

The Role of Switching in Reducing the Number of Electronic Ports in WDM Networks

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

Abstract—We consider the role of switching in minimizing the number of electronic ports [e.g., synchronous optical network (SONET) add/drop multiplexers] in an optical network that carries subwavelength traffic. Providing nodes with the ability to switch traffic between wavelengths, such as through the use of SONET cross-connects, can reduce the required number of electronic ports. We show that only limited switching ability is needed for significant reductions in the number of ports. First, we consider architectures where certain “hub” nodes can switch traffic between wavelengths and other nodes have no switching capability. For such architectures, we provide a lower bound on the number of electronic ports that is a function of the number of hub nodes. We show that our lower bound is relatively tight by providing routing and grooming algorithms that nearly achieve the bound. For uniform traffic, we show that the number of electronic ports is nearly minimized when the number of hub nodes used is equal to the number of wavelengths of traffic generated by each node. Next, we consider architectures where the switching ability is distributed throughout the network. Such architectures are shown to require a similar number of ports as the hub architectures, but with a significantly smaller “switching cost.” We give an algorithm for designing such architectures and characterize a class of topologies, where the minimum number of ports is used. Finally, we provide a general upper bound on the amount of switching required in the network. For uniform traffic, our bound shows that as the size of the network increases, each traffic stream must be switched at most once in order to achieve the minimum port count.

Index Terms—Optical networks, synchronous optical network (SONET), traffic grooming.

I. INTRODUCTION

WAVELENGTH-division-multiplexing (WDM) is increasingly being deployed to provide high-capacity metro core networks. Typically, these networks have a synchronous optical network (SONET) ring architecture, where each node in the ring uses a SONET add/drop multiplexer (ADM) to electronically combine several lower rate streams onto a wavelength, e.g., 16 OC-3 circuits can be multiplexed onto one OC-48 stream. With WDM, multiple SONET rings

can be supported on a single fiber; however, each additional ring will require additional ADMs. The cost of these electronic ports dominates the cost of such a network. Managing this cost is particularly important in the cost-sensitive metro environment. To reduce the number of electronic ADMs, WDM add/drop multiplexers (WADMs) can be employed; WADMs allow a wavelength to either be dropped at a node or to optically bypass a node. When a wavelength is not dropped at a node, an electronic ADM is not required for that wavelength. The required number of SONET ADMs can be further reduced by *grooming* the lower rate traffic so that the minimum number of wavelengths need to be dropped at each node.

The benefits of grooming with WADMs have been looked at in a number of recent papers including [1]–[17]. The general grooming problem is NP-complete [1]. However, for several special cases, algorithms have been found that significantly reduce the required number of ADMs. For example, for uniform all-to-all traffic, algorithms have been found for both bidirectional rings [2], [4], [5] and unidirectional rings [1]. Heuristic algorithms for general (nonuniform) traffic have also been presented in [8]–[11]. In much of the work on grooming, such as [1], [3], [8]–[10], it is assumed that each low-rate circuit must stay on the same wavelength between the source and destination. This assumption can be relaxed when a node is equipped with a SONET digital cross-connect (DXC), which allows for the electronic switching of low-rate streams between SONET rings (i.e., wavelengths). The added flexibility provided by DXCs can enable the traffic to be more efficiently groomed, leading to a reduction in the required number of ADMs. An example of this is given in [1], where it is shown that even a single-*hub* node with a DXC can reduce the required number of ADMs over a network with no switching capability, even when the hub node is required to have an ADM on every wavelength. In [5], it was shown that the cost savings, in terms of ADMs, with a single-hub architecture can be as high as 37.5%. In other work, such as [4], it is assumed that every node can cross-connect every wavelength that is dropped at that node. Clearly, more switching capability will not increase the required number of ADMs. However, there is a nonnegligible cost associated with providing this electronic switching. Therefore, in addition to minimizing the required number of ADMs, it is also desirable to limit the amount of switching in the network.

