

Reducing Electronic Multiplexing Costs in SONET/WDM Rings with Dynamically Changing Traffic

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

Abstract—In this paper, we consider traffic grooming in WDM/SONET ring networks when the offered traffic is characterized by a set of traffic matrices. Our objective is to minimize the cost of electronic add/drop multiplexers (ADMs) in the network, while being able to support any offered traffic matrix in a rearrangeably nonblocking manner. We provide several methods for reducing the required number of ADMs for an arbitrary class of traffic matrices. We then consider the special case where the only restriction on the offered traffic is a constraint on the number of circuits a node may source at any given time. For this case, we provide a lower bound on the number of ADMs required and give conditions that a network must satisfy in order for it to support the desired set of traffic patterns. Circuit assignment and ADM placement algorithms with performance close to this lower bound are provided. These algorithms are shown to reduce the electronic costs of a network by up to 27%. Finally, we discuss extensions of this work for supporting dynamic traffic in a wide-sense or strict sense nonblocking manner as well as the benefits of using a hub node and tunable transceivers. Much of this work relies on showing that these grooming problems can often be formulated as standard combinatorial optimization problems.

Index Terms—Optical network design, SONET add/drop multiplexers (ADMs), SONET rings, topology design, traffic grooming, wavelength division multiplexing.

I. INTRODUCTION

WAVELENGTH division multiplexing (WDM) is increasingly used to expand the available capacity of an optical network. Typically these networks have a SONET/SDH ring architecture. In particular, the network nodes are arranged in a ring and interconnected by fiber (typically multiple fibers for protection purposes). Furthermore, each node in the ring uses a SONET add/drop multiplexer (ADM) to electronically combine lower rate streams onto a wavelength. For example, if a wavelength supports OC-48 (2.5 Gb/s) traffic, then 16 OC-3 (155 Mb/s) circuits can be multiplexed onto this wavelength. Using WDM technology, each fiber in a ring can support mul-

multiple wavelengths; each additional wavelength can be used to add an additional SONET ring among the nodes. This results in a substantial increase of the network capacity. However, every additional SONET ring will require additional ADMs. The dominant cost for this increased capacity is the cost of these ADMs.

Often, the number of electronic ADMs can be reduced by employing WDM add/drop multiplexers (WADMs) which allow a wavelength to be either dropped at a node or to optically bypass the node. When a wavelength is not dropped at a node, an electronic ADM is not required for that wavelength; however, the node cannot access any of the traffic on the bypassed wavelength. Thus, for a wavelength to bypass a node, the traffic for that node must be routed on the remaining wavelengths. Such a routing is referred to as a *grooming* of the traffic. Grooming with WADMs has been the topic of several recent papers including [1]–[9], [20]. Traffic grooming itself predates the work on WDM SONET rings. For example, there has been work (e.g., [10]) on grooming low rate traffic in SONET rings (without WDM) to reduce the number of required line cards. The ideas behind traffic grooming are also applicable to other WDM networks. For example, there has been recent interest in networks using IP directly on top of an optical layer. In this situation, traffic grooming can be used to reduce the number of IP ports instead of SONET ADMs [11]. Notice that traffic grooming is a special case of a virtual topological design problem. In our case the cost is measured in terms of ADMs; other work in this area considers designing virtual topologies to minimize such quantities as the number of wavelengths or the blocking probability (see, e.g., [12]–[15]).

Assume that the traffic requirement of each node in a ring is given. Consider the problem of finding a grooming which minimizes the required number of ADMs needed to support this traffic requirement; we refer to this as the *grooming problem*. In [1], [20] it is shown that the general grooming problem is NP-complete. However, for several special classes of traffic requirements, either optimal algorithms or heuristics with good performance have been found (see, e.g., [1], [6], or [8]). In each of these cases, the traffic requirement is characterized by a single static traffic matrix. Often this is not the best description of the traffic requirements, for example, if traffic changes throughout the day. A better description of the traffic may be as a set of traffic matrices. In this case, we want to minimize the number of ADMs needed to support any traffic matrix in this set. We refer to this as a *dynamic grooming problem*, since the resulting ring can support, in a nonblocking manner, traffic which dynamically changes within the given set. This problem is the main

Manuscript received October 8, 1999; revised May 15, 2000. This work was supported by the Defense Advanced Research Projects Agency under the Next Generation Internet initiative. The material in this paper was presented in part at the Optical Fiber Communication Conference, San Diego, CA, February 21–26, 1999 and at the IEEE International Conference on Communications, Vancouver, BC, Canada, June 6–10, 1999.

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

Publisher Item Identifier S 0733-8716(00)08998-8.

emphasis of this paper. A similar formulation of the dynamic grooming problem was considered in [3]–[5] for a bidirectional ring employing digital cross-connects. In addition, in [16] and [17], a similar nonblocking approach is considered but for wavelength allocation problems as opposed to traffic grooming.

We now give a precise description of the network model to be considered. Let $N = \{1, 2, \dots, N\}$ denote a set of N nodes in a WDM SONET ring. Unless otherwise noted, we consider unidirectional rings in the following. In other words, all traffic must propagate in one direction around the ring. Unidirectional path-switched rings (UPSR) are the primary example of unidirectional rings in practice. Bidirectional rings, such as BLSR/2 and BLSR/4, are also of interest in practice. Some of the following results also apply to bidirectional architectures; when this is true, we point it out below. We assume that each node has one WADM and a SONET ADM for every wavelength dropped at that node. The WADMs are static, i.e., the wavelengths that are dropped at each node are fixed. The SONET ADMs multiplex g low rate streams onto a single wavelength; g is referred to as the traffic granularity. All traffic is assumed to be duplex and consists of circuits with granularity g . A duplex circuit between two distinct nodes i and j is represented by i - j . A traffic requirement, R , between all the nodes is a multiset (a set with repeated entries) of the form $R = \{i$ - $j | i, j \in N, i \neq j\}$. A traffic requirement can also be represented by the $N \times N$ traffic matrix, $[R_{i,j}]$, where $R_{i,j}$ represents the number of circuits i - j in R . Thus, for all i and j , $R_{i,j} = R_{j,i}$ and $R_{i,i} = 0$.

We assume that nodes do not have a digital cross-connect system (DCS). We also assume nodes do not have optical wavelength changers and that both parts of a duplex connection use the same wavelength. Thus, a connection occupies a portion of the same wavelength around the entire ring. In [1], [20] it is shown that allowing wavelength changers or allowing each part of a duplex connection to use a different wavelength does not result in any improvements with regard to grooming; so, there is no loss in making this assumption and it simplifies the following analysis. On the other hand, as we will see in Section IV, using a DCS may be beneficial. It has been shown in [2] and [6] that for static traffic, a DCS can help reduce the required number of ADMs.

Consider a simple example of grooming for a single traffic matrix. Suppose we have a ring with $N = 5$ nodes and a granularity of $g = 4$, e.g., OC-12's on an OC-48. Consider the traffic requirement $R_1 = \{1$ -2, 1-2, 1-3, 1-3, 1-4, 1-4, 1-5, 1-5\}; this corresponds to each node requesting 2 duplex circuits with node 1. The minimum number of wavelengths required to support R_1 is 2. By using two wavelengths and dropping each wavelength at each node, R_1 can trivially be supported using 10 ADMs. Consider grooming the traffic so that $\{1$ -2, 1-2, 1-3, 1-3\} are placed on one wavelength and $\{1$ -4, 1-4, 1-5, 1-5\} are placed on the second wavelength. In this case the traffic can be supported using only 6 ADMs. It can be seen that this is the minimum number of ADMs needed to support R_1 . In this example, the topology which minimized the number of ADMs also minimized the number of wavelengths. As shown in [2], this is often the case, but, in general, it is not true [2], [4].

