using Lemma 4, this can be transformed into an assignment using $W_{\min} = 5$ wavelengths, as shown in

TABLE VII
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)

TABLE VIII
TIME-SLOT ASSIGNMENT USING SIX 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 IX
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)

Table IX. 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.

Theorem 5 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. Furthermore, at least one node in a ring with an odd number of nodes and uniform traffic may be required to have more than $\lceil W_i \rceil$ tunable ports. For example, consider a ring with $N = 3$ nodes, a granularity of $g = 2$, and a uniform traffic demand of $r = 1$ circuit between each pair of nodes. In this case, $W_i = 1$ and $W_{\min} = 3/2$. However, it can be seen that one of the three nodes will require $\lceil W_i \rceil + 1 = 2$ tunable ports, while the other two nodes will require only $\lceil W_i \rceil$. The following theorem shows that, in general, this is the case in a ring with uniform traffic and N odd, i.e., at most one node will require $\lceil W_i \rceil + 1$ tunable ports.

Theorem 6: In a ring with N odd, granularity g , $\lceil W_{\min} \rceil$ wavelengths, and a uniform traffic requirement of r circuits between each pair of nodes, a feasible time-slot allocation can be found where at most one node requires $\lceil W_i \rceil + 1$ tunable transceivers and all other nodes will require $\lceil W_i \rceil$.

In certain cases, with N odd, every node will only require $\lceil W_i \rceil$ tunable transceivers. The next theorem gives several sufficient conditions for when this may occur.

Theorem 7: For a ring as in Theorem 6, let $Q = (N - 1)r \bmod (2g)$. Any of the following conditions are sufficient for every node to require exactly $\lceil W_i \rceil$ tunable transceivers:

- if $Q > g$ or $Q \leq (2/3)g$;
- if $r = 1$ and $Q \neq g$;
- if $N > g$ and $Q \neq g$.

A single proof for both theorems 6 and 7 is given in Appendix III. As an example, consider a ring with N nodes, a granularity of $g = 4$, and $r = 1$. In this case, $Q = (N - 1) \bmod 8$, and so from Theorem 7 case b) or c), the

only time one node may need an extra port is when $Q = 4$, i.e., when $N = 5, 13, 21 \dots$

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 to or from a “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 8: 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: With egress traffic, G_d will be a bipartite multigraph, where all edges are between the vertex for single hub node and the vertices for the other nonhub nodes. As noted in Section VI, the chromatic index of a bipartite multigraph 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.¹⁴ ■

Theorem 8 extends directly to a (symmetric) 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 duplex traffic graph is bipartite (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 nonegress 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 bidirectional circuits between each nonhub node and the hub. In this case, for each nonhub 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

$$\left\lceil \frac{(N - 1)r}{g} \right\rceil + (N - 1) \left\lceil \frac{r}{g} \right\rceil \quad (12)$$

tunable ports. If $W_i \leq 1$, in [1], it is shown that without any wavelength limits the number of fixed-tuned transceivers needed for egress traffic is minimized if each nonhub node only uses one transceiver. Therefore, if $r > g/2$, the minimum number of fixed-tuned transceivers is $2(N - 1)$. Comparing this with (12), it can be seen that in both cases the nonhub nodes require one port each; however, with tunability the number of ports required at the hub is reduced from $N - 1$ to $\lceil (N - 1)r/g \rceil$, which can be much less. For example, with $r = g/2 + 1$, the number of tunable transceivers required by the hub is approximately $(1/2) + (1/2g)$ times the number of fixed-tuned transceivers required; for large values of g , this approaches a 50% reduction in the number of transceivers at the hub (a 25% reduction in the total number of transceivers).

¹⁴In fact, with a single hub node the time-slot assignment can be found directly using a simple greedy algorithm.

Moreover, with tunable transceivers this is achieved using the minimum number of wavelengths, while this fixed-tuned case may require more wavelengths. If the number of wavelengths are limited to W_{\min} , then with fixed tuned transceivers, some nodes may have to transmit on multiple wavelengths. In this case, tunability can reduce the requirements of the nonhub nodes.