In this paper, we consider architectures that are efficient both in terms of the number of ADMs used, as well as the amount of switching provided. First, we consider *multiple-hub architectures*, which are a generalization of the single-hub architecture in [1]. In this case, we assume that there are two types of nodes in the ring—hub nodes and nonhub nodes. Each hub node

Manuscript received September 7, 2003; revised March 15, 2004. The work of E. Modiano was supported in part by the Defense Advanced Research Projects Agency (DARPA) under Contract MDA972-02-1-0021 and in part by the National Science Foundation (NSF) under Grant ANI-0073730 and Grant ANI-0335217. This paper was presented at the 2001 Allerton Conference on Communication, Control and Computing, October 2001, and the Joint Conference on Information Sciences (JCIS), March 2002.

R. A. 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 02139-4307 USA (e-mail: modiano@mit.edu).

Digital Object Identifier 10.1109/JSAC.2004.830371

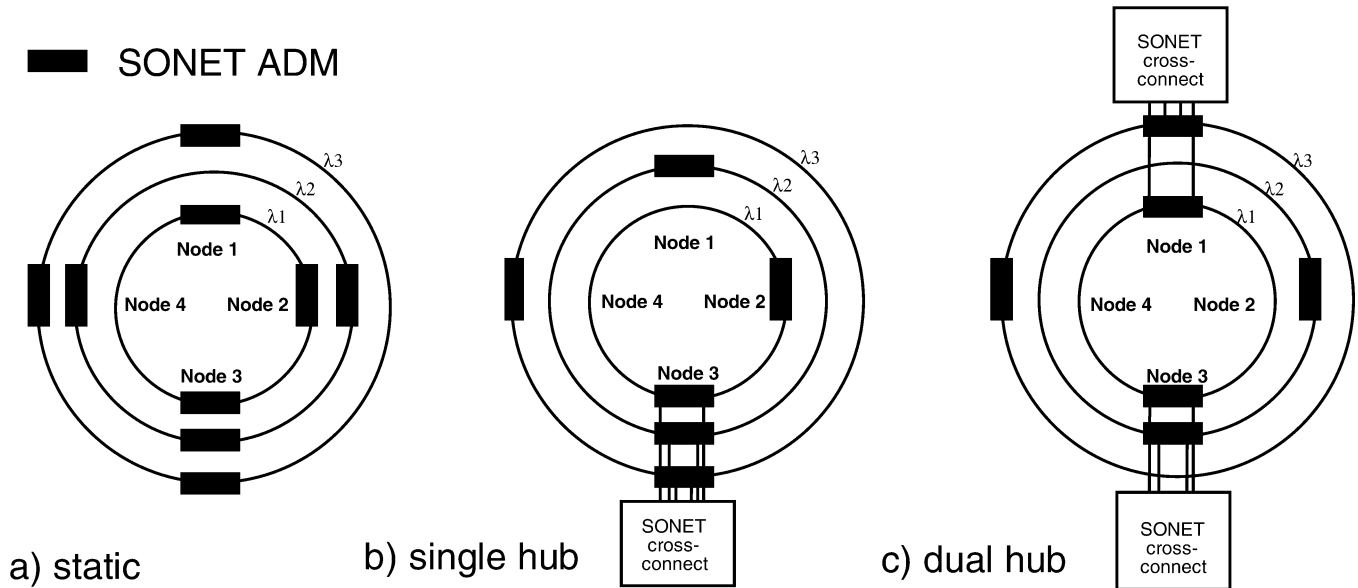


Fig. 1. Possible grooming architectures.

can cross-connect every wavelength dropped at the node, while nonhub nodes have no DXCs. In this case, we study the number of hub nodes needed to minimize the number of ADMs in the network. We bound the number of ADMs needed for a multihub architecture and provide algorithms for traffic grooming in such a ring. Next, we consider *distributed hub architectures*, where each node in the ring may have limited cross-connect capability. In this case, each node is able to switch traffic between only a subset of the wavelengths dropped at that node. We introduce a new notion of “switching cost” that quantifies the amount of switching used in a ring. We give examples to show that a distributed hub architecture can result in an efficient use of ADMs, as well as a smaller switching cost than a multihub architecture. We identify a class of rings, where a distributed hub architecture can be found that requires the minimum number of ADMs. We again provide a heuristic algorithm for grooming and switching in this type of ring. Finally, we consider an upper bound on the average amount of switching needed in a network that minimizes the required number of ADMs. This bound is general in that it applies to an arbitrary topology and is useful in that it provides additional insight into the amount of switching needed in a network.