In the remainder of this paper we consider the dynamic grooming problem, i.e., grooming for a set of traffic allocations.

In Section II, we give a general formulation of this problem. We also present several approaches to reducing the required number of ADMs which are particularly applicable when the number of traffic allocations is small. In Section III, we consider a specific class of traffic requirements. This class is defined by requiring the network to support any traffic matrix such that the number of circuits terminated at each node is less than some constant t . Such a class is natural, for example, if each node is capable of sourcing only t circuits. For such traffic, we lower bound the number of ADMs needed and provide necessary and sufficient conditions that a network must satisfy to support such traffic. We use these conditions to develop algorithms for allocating ADMs in the network. Finally, in Section IV, we develop extensions to the basic model that allow a network to be nonblocking in a strict sense; the use of a hub architecture and tunable lasers in order to achieve further reductions in electronic multiplexing costs is also considered.

II. THE DYNAMIC GROOMING PROBLEM

Suppose we know that, at any time, the traffic requirement belongs to a set $\{R_1, \dots, R_K\}$ of allowable traffic requirements. The dynamic grooming problem is to find a topology with the minimum number of ADMs such that any allowable traffic requirement can be supported with an appropriate grooming. A feasible topology for the grooming problem is a topology which can support any allowable traffic matrix. The minimum number of wavelengths required by a feasible topology is $W_{\min} = \max_{i=1, \dots, K} \lceil |R_i|/g \rceil$, where $|R_i|$ is the cardinality of the set R_i , i.e., the number of circuits that need to be supported. If each of the W_{\min} wavelengths is dropped at each node, then clearly this is a feasible topology which uses NW_{\min} ADMs. Furthermore, this solution requires no grooming of the traffic to support any allowable traffic set; by this we mean that each call can be routed on any wavelength with available capacity. We seek to improve on this obvious no grooming solution.

If a ring can accommodate a particular traffic set R_1 , it can also accommodate any subset of R_1 . Thus an allowable traffic set which is the subset of another can be ignored. For any set of allowable traffic requirements, a new single traffic requirement can be defined which gives the worst-case characterization of this traffic. Specifically, given $\{R_1, \dots, R_K\}$, define $[R_{i,j}^k]$ to be the traffic matrix corresponding to the traffic requirement R_k . Consider a new traffic set R^* , which has the traffic matrix $[R_{i,j}^*]$ defined by $R_{i,j}^* = \max_{k=1, \dots, K} R_{i,j}^k$. Every allowable traffic requirement is then a subset of R^* . Thus we can apply a grooming algorithm designed for a single traffic allocation to R^* (as noted above, grooming for single traffic matrices as been addressed in several previous papers). The resulting topology is clearly feasible for the dynamic grooming problem. As shown in the next two examples, this approach may or may not yield a topology which uses fewer than NW_{\min} ADMs.

Example 1: Consider a ring with $N = 5$ nodes and granularity, $g = 4$. Suppose we have the following two allowable traffic requirements:

$$R_1 = \{1$$
-3, 1-3, 4-2, 4-2, 5-2, 5-2\}

$$R_2 = \{1-3, 1-3, 1-4, 1-4, 5-2, 5-2\}.$$




Fig. 1. Network topology which supports R^* in Example 1. In this figure, horizontal lines correspond to the two wavelengths, and an X indicates that the wavelength bypasses a node.

In this case, $W_{\min} = 2$; thus, with no grooming, both R_1 and R_2 can be supported using 10 ADMs. For this example, $R^* = \{1-3, 1-3, 1-4, 1-4, 4-2, 4-2, 5-2, 5-2\}$. It is easy to see that the topology shown in Fig. 1 can support R^* (and therefore both R_1 and R_2) using only 6 ADMs. So, in this case, grooming for R^* is beneficial.

Example 2: Consider the same ring with $N = 5$ and $g = 4$. Suppose that this time the allowable traffic sets are

$$R_1 = \{1-2, 1-2, 1-3, 1-3, 1-4, 1-4, 1-5, 1-5\}$$

$$R_2 = \{1-2, 1-2, 2-3, 2-3, 2-4, 2-4, 2-5, 2-5\}.$$

Once again, $W_{\min} = 2$ and 10 ADMs are required with no grooming. However, in this case, R^* contains 14 circuits and thus requires at least 4 wavelengths. It can be seen that 11 ADMs are needed to support R^* , which in this case is worse than the no grooming solution.

In Example 2, each of the allowable traffic matrices corresponds to uniform all-to-one traffic. For such traffic matrices, the optimal grooming can be found by an algorithm in [1]. An optimal topology for each of these traffic matrices is shown in Fig. 2. We consider how this information can be used to come up with a good topology for supporting both R_1 and R_2 . Recall that a topology T contains a topology T' if T can be obtained from T' by adding additional ADMs and possibly additional wavelengths, but not removing any. Clearly if T' can support a given traffic set, then T can also support it. Let T_1 and T_2 be the two topologies in Fig. 2 which support R_1 and R_2 , respectively. Consider forming a new topology, which also uses λ_1 and λ_2 , as follows. In the new topology, drop λ_1 at each node which has an ADM on λ_1 in either T_1 or T_2 . Likewise, drop λ_2 at each node which has an ADM on λ_2 in either T_1 or T_2 . The resulting topology contains both T_1 and T_2 and thus can support both R_1 and R_2 . For the topologies in Fig. 2, such a combination results in both wavelengths being dropped at every node, which is the same as the no grooming topology.

This is not the only way to “combine” the two topologies. Consider a second new topology, again using λ_1 and λ_2 . Now drop λ_1 at each node which either has an ADM on λ_1 in T_1 or an ADM on λ_2 in T_2 . Likewise, drop λ_2 at each node which either has an ADM on λ_2 in T_1 or an ADM on λ_1 in T_2 . This new topology is shown in Fig. 3; it contains T_1 and it also contains a topology which is equivalent to T_2 , but with its wavelengths permuted. Such an equivalent topology can clearly support the same traffic as T_2 . Therefore, the new topology can support both R_1 and R_2 . This topology only requires 7 ADMs, which is less than the 10 ADMs required by the no grooming topology.




Fig. 2. Optimal topologies for supporting R_1 (left) and R_2 (right) from Example 2.




Fig. 3. Bipartite graph for the two topologies in Fig. 2 and the resulting topology.

Each of the two preceding combinations of T_1 and T_2 was formed by assigning each wavelength in T_1 to a wavelength in T_2 , and then forming a new topology as above. Suppose that T_1 and T_2 are now two arbitrary topologies which support two corresponding traffic matrices. The above procedure can be generalized as follows. Let W be the maximum number of wavelengths used in either T_1 or T_2 . If either T_1 or T_2 used less than W wavelengths, consider it specified for W , but with no ADMs on the extra wavelengths. There are now $W!$ ways of matching wavelengths in the two topologies. The best combination is one which requires the fewest ADMs. For large values of W , considering each of the $W!$ combinations becomes unattractive, but as shown next such a “brute force” approach is not necessary. Every possible combination can be represented by a bipartite graph (C, D, E) . Recall, a bipartite graph (C, D, E) is a graph with two disjoint sets of nodes, C and D , and a set of edges, E , where each edge is between a node in C and a node in D . Here, C and D will correspond to the sets of wavelengths in T_1 and T_2 , respectively. Between each $c \in C$ and $d \in D$, there is an edge $(c, d) \in E$. We associate with each edge, $(c, d) \in E$, a cost which equals the number of distinct nodes with an ADM on either wavelength c in T_1 or wavelength d in T_2 . The bipartite graph for the two topologies in Fig. 2 is shown in Fig. 3. A matching of all the wavelengths in T_1 to the wavelengths in T_2 correspond to a set of W disjoint edges in this graph; the total cost of these edges gives the number of ADMs required for this combination. Thus we want to find such a matching which has the minimum total cost. This problem can be recognized as a instance of an “assignment problem;” this combinatorial optimization problem has several well-known polynomial algorithms [18]. Such algorithms can be used to find the best combination of wavelengths in polynomial time.