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, G_d will be a simple graph. For simple graphs, from Vizing’s theorem, the chromatic index is at most $\Delta + 1$, where Δ is the maximum degree. Moreover, polynomial time algorithms for finding a proper edge coloring of any simple graph with $\Delta + 1$ colors are known [31]. We use this fact to develop efficient time-slot assignment algorithms for the ring. To begin, we construct G_d based on the traffic requirement. Following a similar approach as in Section VI, we then construct a new graph G'_d by replacing each vertex v_i in G_d with $\lceil R_i/(g - 1) \rceil$ “children” in G'_d , and each edge in G_d with an edge between two of the corresponding children in G'_d so that no node in G'_d has a degree greater than $g - 1$. In this case, G'_d has a proper edge coloring with g colors; this edge coloring can be used for a time-slot assignment, where each node in G'_d 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 $(1/2)(\sum_i \lceil R_i/(g - 1) \rceil)$. This may be greater than $\lceil W_{\min} \rceil$, but, using Lemma 4, it can be transformed into an assignment using $\lceil W_{\min} \rceil$ wavelengths without increasing the number of ports. To summarize, we have shown the following.

Theorem 9: For any 0–1 traffic requirement, a time-slot assignment using $\lceil W_{\min} \rceil$ 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

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

for each node i . This is at most 2 and in many cases will be equal to 1.

Example: Consider a ring with $N = 5$ nodes a granularity of $g = 3$, and assume that the traffic is a uniform demand of $r = 1$ circuit between each pair of nodes. In this case, $W_i = 4/3$ and $W_{\min} = 10/3$. Here, N is odd, and applying Theorem 6, $Q = 0$ and so each node will require $\lceil W_i \rceil = 2$ transceivers. Suppose, we instead apply the heuristic from Theorem 9 to this example. Following this procedure, we split each node into $\lceil R_i/(g - 1) \rceil = 2$ children and form the graph G'_d , where each node has a degree no greater than $(g - 1) = 2$, as shown in




Fig. 6. Example graphs. On the left is the original graph G_d ; on the right is the derived graph G'_d which has a maximum degree of 2. The two “children” of each node i in G_d are labeled by i and i' in G'_d . An edge coloring of G'_d using three colors (c_1 , c_2 , and, c_3) is also shown.

TABLE X
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. A proper edge coloring of G'_d using $g = 3$ colors is also indicated in the figure. This coloring translates into the time-slot assignment shown in Table X, which uses two 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.

For a general traffic requirement, G_d will not be a simple graph. For a multigraph, Vizing’s theorem does not apply, but two other upper bounds on the chromatic index are known. The first, also due to Vizing [31], 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 [34], 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 with G_d and construct a new graph G'_d . This time we replace each vertex v_i in G_d with $\lceil (3R_i)/(2(g - 1)) \rceil$ “children” vertices in G'_d , and each edge in G_d with an edge between two of the corresponding children in G'_d so that no node in G'_d has a degree greater than $(2/3)g$. Hence, G'_d has a proper edge coloring with g colors and this can again be mapped into a time-slot assignment. Therefore, we have the following.

Theorem 10: For any symmetric traffic requirement, a (symmetric) time-slot assignment can be found (in polynomial time), where each node uses $\lceil (3R_i)/(2(g - 1)) \rceil$ tunable transceivers and $\lceil W_{\min} \rceil$ wavelengths.

D. Numerical Examples

We next present some numerical examples. First, we compare the number of tunable ports required to the number of fixed-tuned ports needed in a ring with uniform traffic. Fig. 7 shows the number of ports in a ring with $g = 4$ and a uni-

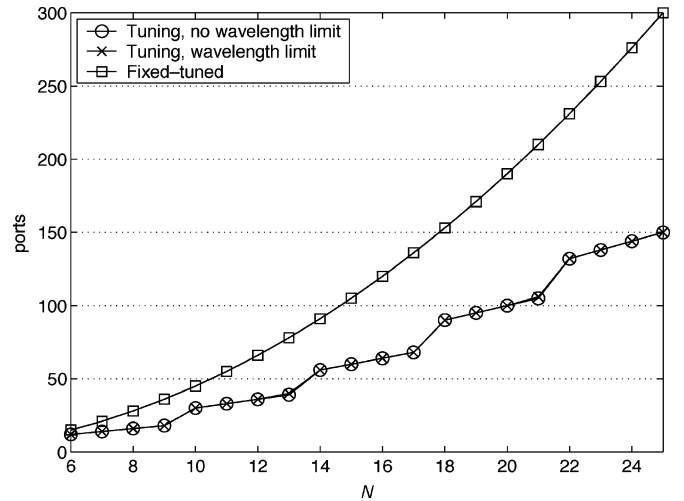


Fig. 7. Number of ports versus N for a ring with uniform demand of $r = 1$ circuits and $g = 4$.