II. MULTIPLE-HUB ARCHITECTURES

In this paper, we primarily consider unidirectional ring networks such as a UPSR SONET ring. This is done mainly to simplify our description; as will be evident, much of the following can be easily generalized to bidirectional rings and, in some cases, to arbitrary mesh networks. Let the network nodes be represented by the set $\mathcal{S}_N = \{1, 2, \dots, N\}$. Also, for simplicity, we assume that all traffic has the same granularity of g , i.e., g low-rate circuits can be combined on each wavelength.

In Fig. 1, three possible ring architectures for a ring with $N = 4$ nodes are shown. Fig. 1(a) is a static ring without cross-connects. In this architecture, no switching is employed, hence, each circuit must be assigned to a single wavelength that must

be processed (dropped) at both the source and the destination. For example, all traffic between nodes 1 and 2 must be assigned to λ_1 . This static architecture is the traditional SONET ring architecture that has been used in the studies of [1]–[3]. Fig. 1(b) depicts a single-hub architecture, where a large cross-connect is located at one hub (node 3). The cross-connect is able to switch any low-rate circuit from any incoming wavelength to any outgoing wavelength. With this architecture, each node can send all of its traffic to the hub node, where the traffic is switched, groomed, and sent back to the destination nodes. Finally, shown in Fig. 1(c) is a multiple-hub architecture, with two hub nodes (nodes 1 and 3). Each hub node has a small cross-connect that can switch traffic among the wavelengths dropped at that node. Each node on the ring can send a fraction of its traffic to one of the hub nodes, where it is properly groomed and relayed to its destination.

To illustrate the potential benefit of the multiple-hub architecture, consider a unidirectional ring with $N = 9$ nodes, where each wavelength supports an OC-48 and traffic demand is uniform with two OC-12s between each pair, i.e., $g = 4$ and the traffic demand is for $r = 2$ low-rate circuits. In this case, each node generates 16 OC-12s or four wavelengths of traffic. With the single-hub solution, each node can send all four wavelengths of traffic to be groomed at the hub at, say, node 1. Thus, with a single hub, each node would use four ADMs, and the hub would need $8 \times 4 = 32$ ADMs for a total of 64 ADMs. In a two-hub architecture, each node would send two wavelengths worth of traffic to each hub (for example, at nodes 1 and 5); an additional wavelength would be used for traffic between the two hubs, resulting in 58 ADMs. Finally, a four-hub architecture can be used, where each node sends one wavelength to each of four hubs and some additional ADMs are used to handle the interhub traffic. Using one of the grooming algorithms that is developed in Section III-B, a four-hub architecture can be found for this ring that requires only 52 ADMs. In the next section, we give a lower bound on the number of ADMs required assuming unlimited switching capability; for this example, that bound would be

48 ADMs. Thus, with four hubs the bound is nearly met, and any further increase in the amount of switching could at best result in only a moderate additional savings of ADMs.¹ Notice that in this case the number of hubs is equal to the number of wavelengths generated by a node. It can also be shown that using the four-hub architecture reduces the required number of wavelengths from 32 to 26. Thus, the four-hub architecture is more efficient both in the use of wavelengths, as well as ADMs.

A. Bounds on the Required Number of ADMs

In the following, we develop a lower bound on the required number of ADMs for a K -hub architecture. We consider the case where there is a uniform traffic demand of $r \leq g$ circuits between each pair of nodes in the ring.² In this case, a lower bound on the number of ADMs needed in a unidirectional ring, assuming unlimited switching ability, is given in the following proposition, first derived in [5].