We wish to make several comments about this approach. First, there are often multiple optimal topologies which are not obtainable from each other by simply permuting the wavelengths. The number of ADMs in the optimal combination may depend on which of these topologies is used. Second, we have considered only a case with two allowable traffic matrices.¹

For a larger number, one can generalize this assignment

¹Note this is not as restrictive as it first appears, since we can also support any subsets of the two allowable traffic sets.

problem. Unfortunately, the generalization results in a much more difficult combinatorial optimization problem.² Of course, one could consider combining topologies sequentially, i.e., first combining two topologies, then adding a third to this combination, etc. For more than a few topologies, this approach is unattractive. Finally we note that nothing in this formulation required a unidirectional ring. Thus, this approach also works in the bidirectional case.

III. t -ALLOWABLE TRAFFIC

In this section we focus on a particular subclass of dynamic grooming problems. Specifically, we consider sets of allowable traffic which are determined only by limiting each node to sourcing at most t duplex circuits at any time, where t is a specified constant. In other words, a traffic matrix $[R_{i,j}]$ is allowable if and only if it satisfies

$$\sum_j R_{ij} \leq t \quad \text{for all } i. \quad (1)$$

This will be the only constraint on the set of traffic patterns which must be supported. We call these allowable traffic matrices t -allowable. Such a constraint is natural in many cases. For example, suppose each node represents a customer site and that a service provider wishes to guarantee the customer the availability of t “switched OC-3 connections” at each site. To satisfy this guarantee, the ring must be provisioned to handle any t -allowable traffic pattern.

A similar class of allowable traffic matrices was studied in [3]–[5]. The model in [3] allows each node i to source at most $t(i)$ circuits, where $t(i)$ is constrained to be a multiple of g (note in [3] the units of $t(i)$ are the number of light-paths as opposed to the number of low-rate circuits as we have defined them here). The work in [3]–[5] focuses on bidirectional rings where every node has a DCS. In contrast, we focus on the unidirectional ring case without using DCSs. Hence, the approaches in [3]–[5] do not directly apply to the problem considered in this paper. Furthermore, in [3] and [4] the grooming problem is divided into two steps. First low speed traffic is groomed into light-paths, and then these light-paths are grouped onto SONET rings. While this simplifies the problem, it may lead to a suboptimal solution [6]. Our approach considers the two problems together.

In the following, assume that we are given a network with N nodes and a traffic granularity of g . We refer to a traffic set as t -maximal if it is t -allowable and if the addition of any other circuit would make it not t -allowable. As noted above, if a network can support every t -maximal traffic matrix it can support every t -allowable one. For any t -allowable traffic set, R , the maximal number of circuits in this set is bounded by

$$|R| \leq \lfloor Nt/2 \rfloor. \quad (2)$$

Furthermore, there exist t -maximal sets which achieve this bound. Therefore, the minimum number of wavelengths, W_{\min} ,

for the set of t -allowable traffic patterns is given by

$$W_{\min} = \left\lceil \frac{\lfloor Nt/2 \rfloor}{g} \right\rceil. \quad (3)$$

Thus, NW_{\min} ADMs are required for the no grooming solution; this gives an upper bound on the required number of ADMs. We will focus on reducing this number of ADMs while still supporting any t -allowable traffic matrix using W_{\min} wavelengths. As noted above, the minimum ADM solution often uses the minimum number of wavelengths, but not always. Hence, restricting our solutions to those using the minimum number of wavelengths is sensible not only because it makes efficient use of wavelengths, but also because it is likely to yield a nearly optimal solution. The problem we address can be stated as follows. For given values of N , g , and t , we wish to specify a topology, i.e., which of the N nodes have ADMs on which of the W_{\min} wavelengths. This topology must be able to support any t -allowable traffic matrix using the minimum number of ADMs. In the previous section, we considered finding topologies for each allowable traffic set and then combining these topologies in order to support every allowable traffic set. When $t = 1$ there are $2^{-N/2} N! / (N/2)!$ different t -maximal traffic matrices. Clearly, finding a topology for each possible t -allowable traffic matrix and then combining these is not a feasible approach. In the following we develop an alternative approach which relies on formulating this problem as a bipartite matching problem. Before looking at this approach, we give a lower bound on the required number of ADMs.

A. Lower Bound on the Number of ADMs

Minimizing the required number of ADMs is equivalent to starting out with every node having an ADM on each wavelength and maximizing the number of ADMs that can be removed while still supporting every t -allowable traffic matrix. In this section we give an upper bound on the number of ADMs that can be removed or equivalently a lower bound on the number of ADMs needed by a feasible topology. First we establish some preliminary results. Throughout this section, assume that N , t , and g are specified.

We look at a particular way to construct a t -allowable traffic set which attains the bound in (2). This construction is useful in proving Lemma 1 below. Denote a permutation of the set of nodes, N , by π , i.e., $\pi(1), \dots, \pi(N)$ is some ordering of the nodes. Let $\pi^{-1}(i)$ denote the i th node in this ordering. Define two sets of pairs of nodes, C_1 and C_2 . The set C_1 contains all pairs of nodes $(\pi^{-1}(i), \pi^{-1}(i+1))$ where i is odd and strictly less than N . The set C_2 contains all pairs of nodes $(\pi^{-1}(i), \pi^{-1}(i+1))$ where i is even and strictly less than N along with the pair $(\pi^{-1}(N), \pi^{-1}(1))$. For a given t , the desired t -maximal set contains $\lceil t/2 \rceil$ circuits between each pair in C_1 and $\lfloor t/2 \rfloor$ circuits between each pair in C_2 . The resulting set is obviously t -maximal and contains

$$\lceil t/2 \rceil \lfloor N/2 \rfloor + \lfloor t/2 \rfloor \lceil N/2 \rceil = \lfloor Nt/2 \rfloor \text{ circuits.}$$

For example, suppose $N = 5$. With the trivial permutation, $\pi(i) = i$, we have $C_1 = \{(1, 2), (3, 4)\}$ and $C_2 = \{(2, 3), (4, 5), (5, 1)\}$. If $t = 3$ the t -maximal set given by the above

²The assignment problem is easy to solve because it is a unimodular problem. Such problems have the characteristic that vertex solutions to the LP relaxation are the same as the solution to the integer program. When we generalize to more than two traffic matrices, the resulting problem is no longer unimodular.

construction is $\{1-2, 1-2, 3-4, 3-4, 2-3, 4-5, 5-1\}$. This set contains $\lfloor (5)(3)/2 \rfloor = 7$ circuits as we desired. If instead we use the permutation $\pi(i) = (i + 1) \bmod N$, then $C_1 = \{(5, 1), (2, 3)\}$ and $C_2 = \{(1, 2), (3, 4), (4, 5)\}$. In this case, we get the 3-maximal set $\{5-1, 5-1, 2-3, 2-3, 1-2, 3-4, 4-5\}$.