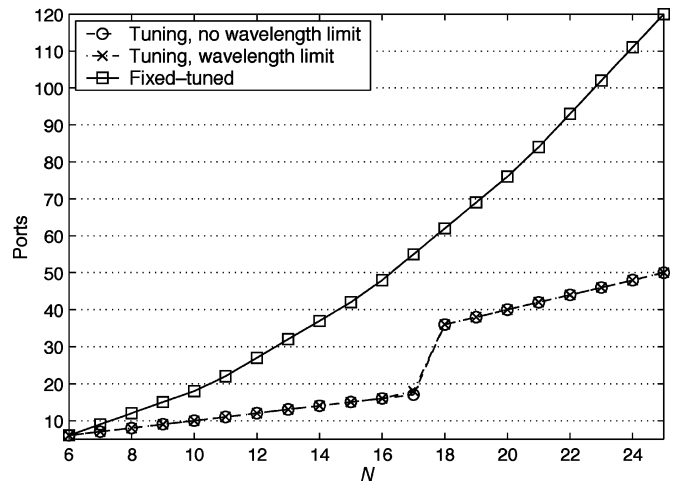


Fig. 8. Number of ports versus N for a ring with uniform demand of $r = 1$ circuits and $g = 16$.

form 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, Section III-B]; we note that in general this bound is not tight. The bottom two curves are the number of ports needed with tunable transceivers and either $\lceil W_{\min} \rceil$ wavelengths or with no wavelength restrictions. The number of ports needed with $\lceil W_{\min} \rceil$ wavelengths is found using theorems 5 and 7; with no wavelengths restrictions, the number of ports is given by Theorem 4. 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 from theorems 4 and 5 that the two cases should be equal; when N is odd, from Theorem 7, they will differ by one when $N = 5, 13, 21, \dots$, and otherwise are 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%.

Next, we present simulation results on the performance of our heuristic algorithm, described in Section VII-C, for arbitrary traffic demands. In particular, we consider the case of 0–1

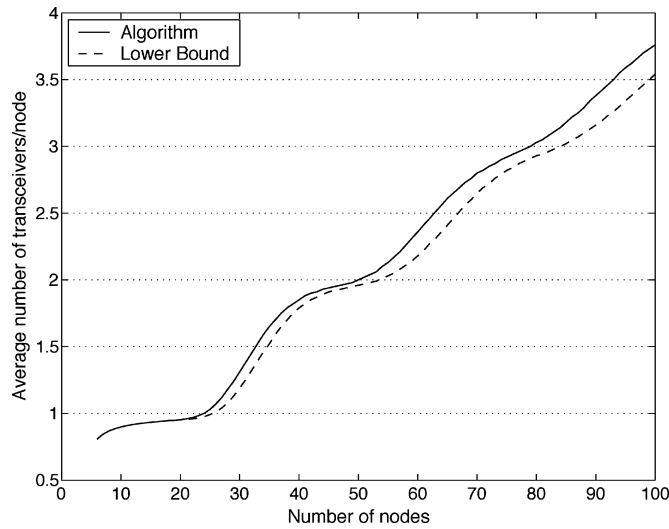


Fig. 9. Average number of transceivers per node for $p = 0.5$ and $g = 16$.

traffic, where $R_{i,j}$ is either 0 or 1 for all i and j . In this case, our algorithm, based on Vizing's theorem, is guaranteed to have a performance ratio that is bounded from the optimal by

$$\frac{\sum_i \left\lceil \frac{R_i}{(g-1)} \right\rceil}{\sum_i \left\lceil \frac{R_i}{g} \right\rceil} \leq \max_i \frac{\left\lceil \frac{R_i}{(g-1)} \right\rceil}{\left\lceil \frac{R_i}{g} \right\rceil}$$

which will be at most 2. To better understand the performance of the algorithm, we simulate its performance with a random traffic matrix, where each entry in the traffic matrix ($R_{i,j}$) is randomly generated according to a Bernoulli random variable with a traffic intensity parameter, p . Our simulations compare the performance of the algorithm to the lower bound (the denominator of the above expression), for various values of N , p , and g . For brevity, here we will present simulation results for $g = 16$. Similar results were obtained for other values of g as well.