Proposition 1 ([5]): The number of ADMs A needed to support uniform traffic in a unidirectional SONET ring with parameters (N, r, g) is bounded by

$$A \geq \frac{2N(N-1)r}{g+r}. \quad (1)$$

A *lightpath* in a ring refers to a single wavelength connection between two nodes, which is not dropped at any intermediate nodes. Thus, when a circuit is carried over n lightpaths, it is either dropped and continued or switched $n-1$ times; in this case, each lightpath is said to carry $1/n$ th of the “full” circuit. The lower bound in Proposition 1 is obtained by recognizing that each lightpath in the network must be terminated with exactly two ports. Thus, a lower bound on the number of lightpaths L needed to support all of the traffic in the network can be translated into a lower bound on the number of ADMs, A . Since the direct traffic between two nodes is equal to r low-rate circuits, each lightpath can at most carry r “full” circuits entirely from their source to their destination. The remaining capacity of that lightpath ($g-r$) can only be used to carry circuits that are also carried on *at least* one other lightpath. Hence, each lightpath can carry at most $Q = r + (g-r)/2$ “full” circuits. Since the total traffic demand under the uniform traffic assumption is equal to $L = N(N-1)r$ circuits, the number of lightpaths required is lower-bounded by L/Q . Each lightpath is terminated at a port; therefore, the number of ports needed is at least $2L/Q$. However, a SONET ADM can be used both as a receiving and transmitting port, hence, $A \geq L/Q$ and the bound follows.³

The bound in Proposition 1 is not tight in general, but it can be achieved in several cases. It is insightful to consider some characteristics of these cases. From the above, it can be seen that for (1) to be tight, each lightpath must be efficiently packed so that it contains Q “full” circuits. This in turn requires the following three conditions to be met:

- each lightpath must be fully utilized;
- no circuit can travel over more than 2 lightpaths;
- each lightpath must carry r full circuits (directly from the source to the destination).

To see that these conditions can indeed be satisfied, consider the case where $(N-1)r = g$, i.e., each node generates a full wavelength worth of traffic. Suppose a single cross-connect hub is chosen and all traffic is sent to the hub, where it is switched and sent back to its destination. In this case, the above conditions are met, and the bound in (1) is tight. However, in general it is not possible to achieve the bound by using this single-hub architecture. This is because each node only has r circuits whose final destination is the hub. Thus, when a node generates more than one wavelength worth of traffic, each wavelength sent to the hub cannot contain r full circuits, as required by condition c) above. When all of the traffic is routed through a single hub, only $2(N-1)r$ circuits can be carried on a single lightpath and the remaining $(N-1)(N-2)r$ circuits must traverse two lightpaths. Since each lightpath can carry at most g circuits, the total number of lightpaths L (and hence, A) is bounded by

$$\begin{aligned} L = A &\geq (2(N-1) + 2(N-1)(N-2)) \left(\frac{r}{g}\right) \\ &= \frac{2(N-1)^2 r}{g}. \end{aligned} \quad (2)$$

Notice that the difference between the right-hand side of the bound in (2) and the bound in (1) is

$$\frac{2(N-1)r}{g(g+r)} ((N-1)r - g).$$

This is strictly positive unless $(N-1)r = g$, i.e., each node generates a single wavelength of traffic.

The above considerations lead us to consider a multiple-hub architecture. We define a *K-hub architecture* to be a ring with K -hub nodes, with the restriction that all traffic between nonhub nodes must be routed to one of K hubs. We do not allow traffic between nonhub nodes to be sent directly (without going through a hub). Our main reason for this restriction is to focus on architectures that are simple to design, implement and analyze. Also, as will be seen in the following, relaxing this restriction cannot result in significant improvements. Assume each of the K hubs has a cross-connect capable of switching any circuit from any input wavelength to any output wavelength. Again, consider a unidirectional ring with N nodes, a traffic granularity of g and uniform traffic with r circuits between each pair. With K hubs (and $N-K$ nonhub nodes), a total of $2(N-K)Kr$ circuits can be routed between the hubs and the nonhubs in one hop. The remaining traffic between the nonhub nodes, of which there are $(N-K)(N-K-1)r$ circuits, will traverse two lightpaths. Therefore, all traffic that is either to or from a nonhub node requires at least

$$\begin{aligned} &\frac{2(N-K)Kr + 2(N-K)(N-K-1)r}{g} \\ &= \frac{2(N-K)(N-1)r}{g} \end{aligned} \quad (3)$$

lightpaths. Additionally, we have to account for the traffic between hub nodes. By the same reasoning as used in deriving

¹Indeed, a four-hub architecture requiring only 49 ADMs can be found; thus the possible savings with more switching capability is at most one ADM.