Define M_i to be the set of nodes with ADMs removed from wavelength i for $i = 1 \cdots W_{\min}$. The following lemmas help to bound the maximum number of ADMs which can be removed.

Lemma 1: If a network with W_{\min} wavelengths can support every t -allowable traffic set, then for all $i = 1 \cdots W_{\min}$, $|M_i| < N/2$.

Proof: To establish a contradiction assume that a network can support every t -allowable set, but for some i , $|M_i| \geq N/2$. In this case, we construct a t -maximal set where every connection involves a node in M_i . Thus, no circuit in this set can be supported on wavelength i . Furthermore, this set will require at least W_{\min} wavelengths, and so cannot be supported. The particular t -maximal set we construct has the form discussed above. Let π be a permutation of N such that for every odd $i \leq N$, $\pi^{-1}(i)$ is in M_i . Such a permutation exists since $|M_i| \geq N/2$. Using this permutation, consider the t -maximal set defined above. This set requires W_{\min} wavelengths and each circuit involves a node from M_i as desired. ■

Lemma 2: If a network with W_{\min} wavelengths can support every t -allowable traffic set, then for all $i = 1 \cdots W_{\min}$, $|M_i| < (W_{\min} - 1)g/t$.

Proof: Again, we prove the lemma by contradiction. Assume that a network can support every t -allowable set, but for some i , $|M_i| \geq (W_{\min} - 1)g/t$. By Lemma 1, $|M_i| < N/2$. Thus, we can pair up each node in M_i with a distinct node in $N - M_i$ and form a t -allowable traffic set by setting up t circuits between each pair. This traffic set consists of $|M_i|t$ circuits, none of which can be placed on wavelength i . This set must be placed on the remaining $W_{\min} - 1$ wavelengths, but these wavelengths can accommodate at most $(W_{\min} - 1)g$ circuits. Thus this t -allowable set cannot be supported, yielding a contradiction. ■

Lemma 3: If a network with W_{\min} wavelengths can support every t -allowable traffic set, then for all $i \neq j$,

$$\min(|M_i|, |M_j|) \leq (W_{\min} - 2)g/t. \quad (4)$$

Proof: We show that we can always construct a t -allowable set with $\min(|M_i|, |M_j|)t$ circuits which cannot be carried on either wavelength i or j and thus must be carried on the other $W_{\min} - 2$ wavelengths. Since each wavelength can accommodate at most g circuits, (4) must be true for this set to be supported. The proof will be completed once we show how to construct the above set. Consider two wavelengths i and j and assume $|M_i| \leq |M_j|$. Let K be the set of nodes removed from both i and j (K may be empty). From Lemma 1, we can assume that $|M_j| < N/2$. Thus, each node in K can be paired with a distinct node in $N - M_j$. Likewise, every node in $M_i - K$ can be paired with a distinct node in $M_j - K$. Placing t circuits between each pair gives the required t -allowable set. ■

An immediate corollary of Lemma 3 is that for every wavelength except one, we must have $|M_i| \leq (W_{\min} - 2)g/t$. Lemma




Fig. 4. Provisioning of ring in Example 3.

2 gives a bound on the ADMs that can be removed on the remaining wavelength. Thus, we have the following upper bound:

$$\begin{aligned} \text{ADMs removed} \leq & (W_{\min} - 1) \left\lfloor \frac{g}{t} (W_{\min} - 2) \right\rfloor \\ & + \left\lfloor \frac{g}{t} (W_{\min} - 1) \right\rfloor. \end{aligned} \quad (5)$$

The next example shows that this bound is tight for some choices of N , t , and g .

Example 3: Suppose we have a network with $N = 5$, $t = 2$, and $g = 2$. For this ring, $W_{\min} = 3$ and the above upper bound yields that at most 4 ADMs can be removed. Fig. 4 shows a topology that achieves this bound. Consider the 2-allowable traffic requirement $\{1-2, 1-3, 2-3, 4-5, 4-5\}$. This can be supported on a ring provisioned as in Fig. 4 by assigning $\{1-2, 2-3\}$ to the first wavelength, $\{4-5, 4-5\}$ to the second wavelength, and 1-3 to the third wavelength. Such an assignment can be found for any other 2-allowable traffic set.

Next we establish a connection between this problem and bipartite matching problems. By exploiting this connection, we come up with necessary and sufficient conditions for a topology to be able to support any t -allowable traffic matrix. This is then used to develop several heuristic algorithms for removing ADMs from wavelengths.

B. Bipartite Matching Formulation

For a given ring network, we want to construct a bipartite graph (C, D, E) which represents the possible placements for each call from a given t -allowable traffic set, R . We will denote one set of nodes in the graph by $D = \{\lambda_{1,1}, \dots, \lambda_{1,g}, \lambda_{2,1}, \dots, \lambda_{W_{\min},g}\}$. This set contains g elements for each of the W_{\min} wavelengths corresponding to possible circuit assignments on that wavelength. The other set of nodes, C , will correspond to the traffic requirement R . There is an edge in the graph between $\lambda_{i,j}$ and a circuit $k-l \in R$ if both nodes k and l have an ADM on wavelength i . For example, the bipartite graph corresponding to the topology and the traffic requirement from Example 3 is shown in Fig. 5.

A matching M in a bipartite graph is a set of disjoint edges. Being able to accommodate a traffic matrix in a given topology is equivalent to being able to find a matching in the corresponding bipartite graph which uses all the nodes in the set of requested circuits C . Such a matching is called a C -saturating matching. A necessary and sufficient condition for the existence of such a matching is given by Hall's theorem which we state below. First we need the following definition. For a bipartite graph (C, D, E) , if S is a subset of nodes in C , then the open neighborhood of S , $N(S)$, is a subset of nodes in D such that d is in $N(S)$ if and only if there is an edge between d and a node in S , i.e., $d \in N(S) \Leftrightarrow (c, d) \in E$ and $c \in S$.




Fig. 5. Bipartite graph corresponding to traffic set and topology in Example 3.

Hall's Theorem: Let $G = (C, D, E)$ be a bipartite graph. There exists a C -saturating matching if and only if for all subsets S of C , $|N(S)| \geq |S|$.

A proof of this theorem can be found in many texts on combinatorics (e.g., [19]). As stated, this theorem is useful to check that a single traffic matrix can be supported. We are interested in supporting every t -allowable traffic matrix; the following theorem provides a necessary and sufficient condition for this. We say that a circuit i - j must be routed on a set of wavelengths if either i or j does not have an ADM on any wavelength not in this set.

Theorem 1: For a given topology with W_{\min} wavelengths, any t -allowable traffic matrix can be supported if and only if the following two conditions are satisfied.

- A. For every pair of nodes (i, j) , there exists a wavelength on which both i and j have an ADM.
- B. For any group of m wavelengths and any t -allowable set, there exist at most gm circuits which must be routed on this group.

Proof: We first show that these conditions are necessary. Clearly, if **A** is not satisfied then any t -allowable set containing i - j cannot be accommodated. If **B** is not satisfied then there exists a set of m wavelengths on which we must route more than gm circuits in some t -allowable set, C . Consider the bipartite graph corresponding to C . Let S be the subset of C containing the above circuits, then $|N(S)| = gm$ and $|S| > gm$. Thus by Hall's theorem there exists no C -saturated matching, and this traffic matrix cannot be accommodated.