We start with the case of $g = 16$ corresponding to 16 low-rate circuits per wavelength. Fig. 9 shows the number of transceivers per node when the traffic intensity is 0.5 and the number of nodes is varied from 6 to 100. The figure shows both the average number of transceivers used by the algorithms, as well as the lower bound. Fig. 10 shows the ratio of the two averages. The ratio of the algorithm to the lower bound oscillates, but tends toward an average value of $g/(g-1) = 16/15 = 1.0666$. The oscillatory behavior is a result from the fact that for certain values of N , R_i tends to be such that $\lceil R_i/g \rceil$ and $\lceil R_i/(g-1) \rceil$ are the same; while other values of N maximize the probability that the two are different. For example, note that the "peak" values are reached when Np is approximately equal to a multiple of g ; intuitively, this is the case where $\lceil R_i/g \rceil$ and $\lceil R_i/(g-1) \rceil$ are most likely to differ. The "lows" are reached at values of Np that are as far apart from being a multiple of g as possible. As can be seen from the figure, this effect diminishes as N increases; as the integer rounding becomes less significant of a factor in the overall number of transceivers. Fig. 11 shows the ratio of the average number of transceivers to the lower bound for $p = 0.7$; as can be seen from the figure, this ratio again oscillates with "peak" values that are around Np being a multiple of




Fig. 10. Ratio of the average number of transceivers used by the algorithm to the lower bound for $p = 0.5$ and $g = 16$.




Fig. 11. Ratio of the average number of transceivers used by the algorithm to the lower bound for $p = 0.7$ and $g = 16$.

g . Also, note that from Theorem 4, the lower bound can always be achieved with a sufficient number of available wavelengths. Hence, the above plots can also be viewed as a bounds on the average number of additional ports required to use the minimum number of wavelengths.

VIII. DYNAMIC TRAFFIC

In the previous sections, we assumed a single fixed traffic requirement. We now consider the case where there is a set of *allowable* traffic requirements, and our objective is to provision the ring with the minimum number of tunable ports to support any requirement in this set. The traffic demand may change over time by circuits arriving and departing; however, we assume that the resulting traffic at each time is still allowable. The set of allowable traffic is determined by limiting each node i to terminating at most t_i duplex circuits at any time; i.e., $R = [R_{i,j}]$ is allowable if it satisfies

$$\sum_i R_{i,j} \leq t_j, \quad \text{for all } j.$$

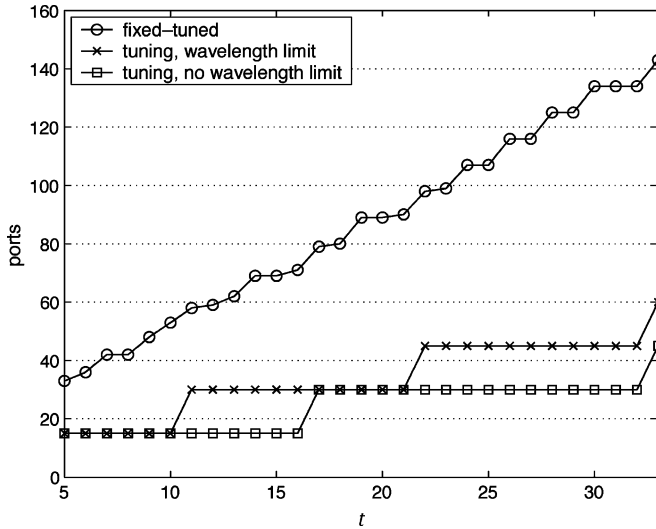


Fig. 12. Number of ports for t -allowable traffic versus t for ring with $N = 15$ and $g = 16$.

We refer to this set of traffic requirements as (t_1, \dots, t_N) -allowable or simply t -allowable when $t_i = t$ for all i . For (t_1, \dots, t_N) -allowable traffic, we define W_{\min} to be the maximum value of W_{\min} for any allowable traffic matrix. Traffic grooming for t -allowable traffic with fixed-tuned transceivers was studied in [3].

A. Rearrangeably Nonblocking Rings

We first consider the case where after each change in the traffic matrix, any remaining circuits can be reassigned to a different wavelength/time-slot. In this case, the ring is said to be rearrangeably nonblocking for all (t_1, \dots, t_N) -allowable traffic. A sufficient number of tunable transceivers for a ring to be rearrangeably nonblocking is given by the following theorem.

Theorem 11: A unidirectional ring with $\lceil W_{\min} \rceil$ wavelengths can support any (t_1, \dots, t_N) -allowable traffic requirement in a rearrangeably nonblocking manner if each node i uses $\lceil (3t_i)/(2(g-1)) \rceil$ tunable transceivers. Without wavelength limits, this traffic can be supported using $\lceil t_i/g \rceil$ tunable transceivers per node.

The first part of this theorem follows directly from Theorem 10; the second part from Theorem 4. Also, for each allowable traffic requirement, a corresponding time-slot assignment can be found (in polynomial time) by using the algorithms in the previous sections.