²If there is more than a full wavelength of traffic between a pair of nodes, then assigning each full wavelength to a direct lightpath is clearly optimal. Therefore, this traffic can be ignored for our purposes.

³The bound can clearly be made tighter by including a ceiling; for large g/r the bound can also be tightened by taking the maximum of L/Q and N , since each node must have at least one ADM.

(1), this traffic requires at least $2K(K-1)r/(g+r)$ lightpaths. Hence

$$A \geq \frac{2(N-K)(N-1)r}{g} + \frac{2K(K-1)r}{g+r}. \quad (4)$$

Since the bound in (1) does not depend on the number of hubs, (4) can be tightened by combining it with (1). This yields the following bound on the number of ADMs.

Proposition 2: For a K -hub architecture, the number of ADMs A satisfies

$$A \geq \max \left\{ \frac{2(N-K)(N-1)r}{g} + \frac{2K(K-1)r}{g+r}, \frac{2N(N-1)r}{g+r} \right\}. \quad (5)$$

Some insight can be gained from examining the behavior of (5) as K , the number of hubs, varies. Notice that only the first quantity inside the maximization in (5) varies with K , we denote this quantity by $A(K)$. When $K = (N-1)r/g$, $A(K) = 2N(N-1)r/(g+r)$, i.e., the two quantities in the maximization in (5) are equal. Hence, when $(N-1)r/g$ is an integer, this number of hubs minimizes the bound in (5). To address the case where $(N-1)r/g$ is not an integer, we note that for

$$K \leq 0.5 \left[(N-1) \left(1 + \frac{r}{g} \right) + 1 \right] \quad (6)$$

$A(K)$ can be shown to be decreasing in K , otherwise, it is increasing. Furthermore, for $r \leq g$

$$\left[(N-1) \frac{r}{g} \right] < 0.5 \left[(N-1) \left(1 + \frac{r}{g} \right) + 1 \right]. \quad (7)$$

Thus, we have that for $r \leq g$, the number of hubs, K^* , that results in the smallest bound on the number of ADMs is given by $K^* = \lceil (N-1)r/g \rceil$; i.e., K^* is equal to the number of wavelengths of traffic generated by each node. Notice that when using K^* hubs, the lower bound in (5) is equal to the lower bound in (1), which did not have the restriction that traffic had to be routed through a hub. The above example, where $(N-1)r = g$, provides one case where this bound is tight using $K^* = 1$ hub. As another example, consider the case where $r = g$, i.e., there is a full wavelength traffic demand between each pair of nodes. Setting up one lightpath between every pair of nodes is clearly the optimal way to route this traffic. This requires $N(N-1)$ ADMs, which meets the bound in (5) with $K^* = N$ hubs, i.e., each node is essentially a hub. We note in this case, however, that no switching is required at the hubs.

At this point we have bounded the number of ADMs in a K -hub architecture, and we have shown that the number of hubs that optimizes this bound is given by K^* . This does not tell us how to groom traffic or, in general, how tight this bound will be. In the next section, we develop some simple grooming algorithms for a K -hub architecture, where each nonhub node sends its traffic to one or more of the hubs. For these algorithms, we will see that, indeed, K^* is (approximately) the optimal number of hubs and that the bound in (5) can be approached closely in many cases.

B. K -Hub Grooming Algorithms

We consider several simple grooming algorithms for a K -hub architecture in a ring with $N > K$ nodes. For the purpose of describing these algorithms, the exact location of the hubs is irrelevant.