Next we show that these conditions are sufficient. Assume that they are not sufficient, so that there exists an assignment of ADMs to W_{\min} wavelengths which satisfies both of the above conditions, but which cannot support some t -allowable traffic set C . Since C cannot be supported, by Hall's Theorem there exists a subset S of C such that $|N(S)| < |S|$. Let k be a nonnegative integer such that $(k-1)g < |S| \leq kg$. For a bipartite graph corresponding to an allocation of ADMs, $|N(S)|$ will always be a multiple of g . Thus $|N(S)| < |S|$ implies that $|N(S)| \leq (k-1)g$. Therefore, this set of more than $(k-1)g$ calls must be routed on a set of $k-1$ or fewer wavelengths, which contradicts condition **B**, completing the proof. ■

C. Algorithms for Removing ADMs

We now use the results from Theorem 1 to develop algorithms for removing ADMs from wavelengths. The resulting topologies will support any t -allowable traffic requirement. Given such a topology, one then needs to know how to groom the traffic for each allowable traffic set. For each traffic matrix one can set up a maximum matching problem as in

Section III-B. Polynomial algorithms for solving this problem are known (see, e.g., [18]). In many cases an assignment can be found by inspection. These assignments can be all computed off-line and stored in a look-up table. Also, in some cases, the assignments can be stored in a more compact form than simply listing every possible assignment. Alternatively, if the traffic changes slowly, the assignment for the current traffic set can be computed on-line.

If $W_{\min} = 1$, no ADMs can be removed in any feasible topology. If $W_{\min} = 2$, every node must have an ADM on one wavelength, and at most $\lfloor g/t \rfloor$ nodes can be removed from the other wavelength. This follows directly from Lemmas 2 and 3. Furthermore, if $\lfloor g/t \rfloor$ nodes are removed from the other wavelength, the resulting topology is feasible. To see this, note that the most circuits that will be forced onto one wavelength is $\lfloor g/t \rfloor t \leq g$. So, by the Theorem 1 we can accommodate all t -allowable circuits. Thus for $W_{\min} \leq 2$ we have a trivial algorithm which yields the minimum number of ADMs. Therefore, in the following, we shall only consider the case where $W_{\min} \geq 3$.

To use Theorem 1 to verify that a topology can support every t -allowable traffic pattern, condition **B** must be checked for every subset of wavelengths. There are $2^{W_{\min}}$ possible subsets; checking each set is not an appealing prospect. In the following we avoid this by removing ADMs in certain symmetric patterns which require us to check many fewer cases.

For a circuit i - j to be forced on a set of n wavelengths, either i or j must have an ADM removed from each of the remaining $W_{\min} - n$ wavelengths. When $W_{\min} - n > 2$, this can only occur if at least one of the two nodes has an ADM removed from more than one wavelength. So if we remove at most one ADM for each node, we only have to check **B** for sets of $W_{\min} - 1$ and $W_{\min} - 2$ wavelengths. Clearly, we can remove $\lfloor N/W_{\min} \rfloor$ nodes from each wavelength so that no node will be removed from more than one wavelength. Also if we remove $\lfloor (W_{\min} - 2)g/t \rfloor$ or fewer nodes per wavelength, then no more than $(W_{\min} - 2)g$ circuits will be forced on any set of $W_{\min} - 2$ or $W_{\min} - 1$ wavelengths. Thus if we remove $\min(\lfloor (W_{\min} - 2)g/t \rfloor, \lfloor N/W_{\min} \rfloor)$ ADMs per wavelength and no more than one ADM per node, condition **B** is satisfied. Condition **A** is also easily satisfied in this case. Thus we have proved the following lemma which immediately yields an algorithm for allocating ADMs.

Lemma 4: For $W_{\min} > 2$, one can always remove $\min(\lfloor (W_{\min} - 2)g/t \rfloor, \lfloor N/W_{\min} \rfloor)$ ADMs from each of W_{\min} wavelengths such that no node has more than one ADM removed and any t -allowable traffic matrix can be supported.

Recall that according to Lemma 3 we can remove more than $\lfloor (W_{\min} - 2)g/t \rfloor$ nodes from at most one wavelength. Thus if

$$\lfloor (W_{\min} - 2)g/t \rfloor < \lfloor N/W_{\min} \rfloor \quad (6)$$

the above algorithm removes the most nodes possible from every wavelength except possibly one. When W_{\min} becomes large for a given N , the inequality in (6) is reversed. When this occurs, the procedure in Lemma 6 will remove only a small percentage of the ADMs. In such cases, to get further reductions in ADMs, we have to consider removing nodes from more than one wavelength. In the following we first consider the case

where a node can be removed from at most 2 wavelengths; then we generalize to an arbitrary number of wavelengths.

Suppose we allow a node to be removed from at most two wavelengths. Assume (6) does not hold, then we can remove at least $\lfloor N/W_{\min} \rfloor$ ADMs from each wavelength. Consider removing ADMs in the following manner: from wavelength i , remove ADMs for nodes $(i - 1)\lfloor N/W_{\min} \rfloor$ to $(i\lfloor N/W_{\min} \rfloor + k) \bmod N$, where k is a constant to be determined. For a node to be removed from at most 2 wavelengths, we must have $k \leq \lfloor N/W_{\min} \rfloor$. If a node is removed from wavelength i , the only other wavelengths it can be removed from are $(i-1) \bmod N$ and $(i + 1) \bmod N$. Since a node is removed from no more than 2 wavelengths, traffic can only be forced onto groups of $W_{\min} - 4$ or more wavelengths. Thus condition **B** need only be checked for sets of $W_{\min} - 4$, $W_{\min} - 3$, $W_{\min} - 2$, or $W_{\min} - 1$ wavelengths. For a given choice of k , the number of circuits forced on a set of $W_{\min} - 4$ wavelengths is at most kt . The number forced on a set of $W_{\min} - 3$ is at most $2kt$, and the number forced on sets of $W_{\min} - 1$ and $W_{\min} - 2$ is at most $(\lfloor N/W_{\min} \rfloor + k)t$. Thus for condition **B** to hold, the following inequalities must be satisfied:

$$(W_{\min} - 4)g \geq kt \tag{7}$$

$$(W_{\min} - 3)g \geq 2kt \tag{8}$$

$$(W_{\min} - 2)g \geq (\lfloor N/W_{\min} \rfloor + k)t \tag{9}$$

$$(W_{\min} - 1)g \geq (\lfloor N/W_{\min} \rfloor + k)t. \tag{10}$$

For (7) to be satisfied for a positive value of k , it must be that $W_{\min} > 4$. Also note that for condition **A** to fail, there must be a circuit which is blocked from every wavelength. When ADMs are removed in the above manner, a circuit can be blocked from at most 4 wavelengths. Thus, when $W_{\min} > 4$, **A** is always satisfied. Therefore, when $W_{\min} > 4$, the largest $k < \lfloor N/W_{\min} \rfloor$ which satisfies (7)–(10) yields the most ADMs which can be removed in this manner.

Assuming that $W_{\min} > 4$, it is sufficient to check only (8) and (9) out the four inequalities above; this is shown next. First note that if (9) is satisfied, then clearly (10) must also be. Let $l = W_{\min} - 4$, then by assumption $l \geq 1$. Inequality (8) can be rewritten as $(l + 1)g/2 \geq kt$ and (7) can be written as $lg \geq kt$. Note that for $l \geq 1$, $(l + 1)/2 \leq l$. Thus if (8) is satisfied, then (7) must also be. If either N or t is even, it can be shown that if (8) is satisfied, then (9) must also be; thus in this case we only need to check that (8) is satisfied. In this case it also follows that (6) is never satisfied. These results are summarized in the following lemma.