In Figs. 12 and 13, we compare the number of ports needed for a ring to be rearrangeably nonblocking for t -allowable traffic with and without tunable ports. Fig. 12 is for a ring with $N = 15$ nodes and $g = 16$. Fig. 13 is for a ring with $N = 8$ and $g = 16$. The top curve in each figure is a lower bound from [3, Lemma 3] on the number of fixed-tuned ports needed for $\lceil W_{\min} \rceil$ wavelengths (this is not tight in general). The lower curves are the number of tunable ports given by Theorem 11 for the wavelength limited and no wavelength limit cases. In the $N = 15$ case, it can be seen that more than a 60% reduction in the number of ports is achieved when compared with the bound




Fig. 13. Number of ports for t -allowable traffic versus t for ring with $N = 8$ and $g = 16$.

in [3]. Since the bound is not tight in general, it is likely that the actual reduction is even more significant.

B. Nonrearrangeable Nonblocking Rings

Next, we consider the case where the existing traffic cannot be rearranged when the traffic requirement changes. In this case, the ring is said to be either *strict-sense nonblocking* (if a new circuit can be assigned to any available wavelength/time-slot) or *wide-sense nonblocking* (if new circuits must be assigned according to a specified algorithm). The following gives a simple upper bound on the number of tunable ports needed for a ring to be strict-sense nonblocking with $\lceil W_{\min} \rceil$ wavelengths.

Theorem 12: A ring can support all (t_1, \dots, t_N) -allowable traffic without rearrangements on $\lceil W_{\min} \rceil$ wavelengths if each node i has $\min(t_i, \lceil W_{\min} \rceil)$ tunable transceivers.

Proof: Clearly, with t_i tunable ports a node will always have a port available when a new circuit arrives and it can use this port on any available time-slot/wavelength. Furthermore, a node will never need more tunable transceivers than the number of wavelengths in the ring. ■

Though this is a rather simple bound, it is still sufficient to show significant gains over the number of fixed-tuned ports needed. For example, in [3], it is shown that with fixed-tuned ports a strict-sense blocking ring with $\lceil W_{\min} \rceil$ wavelengths requires every node to have a port on every wavelength. Suppose that $t_i = 1$ for all i , then each node only requires one tunable port, but requires as many fixed-tuned ports as there are wavelengths; this can be arbitrarily large.

Next, we consider the case without wavelength restrictions; for this case, we restrict our attention to t -allowable traffic.

Theorem 13: With no wavelength restrictions, a ring can support all t -allowable traffic without rearrangements if each node has $\lceil (2t-1)/g \rceil$ tunable ports.

This follows because for any multigraph with maximum degree Δ , a "greedy" (proper) edge coloring algorithm that simply assigns each edge any available colors requires at most $2\Delta - 1$ colors [35]. With t -allowable traffic, at any time the duplex

traffic graph G_d for current traffic requirement will have a maximum degree of t and so $2t - 1$ colors are sufficient to find a proper edge coloring of G_d using a greedy algorithm. An edge coloring will again correspond to a time-slot allocation; this will require $\lceil (2t - 1)/g \rceil$ tunable ports per node. Furthermore, since we are using a greedy coloring algorithm, when the traffic changes, existing traffic need not be rearranged.

For certain choices of N and t , the number of ports in Theorem 13 can be shown to be the minimum possible for either a wide-sense or strict-sense nonblocking ring. This follows from a result for the on-line version of the edge-coloring problem given in [35]. In the on-line version of EDGE COLORING, the edges of the graph are revealed one edge at a time; an on-line algorithm must decide on a color for each edge as it is revealed based only on the past history. In [35], it is shown that for a graph with maximum degree Δ , there exists graphs for which an on-line algorithm requires $2\Delta - 1$ colors; a t -allowable traffic requirement that corresponds to such a graph will require $2t - 1$ colors and, hence, $\lceil (2t - 1)/g \rceil$ tunable ports. We also note that the above approach requires at most

$$W = \left\lceil \left(\frac{N}{2} \right) \left\lceil \frac{2t - 1}{g} \right\rceil \right\rceil$$

wavelengths.

IX. CONCLUSION

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 result 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 about 60% as compared with fixed tuned transceivers. Similar savings were found in the case where the offered traffic can change among an allowable set of traffic matrices.

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 tradeoff 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 tradeoff may become extremely beneficial.

APPENDIX I PROOF OF LEMMA 5

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 bidirectional 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

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

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

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

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 3

Proof: To prove this theorem, we again use a reduction from EDGE COLORING. Given a graph $G = (V, E)$ with maximum degree Δ , similar to the proof of Lemma 2, we view G as the duplex traffic graph for an instance of the MTP problem in a ring with $g = \Delta$. Hence, $W_{\min} = \sum_{v \in V} (d(v)/2g)$.

First, consider the case where W_{\min} is an integer and assume that there are W_{\min} wavelengths for the MTP problem. In this case, from Lemma 5, 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 Lemma 2, G has a chromatic index of g if and only if the solution to the MTPS problem with $W = \lfloor N/2 \rfloor$ wavelengths is N . Finally, from Lemma 4, 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 5. Instead, define a new graph G' as follows. Let $G' = (V', E')$ 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 $G_d = G'$ in a ring with granularity g and $W = W'_{\min} = \sum_{v \in V'} d(v)/(2g)$ wavelengths. 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 . ■

APPENDIX III PROOF OF THEOREMS 6 AND 7

Proof: For N odd, G_d will again be a complete multigraph with r parallel edges between each pair of vertices. In Theorem 5, we used a proper edge coloring of G_d to find a time-slot assignment. Here, we will instead use a decomposition of G_d into disjoint subgraphs, where each subgraph is a Hamiltonian cycle.¹⁵ A complete multigraph with an odd number of vertices can always be decomposed into $(N-1)r/2$ disjoint Hamiltonian cycles and this decomposition can be found in polynomial time [33]. First, consider assigning the circuits corresponding to each Hamiltonian cycle to one time-slot. In each cycle, each vertex will have two edges, and so will require two tunable ports. These ports can be reused for a total of g cycles, and so this requires

$$X'_i = 2 \left\lceil \frac{(N-1)r}{2g} \right\rceil \quad (13)$$

tunable ports per node.

Let $Q = (N-1)r \bmod (2g)$ and $K = \lfloor (N-1)r/2g \rfloor$, so that

$$(N-1)r = K(2g) + Q. \quad (14)$$

Note that both K and Q are integers and $Q < 2g$. Substituting this into (13), we have

$$X'_i = 2K + 2 \left\lceil \frac{Q}{2g} \right\rceil.$$

Also, note that

$$\lceil W_i \rceil = \left\lceil \frac{(N-1)r}{g} \right\rceil = 2K + \left\lceil \frac{Q}{g} \right\rceil.$$

¹⁵A Hamiltonian cycle in a graph is a cycle which spans every node in the graph.

Hence, if $Q > g$, we have $X'_i = \lceil W_i \rceil = 2K + 2$, and likewise if $Q = 0$, we have $X'_i = \lceil W_i \rceil = 2K$. Therefore, if $Q > g$ or $Q = 0$, then $X'_i = \lceil W_i \rceil$, as specified by Theorem 7 a).

Next, we consider the case where $0 < Q \leq g$. Here, $X'_i = 2K + 2$, but $\lceil W_i \rceil = 2K + 1$, i.e., the preceding approach results in each node using $\lceil W_i \rceil + 1$ tunable ports. In this case, we consider a different assignment of circuits to wavelength/time-slots which will require fewer ports. Specifically, we assign the circuits for Kg Hamiltonian cycles as above. Since there are $(N-1)r/2$ total Hamiltonian cycles, from (14), this will leave $(N-1)r/2 - Kg = Q/2$ Hamiltonian cycles for which the corresponding circuits have not yet been assigned. Let H_Q be the subgraph of G_d consisting of only those edges in the remaining $Q/2$ cycles; H_Q will be a multigraph, where each vertex has degree Q . It is known that the chromatic index of a multigraph is upper bounded by $(3/2)\Delta$, where Δ is the maximum degree [34]. Thus, if $Q \leq (2/3)g$, we can find a proper edge coloring of H_Q using g colors. This edge coloring gives a time-slot assignment of these circuits using one port per node for a total of

$$X''_i = 2 \left\lceil \frac{Kg}{g} \right\rceil + 1 = 2K + 1$$

tunable ports per node. Thus, for $Q \leq (2/3)g$, $X''_i = \lceil W_i \rceil$. Combined with the above, this shows that condition a) of Theorem 7 is sufficient for each node to require $\lceil W_i \rceil$ tunable transceivers.

Likewise, from this argument, whenever the subgraph H_Q has a proper edge coloring with g colors, then each node will require $\lceil W_i \rceil$ transceivers. If either $r = 1$ [Theorem 7b)] or $N > g$ [Theorem 7c)], then we can ensure that H_Q will be a simple graph with maximum degree $Q < g$. From Vizing's theorem, H_Q will have a proper edge coloring with g colors as desired.