Lemma 5: If $W_{\min} > 4$ and (6) is not satisfied, then $\lfloor N/W_{\min} \rfloor + k$ ADMs can be removed per wavelength in the above manner for any $k < \lfloor N/W_{\min} \rfloor$ which satisfies (8) and (9). If $W_{\min} > 4$ and either N or t is even, then $\lfloor N/W_{\min} \rfloor + k$ ADMs can be removed per wavelength in the above manner for any $k < \lfloor N/W_{\min} \rfloor$ which satisfies (8).

Next we generalize the above procedure to allow nodes to be removed from an arbitrary number of wavelengths. For now we assume that $N \geq W_{\min}$. For given integers x and k , suppose we remove nodes $(i-1)\lfloor N/W_{\min} \rfloor + 1$ to $((i+x-2)\lfloor N/W_{\min} \rfloor + k) \bmod N$ from wavelength i , where $0 \leq k \leq \lfloor N/W_{\min} \rfloor$. Thus we remove $(x-1)\lfloor N/W_{\min} \rfloor + k$ nodes from each wave-

length, and a node is removed from at most x wavelengths. Traffic is then only forced onto groups of $W_{\min} - 2x$ or more wavelengths. For an arbitrary value of x , as long as x is less than $W_{\min}/2$, condition **A** is satisfied. We only need to check condition **B** for sets of $W_{\min} - 2x$ or more wavelengths. The most circuits that can be forced on a set of $W_{\min} - 2x$ wavelengths is kt . To see this, consider $2x$ adjacent wavelengths and note that there are k nodes without an ADM on any of the first x wavelengths and k other nodes without an ADM on the next x wavelengths. By similar reasoning, we can find the most circuits that can be forced on sets of $W_{\min} - (2x - 1)$ to $W_{\min} - 1$ wavelengths. In this manner, we get the following set of inequalities which must be satisfied for **B** to hold.

$$\begin{aligned} (W_{\min} - 2x)g &\geq kt \\ (W_{\min} - (x+i))g &\geq ((x-1-i)\lfloor N/W_{\min} \rfloor + 2k)t, \quad \forall i = 1, \dots, x-1 \\ (W_{\min} - i)g &\geq ((x-1)\lfloor N/W_{\min} \rfloor + k)t, \quad \forall i = 1 \dots, x. \end{aligned}$$

When $x = 2$, these inequalities are the same as (7) and (8). In the $x = 2$ case, we were able to reduce this set of inequalities to a smaller subset. A similar reduction can be shown for an arbitrary choice of x . Specifically, out of this set of $2x$ inequalities, it can be shown via algebraic manipulations that if the following three inequalities are satisfied, then the entire set of $2x$ must also be.

$$\begin{aligned} (W_{\min} - (2x - 1))g &\geq 2kt \\ (W_{\min} - (x + 1))g &\geq ((x - 2)\lfloor N/W_{\min} \rfloor + 2k)t \\ (W_{\min} - x)g &\geq ((x - 1)\lfloor N/W_{\min} \rfloor + k)t. \end{aligned}$$

From this it follows that the most ADMs that be removed in this manner is given by the solution to the following integer program:

$$\begin{aligned} \text{maximize} \quad & (x-1)\lfloor N/W_{\min} \rfloor + k \\ \text{subject to:} \quad & (W_{\min} - (2x-1))g \geq 2kt \\ & (W_{\min} - (x+1))g \geq ((x-2)\lfloor N/W_{\min} \rfloor + 2k)t \\ & (W_{\min} - x)g \geq ((x-1)\lfloor N/W_{\min} \rfloor + k)t \\ & 0 \leq k \leq \lfloor N/W_{\min} \rfloor \\ & 1 \leq x \leq \lfloor W_{\min}/2 \rfloor \end{aligned} \tag{P}$$

where x and k are constrained to be integers. This optimization problem can be solved in the following manner. First, set $k = 0$ and find the largest value of x which satisfies the constraints. Next, fix x at this value and find the largest value of k satisfying the constraints. Again we summarize these results in the following lemma which immediately yields an algorithm for removing ADMs.

Lemma 6: Consider a ring with $W_{\min} > 2$. Then we can remove $(x-1)\lfloor N/W_{\min} \rfloor + k$ ADMs per wavelength in the above manner where x and k are solutions to the integer program (P) and still support every t -allowable traffic matrix.

Example 4: The following provides an example of the algorithms in Lemmas 4 and 6. Consider a ring with 15 nodes, $g = 16$, and $t = 10$. For this ring, $W_{\min} = 5$ and $\lfloor N/W_{\min} \rfloor = 3$.




Fig. 6. Topology corresponding to Lemma 4.




Fig. 7. Topology corresponding to Lemma 6.


Using the algorithm from Lemma 4 we can remove 15 ADMs. The resulting allocation is shown in Fig. 6. Using the algorithm in Lemma 6, one finds that $x = 2$ and $k = 1$, and thus one can remove 4 nodes per wavelength, for a total of 20 ADMs removed. The resulting topology is shown in Fig. 7. For comparison, the upper bound on the number of ADMs removed from (5) is 22.

When W_{\min} is larger than N , then $\lfloor N/W_{\min} \rfloor = 0$ and the above algorithms as stated will not remove any ADMs. We describe a way that these algorithms can be modified to be useful in this case. First note that for positive integers s and r , a traffic set is $(s+r)$ -allowable if and only if it can be written as the union of an s -allowable set and an r -allowable set. When $W_{\min} > N$, we can use this to decompose the allowable traffic into smaller sets such that each set will fit on N or fewer wavelengths. Suppose we want to support all t -allowable traffic, and this requires more than N wavelengths. Let $k = \lfloor t/2g \rfloor$ and let $t' = t - 2kg$. Decompose each t -allowable traffic set into k $2g$ -allowable sets and one t' -allowable set. Each $2g$ -allowable set can be accommodated on N wavelengths and the remaining set requires $\lceil \lfloor Nt'/2 \rfloor \lceil 1/g \rceil$ wavelengths. Note that

$$kN + \left\lceil \lfloor Nt'/2 \rfloor \frac{1}{g} \right\rceil = \left\lceil \lfloor Nt/2 \rfloor \frac{1}{g} \right\rceil \quad (11)$$

i.e., decomposing traffic in this way requires no more wavelengths. Since the number of wavelengths needed for each set in this decomposition is less than or equal to N , we can apply the above algorithms to remove ADMs from each set. The resulting topology will support all t -allowable traffic. This is illustrated next.

Example 5: Consider a ring with $N = 5$, $g = 2$, and $t = 6$ so that $W_{\min} = 8$. Applying the above procedure, we get one set of 5 wavelengths which must support 4-allowable traffic and one set of 3 wavelengths which must support 2-allowable traffic. Applying Lemma 4 to both of these sets, we find we can remove

Fig. 8. Number of ADMs in topology generated by Lemma 6 for $N = 15$, $g = 16$.

1 ADM from each wavelength and thus eliminate a total of 8 ADMs.

As in Example 4, consider a ring with $N = 15$ and $g = 16$. In Fig. 8 we have plotted the number of ADMs resulting from the algorithm in Lemma 6 as t ranges from 1 to 30. The number of ADMs with no grooming is also plotted along with the lower bound from (5). With the grooming, the number of ADMs is reduced by up to 27%. In [1], [20] it was found that approximately 60 ADMs were needed to support uniform all-to-all traffic in this network. Uniform all-to-all traffic is $(N - 1)$ -maximal. In this example, supporting all $(N - 1)$ -allowable traffic requires 77 ADMs, but this is a much less restrictive set of traffic.

We note that for any $W' \geq W_{\min}$, the algorithms in this section still work; in this case they generate a topology using W' wavelengths which supports all t -allowable traffic. As noted above, at times using more than W_{\min} wavelengths can reduce the required number of ADMs.

D. Hot Spot Node

In this subsection we consider a generalization the set of t -allowable traffic. Suppose there is one node in the network which has no restriction on the number of circuits it can source, while all other nodes are still restricted to t circuits. We refer to the unconstrained node as a *hot spot*. For example, this node can be used to model a central office node. In this case, the minimum number of wavelengths required to support all allowable traffic is $W_{\min} = \lceil (N - 1)t/g \rceil$. Clearly the hot spot node needs an ADM on each wavelength. Consider applying the grooming algorithms to the set of nodes not including the hot spot with the above number of wavelengths. The resulting topology will handle all t -allowable traffic between these nodes. This allocation is also sufficient for all allowable traffic including the hot spot node. To see this, note that by including the hot spot node, no additional calls are forced onto any group of wavelengths. Thus, by Theorem 1, any allowable traffic matrix can be supported. This procedure applies with an arbitrary number of hot spots.

IV. EXTENSIONS TO THE BASIC MODEL

In this section we will describe a number of extensions to our basic model. First, we discuss the use of a strict sense or wide sense nonblocking network to support rapidly changing traffic. Then we discuss the benefits of using a hub node and tunable lasers.

A. Blocking Properties

In the previous sections we found topologies which can support any allowable traffic set. In this section we want to study the properties of the rings as traffic changes from one allowable set to the other. Specifically, suppose traffic changes from one allowable traffic set to another, while some subset of the circuits stay active. We look at whether we can support the new traffic set without rerouting the existing calls. When discussing such properties, we will use some standard definitions from switching theory which we repeat here. A ring is *strict sense* or *strictly nonblocking* if any allowable circuit between nodes can be established without interference from any other existing allowable circuits. A ring is *wide-sense nonblocking* if any allowable circuit between nodes can be established without interference from any other existing allowable circuits, provided that the existing circuits have been established according to some algorithm. A ring is *rearrangeably nonblocking* if any allowable circuit can be established by possibly rerouting any existing circuits. Clearly,

$$\text{Strict sense} \Rightarrow \text{Wide sense} \Rightarrow \text{Rearrangeable.}$$

The converse implications do not, in general, hold.

A ring provisioned according to the algorithms in Sections II and III is rearrangeably nonblocking but not necessarily strictly or wide-sense nonblocking. If traffic changes frequently, then the control overhead associated with rearranging existing circuits may not be acceptable. In such a case, one may prefer a ring that is either wide-sense or strictly nonblocking. If every node has an ADM on each of the W_{\min} wavelengths, then the resulting ring is strictly nonblocking. Similarly, when a ring is provisioned to support the traffic matrix R^* defined in Section II, it will also be strictly nonblocking. For any of the other cases looked at, the resulting ring will not necessarily be strictly nonblocking. For the case of t -allowable traffic, any ring with W_{\min} wavelengths must have NW_{\min} ADMs in order to be strictly nonblocking. In other words, in this case one cannot save on the cost of ADMs by grooming. We prove this for the case of $t = 1$, but it can be modified for an arbitrary t .

Theorem 2: For t -allowable traffic, a strictly nonblocking ring with $t = 1$ and W_{\min} wavelengths must have an ADM for each node on each wavelength.

Proof: When $W_{\min} = 1$, the theorem is clearly true. For $W_{\min} = 2$, we know that all the nodes must be on one of the wavelengths, if we remove only one node, say node j , from wavelength 1. We can find a set of g circuits not involving node j and place them on wavelength 2. Then any additional circuit involving node j cannot be established without rearranging these existing circuits, and so the ring is not strictly nonblocking.

For $W_{\min} > 2$, we proceed by induction. First note, from Lemma 1 there must be at least $(N/2) + 1$ nodes on each wavelength for the ring to be even rearrangeably nonblocking. When $W_{\min} > 2$ and $t = 1$, it follows from the definition of W_{\min} that $N/2 \geq 2g + 1$. Thus there must be more than $2g + 2$ nodes on each wavelength. Now assume that the theorem is true for $W_{\min} = k$ wavelengths, and consider the case when $W_{\min} = k + 1$. Without loss of generality, we can assume that nodes $1, \dots, 2g + 2$ are on wavelength 1. Thus we can consider any 1-allowable set of g circuits between $2g$ of these nodes and

place these circuits on wavelength 1. Then any other 1-allowable set of calls between the remaining $N - 2g$ nodes must be placed on the remaining k wavelengths. If we consider a ring with these $N - 2g$ nodes, then the minimum number of wavelengths for this ring is k . Therefore, by the induction hypothesis, we cannot remove any of these $N - 2g$ nodes from the remaining k wavelengths. The original $2g$ nodes were picked arbitrarily from the set of $2g + 2$ nodes that must be on wavelength 1, and by choosing different sets and repeating this argument we have that every node must be on the remaining k wavelengths. Likewise, by repeating this argument, but starting with a different initial wavelength, we see that every node must be on every wavelength. Thus the theorem is true for $W_{\min} = k + 1$, and, by induction, for any ring with $t = 1$. ■

Next we consider wide-sense nonblocking rings. This case is more difficult than the other cases due to the fact that a routing algorithm must also be considered. The following gives an upper bound on the ADMs that can be removed for a wide-sense nonblocking ring with t -allowable traffic.

Lemma 7: Consider a unidirectional ring with W_{\min} wavelengths. Let M_i be the set of nodes removed from wavelength i . For the ring to be wide-sense nonblocking for t -allowable traffic, where t is even, we must have for all i

$$|M_i| \leq \max(2W_{\min}g/t - N, 1).$$

Proof: First note that $|M_i| \leq N - 2$ for all i . We show that if $|M_i| > 1$, then it must be that $|M_i| \leq 2W_{\min}g/t - N$. The lemma then follows. If $2 \leq |M_i| \leq N - 2$, then we can form the following t -maximal set which also has the maximal link load. This set consists of two groups of traffic. One group consists of $|M_i|t/2$ of circuits which are only between nodes in M_i . The other group consists of circuits only between nodes in $N - M_i$. Let X be the subset of the circuits in $N - M_i$ which are routed on wavelength i (X cannot be empty since it is a t -maximal set and we are using W_{\min} wavelengths).

First we prove if the ring is wide-sense nonblocking, then

$$|X| \geq |M_i|/2. \quad (12)$$

Assume this is not true. Suppose the circuits in X were disconnected as well as $|X|$ of the circuits involving the nodes in M_i . We can find a set of $2|X|$ new circuits where each circuit involves only one node in M_i and one node which previously was in a circuit in X . Adding this set of circuits to the remaining calls results in a new t -maximal set, and none of these new calls can be routed on wavelength i . This new set will also have the maximum link load and thus requires all W_{\min} wavelengths. Thus, these calls cannot be accepted without rearranging some of the other active calls. This is a contradiction and so (10) must be true.

If (12) is true, beginning with a t -maximal set as above, assume that the circuits involving the nodes in M_i are disconnected along with $|M_i|t/2$ circuits involving nodes from X . Then we can form $|M_i|t$ circuits, as above, where each circuit is between one node from M_i and one node that was in a circuit in X . These additional circuits must be routed on the remaining $W_{\min} - 1$ wavelengths without rearranging the active calls. This means that at most $(W_{\min} - 1)g - |M_i|t$ calls not involving

the nodes in M_i can be routed on these wavelengths. There are $(N - |M_i|)t/2$ circuits in the original maximal set not involving nodes in M_i , thus we must have

$$|X| \geq (N - |M_i|)t/2 - (W_{\min} - 1)g + |M_i|t.$$

Also $|X| \leq g$. Combining these and performing some algebra yields the desired result. ■

We assumed that t was even in this lemma just to simplify the proof; a similar bound can be found for t odd. Consider our previous example with $N = 15$, $g = 16$, and $t = 10$. In this case, the above bound is $|M_i| \leq 1$; so, for the network to be wide-sense nonblocking, at most 5 ADMs can be removed. Compare this to 20 ADMs that can be removed for a rearrangeably nonblocking ring. For this example, with t taking on any even value between 2 and 14, $|M_i|$ is always bounded to be less than or equal to 1, resulting in at most an 8% reduction in ADMs. These results suggest that to get great benefits from grooming for t -allowable traffic, some rerouting of existing traffic is needed, at least within the unidirectional ring model considered here.