Finally, to prove Theorem 6, assume that $(2/3)g \leq Q \leq g$ so that Theorem 7a) does not apply. In this case, H_Q may not have a proper edge coloring with g colors. Pick a vertex v_j in G_d . We can then find an edge coloring of each Hamiltonian cycle using two colors so that in each cycle only v_j will have two adjacent edges of the same color. This can be done by simply alternating colors around the cycle starting with an edge incident to v_j . Using two distinct colors for each Hamiltonian cycle requires a total of $2(Q/2) = Q \leq g$ colors. The corresponding edge coloring of H_Q again provides a time slot assignment of the circuits in the $Q/2$ remaining Hamiltonian cycles using one transceiver per node for every node except of v_j , which will require two tunable transceivers. This requires $X''_i = \lceil W_i \rceil$ tunable transceivers for all nodes $i \neq j$ and $X''_j = \lceil W_i \rceil + 1$, which completes the proof of Theorem 6. ■

REFERENCES

- [1] E. Modiano and A. Chiu, "Traffic grooming algorithms for minimizing electronic multiplexing costs in unidirectional SONET/WDM ring networks," in *Proc. CISS 98*, Princeton, NJ, Feb. 1998, pp. 653–658.
- [2] J. Simmons, E. Goldstein, and A. Saleh, "Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic," *IEEE J. Lightw. Technol.*, vol. 17, no. 1, pp. 48–57, Jan. 1999.
- [3] R. Berry and E. Modiano, "Reducing electronic multiplexing costs in SONET/WDM rings with dynamically changing traffic," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1961–1971, Oct. 2000.

- [4] —, “The role of switching in reducing the number of electronic ports in WDM networks,” *IEEE J. Sel. Areas Commun.*, vol. 22, no. 8, pp. 1396–1405, Oct. 2004.
- [5] O. Gerstel and R. Ramaswami, “Cost effective grooming in WDM rings,” in *Proc. IEEE INFOCOM 98*, San Francisco, CA, Apr. 1998, pp. 69–77.
- [6] O. Gerstel, P. Lin, and G. Sasaki, “Wavelength assignment in a WDM ring to minimize the cost of embedded SONET rings,” in *Proc. IEEE INFOCOM*, San Francisco, CA, Apr. 1998, pp. 94–101.
- [7] —, “Combined WDM and SONET network design,” in *Proc. IEEE INFOCOM*, New York, Mar. 1999, pp. 734–743.
- [8] T. Y. Chow and P. J. Lin, “The ring grooming problem,” *Networks*, vol. 44, pp. 194–202, 2004.
- [9] V. R. Konda and T. Y. Chow, “Algorithms for traffic grooming in optical networks to minimize the number of transceivers,” in *Proc. IEEE Workshop High Perform. Switch. Routing*, May 2001, pp. 218–221.
- [10] K. Zhu and B. Mukherjee, “Traffic grooming in an optical WDM mesh network,” *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 122–133, Jan. 2002.
- [11] —, “Traffic grooming in an optical WDM mesh network,” in *Proc. IEEE ICC*, Jun. 2001, pp. 721–725.
- [12] E. Modiano and P. Lin, “Traffic grooming in WDM networks,” *IEEE Commun. Mag.*, vol. 39, pp. 124–129, Jul. 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. Netw.*, vol. 8, no. 5, pp. 608–617, Oct. 2000.
- [14] H. Ghafouri-Shiraz, G. Zhu, and Y. Fei, “Effective wavelength assignment algorithms for optimizing design costs in SONET/WDM rings,” *IEEE J. Lightw. Technol.*, vol. 19, no. 10, pp. 1427–1439, Oct. 2001.
- [15] R. Dutta and G. N. Rouskas, “On optimal traffic grooming in WDM rings,” *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 110–121, Jan. 2002.
- [16] P. Wan, G. Calinescu, L. Liu, and O. Frieder, “Grooming of arbitrary traffic in SONET/WDM BLSRs,” *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1995–2003, Oct. 2000.
- [17] G. N. Rouskas and M. H. Ammar, “Dynamic reconfiguration in multihop WDM networks,” *J. High Speed Netw.*, vol. 4, no. 3, pp. 221–238, Jun. 1995.
- [18] A. Narula-Tam and E. Modiano, “Dynamic load balancing in WDM packet networks with and without wavelength constraints,” *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1972–1979, Oct. 2000.
- [19] C. Qiao and M. Yoo, “Optical burst switching—a new paradigm for an optical internet,” *J. High Speed Netw.*, vol. 8, no. 1, pp. 69–84, Jan. 1999.
- [20] K. Shrikhande *et al.*, “Performance demonstration of a fast-tunable transmitter and burst mode packet receiver for HORNET,” in *Proc. OFC*, vol. 4, 2001, pp. ThG2-1–ThG2-3.
- [21] M. Kauer *et al.*, “16-channel digitally tunable packet switching transmitter with subnanosecond switching time,” in *Proc. ECOC*, 2002, pp. 371–373.
- [22] D. de Werra, “On the use of alternating chains and hypergraphs in edge coloring,” *J. Graph Theory*, vol. 3, pp. 175–182, 1979.
- [23] —, “On the existence of generalized good and equitable edge colorings,” *J. Graph Theory*, vol. 5, pp. 247–258, 1981.
- [24] A. J. Hilton and D. de Werra, “A sufficient condition for equitable edge colorings of simple graphs,” *Discrete Math.*, vol. 128, pp. 179–201, 1994.
- [25] I. Holyar, “The NP-completeness of edge-coloring,” *SIAM J. Comput.*, vol. 10, no. 4, pp. 718–720, 1981.
- [26] Z. Galil and D. Leven, “NP-completeness of finding the chromatic index of regular graphs,” *J. Algorithms*, vol. 4, no. 1, pp. 35–44, Mar. 1983.
- [27] D. König, “Graphok és alkalmazásuk a determinánsok és a halmazok elméletére,” *Mathematikai és Természettudományi Értesítő*, vol. 34, pp. 104–119, 1916. [Hungarian].
- [28] H. N. Gabow and O. Kariv, “Algorithms for edge coloring bipartite graphs and multigraphs,” *SIAM J. Comput.*, vol. 11, pp. 117–129, 1982.
- [29] R. Cole and J. Hopcroft, “On edge coloring bipartite graphs,” *SIAM J. Comput.*, vol. 11, pp. 540–546, 1982.
- [30] A. Schrijver, “Bipartite edge-coloring in $O(\Delta m)$ time,” *SIAM J. Comput.*, vol. 28, no. 1999, pp. 841–846.
- [31] V. G. Vizing, “On an estimate of the chromatic class of a p-graph,” *Diskret. Analiz*, vol. 3, pp. 25–30, 1964.
- [32] F. Harary, *Graph Theory*. Reading, MA: Addison-Wesley, 1971.
- [33] J. Bosák, *Decompositions of Graphs*. Norwell, MA: Kluwer, 1990.
- [34] C. E. Shannon, “A theorem on coloring the lines in a network,” *J. Math. Phys.*, vol. 28, pp. 148–151, 1949.
- [35] A. Bar-Noy, R. Motwani, and J. Naor, “The greedy algorithm is optimal for on-line edge coloring,” *Inf. Process. Lett.*, vol. 44, pp. 251–253, 1992.



Randall Berry (S'93–M'00) received the B.S. degree in electrical engineering from the University of Missouri, Rolla, in 1993 and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1996 and 2000, respectively.

He is currently an Assistant Professor in the Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL. During the Summer of 2000, he was a Postdoctoral Associate in the Laboratory for Information and Decision Systems, MIT. In 1998, he was on the Technical Staff at MIT Lincoln Laboratory in the Advanced Networks Group. His primary research interests include wireless communications, data networks, and information theory. He

Dr. Berry is the recipient of a 2003 National Science Foundation (NSF) CAREER Award.



Eytan Modiano (S'90–M'93–SM'00) received the B.S. degree in electrical engineering and computer science from the University of Connecticut, Storrs, in 1986, and the M.S. and Ph.D. degrees from the University of Maryland, College Park, in 1989 and 1992, respectively, both in electrical engineering.

He was a Naval Research Laboratory Fellow from 1987 to 1992 and a National Research Council Post Doctoral Fellow during 1992–1993, while he was conducting research on security and performance issues in distributed network protocols. From 1993 to 1999, he was with the Communications Division, MIT Lincoln Laboratory, where he designed communication protocols for satellite, wireless, and optical networks and was the Project Leader for MIT Lincoln Laboratory's Next-Generation Internet (NGI) Project. He joined the MIT faculty in 1999, where he is presently an Associate Professor in the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems (LIDS). His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks.

Dr. Modiano is currently an Associate Editor for Communication Networks for the IEEE TRANSACTIONS ON INFORMATION THEORY and for the *International Journal of Satellite Communications*. He has served as a Guest Editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (Special Issue on WDM Network Architectures), the *Computer Networks Journal* (Special Issue on Broadband Internet Access), the *Journal of Communications and Networks* (Special Issue on Wireless Ad Hoc Networks), and the IEEE JOURNAL OF LIGHTWAVE TECHNOLOGY (Special Issue on Optical Networks). He is the Technical Program Co-Chair for WiOpt 2006 and Vice-Chair for INFOCOM 2007.