B. Using a Hub Node and Tunable Lasers

Investing in more sophisticated components elsewhere in the network can yield further reductions in the cost of the electronic layer multiplexing. We consider two examples—the use of a hub node and the use of tunable lasers. First we consider a hub node. By a hub node, we mean a node which has ADMs on every wavelength and has a SONET DCS. The benefits of a hub architecture in reducing the required number of ADMs has been pointed out previously (see, e.g., [2] or [3]). For bidirectional rings, several hub architectures are given in [3] which support dynamic traffic. By similar arguments to those used in [2], we can show that making one node in the ring such a hub node will not require any more ADMs than were required without the hub. It can easily be shown that the minimum number of ADMs needed to support all t -allowable traffic with a single hub node is given by

$$\left\lceil \frac{N}{\lceil g/t \rceil} \right\rceil + (N - 1) \text{ ADMs.} \quad (13)$$

For example, consider a ring with $N = 7$, $g = 2$, and $t = 1$. Using the algorithm from Lemma 6, 12 ADMs are needed to support all t -allowable traffic. By making one node a hub node, this traffic can be supported using only 10 ADMs.

We can also reduce the required number of ADMs if, instead of having fixed tuned lasers, each node is equipped with tunable lasers. For example, again consider the ring with $N = 7$, $g = 2$ and suppose we want to support all 1-allowable traffic. If nodes are equipped with tunable lasers, then each node only needs one ADM, and thus only 7 ADMs are needed for the entire ring. In this case, using tunable lasers reduced the required number of ADMs by 58%. Clearly, with tunability, a node needs no more than t ADMs to support t -allowable traffic. Thus, when t is small, there is a clear advantage to tunability. On the other hand, for larger values of t , the gain from tunability is not as obvious and is an open issue. Both tunability and the use of a

hub node can also reduce the required number of ADMs for an arbitrary set of allowable traffic as studied in Section II.

V. CONCLUSION

In this paper we examine the problem of designing a WDM ring network to support dynamic SONET traffic. The goal of our design is to minimize the number of electronic multiplexers (e.g., SONET ADMs) used in the network. We developed a number of algorithms for assigning ADMs to wavelengths in a way that supports every allowable traffic matrix in a non-blocking manner. For the special case of t -allowable traffic, these algorithms are shown to reduce the number of ADMs needed by up to 27%. We also derive a lower bound on the number of ADMs required to support all t -allowable traffic, and show that in some cases our algorithms perform close to this bound. Finally, we discuss extensions of our model to include supporting traffic in a strictly nonblocking manner. Additionally, we discuss the use of a hub node and tunability to further reduce the number of ADMs.

REFERENCES

- [1] E. Modiano and A. Chiu, "Traffic grooming algorithms for minimizing electronic multiplexing costs in unidirectional SONET/WDM ring networks," in *Proc. CISS*, Feb. 1998.
- [2] A. Chiu and E. Modiano, "Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming," in *Proc. Globecom*, Nov. 1998.
- [3] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost effective traffic grooming in WDM rings," in *Proc. Infocom*, Apr. 1998.
- [4] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength assignment in a WDM ring to minimize the cost of embedded SONET rings," in *Proc. Infocom*, San Francisco, CA, Apr. 1998, pp. 94–101.
- [5] G. Sasaki, O. Gerstel, and R. Ramaswami, "A WDM ring network for incremental traffic," in *Proc. 36th Annu. Allerton Conf. Commun., Contr., Comput.*, Sept. 1998, pp. 662–672.
- [6] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET network design," in *Proc. Infocom*, New York, Mar. 1999, pp. 734–743.
- [7] J. Simmons, E. Goldstein, and A. Saleh, "On the value of wavelength add/drop in WDM rings with uniform traffic," in *Proc. OFC*, Feb. 1998.
- [8] —, "Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic," *IEEE J. Lightwave Technol.*, vol. 17, pp. 48–57, Jan. 1999.
- [9] X. Zhang and C. Qiao, "An effective and comprehensive approach to traffic grooming and wavelength assignment in SONET/WDM rings," in *SPIE Proc. Conf. All-Opt. Networking*, vol. 3531, Boston, MA, Sept. 1998.
- [10] T. Wu, D. Kong, and R. Lau, "An economic feasibility study for a broadband virtual path SONET/ATM self-healing ring architecture," *IEEE J. Select. Areas Commun.*, vol. 10, pp. 1459–1473, Dec. 1992.
- [11] J. Simmons and A. Saleh, "The value of optical bypass in reducing router size in gigabit networks," in *Proc. ICC*, Vancouver, BC, Canada, June 6–10, 1999, pp. 591–596.
- [12] R. Ramaswami and K. Sivarajan, "Design of logical topologies for wavelength routed optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 840–851, June 1996.
- [13] J. Bannister, L. Fratta, and M. Gerla, "Topological design of WDM networks," in *Proc. Infocom*, Los Alamitos, 1990.
- [14] S. Subramanian and R. Barry, "Wavelength assignment in fixed routing WDM networks," in *Proc. ICC*, Montreal, PQ, Canada, June 1997.
- [15] Y. Zhu, G. N. Rouskas, and H. G. Perros, "Blocking in wavelength routing networks, Part I: The single path case," in *Proc. Infocom*, New York, Mar. 1999.
- [16] O. Gerstel and S. Kutten, "Dynamic wavelength allocation in all-optical ring networks," in *Proc. ICC*, June 1997, pp. 432–436.
- [17] O. Gerstel, G. Sasaki, S. Kutten, and R. Ramaswami, "Worst-case analysis of dynamic wavelength allocation in optical networks," *IEEE/ACM Trans. Networking*, Dec. 1999.
- [18] D. Bertsimas and J. Tsitsiklis, *Introduction to Linear Optimization*. Belmont, MA: Athena Scientific, 1997.

- [19] F. Roberts, *Applied Combinatorics*. Englewood Cliffs, NJ: Prentice-Hall, 1984.
 [20] E. Modiano and A. Chiu, *IEEE J. Lightwave Technol.*, vol. 18, Jan. 2000.



Randall Berry 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.

In 1998 he was on the technical staff at MIT Lincoln Laboratory in the Advanced Networks group. Since Fall 2000, he has been with Northwestern University, Evanston, IL. He is currently an Assistant Professor in the Department of Electrical and

Computer Engineering. His research interests include wireless communication, data communication networks, and information theory.



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, both in electrical engineering, from the University of Maryland, College Park, in 1989 and 1992, respectively.

He was a Naval Research Laboratory Fellow between 1987 and 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.

Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory where he worked on communication protocols for satellite, wireless, and optical networks, and was the project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999, he has been a member of the faculty of the Aeronautics and Astronautics Department and the Laboratory for Information and Decision Systems (LIDS) at MIT, where he conducts research on communication networks and protocols with emphasis on satellite and hybrid networks, and high speed networks.