1) *Group Algorithms:* The first type of multihub architectures we discuss involves grouping the N nodes in the ring into K distinct groups, each of approximate equal size N/K . Of course, when K does not divide N , group sizes may differ by one. Each hub node is associated with exactly one group. Given such a division of the nodes, several possible grooming/routing algorithms are possible. One natural approach would be for all nonhub nodes within a group to send and receive all of their traffic from the hub node associated with the group. The hub nodes would then exchange all traffic between groups.⁴ This requires $\lceil (N-1)r/g \rceil$ ADMs at each nonhub node; a corresponding number of ADMs is also required at each hub node for the traffic to and from the nonhub nodes. The intergroup traffic can be handled by making one hub a “super-hub” which switches and distributes all intergroup traffic. With this approach, the total ADMs requirement for this architecture can be upper-bounded by

$$A \leq \left\lceil \frac{(N - \lfloor \frac{N}{K} \rfloor) \lfloor \frac{N}{K} \rfloor r}{g} \right\rceil 2K + 2(N-K) \left\lceil \frac{(N-1)r}{g} \right\rceil. \quad (8)$$

Notice that with this architecture, traffic between nonhub nodes in different groups needs to be switched at the hub for each group. Such traffic would then be carried over three lightpaths. As discussed in Section III, this precludes such architecture from ever attaining the bound in (1). We consider a variation of this architecture where all traffic travels over at most two lightpaths. Specifically, assume that every node, including the hub nodes, now sends all traffic destined to any node in a group to the respective hub node.⁵ The hub nodes once again distribute the traffic to the nonhub nodes in their group. Exact computation of the ADM requirement for this architecture is cumbersome because of the fact that K does not always divide N . Instead, we proceed with the following approximate, yet insightful, analysis. Assume each node sends $1/K$ of its total traffic to each hub (this assumption would be exact if all groups were of equal size). Hence, each node sends $\lceil (N-1)r/Kg \rceil$ wavelengths of traffic to each of the K hubs. In addition, each hub node must send the groomed traffic to its subsidiary nodes. Each subsidiary node must receive a total of $(N-1)r$ circuits using $\lceil (N-1)r/g \rceil$ wavelengths. Hence, each nonhub node generates $K \lceil (N-1)r/Kg \rceil$ wavelengths worth of traffic and receives $\lceil (N-1)r/g \rceil$ wavelengths. This can be accomplished using no more than $K \lceil (N-1)r/Kg \rceil$ ADMs at each nonhub node. Now, each hub node receives $\lceil (N-1)r/Kg \rceil$ wavelengths of traffic from each of $(N-1)$ nodes and each hub node sends $(K-1) \lceil (N-1)r/Kg \rceil$ to the other hub nodes. Also, each hub node must send $\lceil (N-1)r/g \rceil$ wavelengths of

⁴The “hierarchical ring” proposed in [4] is similar to this type of architecture.

⁵An analogous architecture can be considered, where all nodes in a group send their traffic to the hub node for the group, and the hub node then forwards the traffic to the destination.

traffic to each of its subsidiary nodes. Hence, the number of wavelengths sourced and terminated at each hub node is approximately the same and equal to $(N - 1) \lceil (N - 1)r/Kg \rceil$. Summing over all of the nodes, the total number of ADMs required is equal to

$$(N - K)K \left\lceil \frac{(N - 1)r}{Kg} \right\rceil + K(N - 1) \left\lceil \frac{(N - 1)r}{Kg} \right\rceil = K(2N - K - 1) \left\lceil \frac{(N - 1)r}{Kg} \right\rceil. \quad (9)$$

With this algorithm each circuit travels over at most two lightpaths. However, notice that each nonhub node receives all its traffic from the corresponding hub. Thus, when every node generates more than one wavelength worth of traffic, each lightpath terminated at a nonhub node cannot contain r direct circuits, which is another requirement for the bound in (1) to be met. In the next section, we consider an algorithm, where each nonhub node sends and receives traffic from all of the hub nodes. This approach allows traffic to more closely emulate the characteristics for achieving the bound in (1).

2) *Symmetric Algorithm:* In this algorithm, each nonhub node divides its traffic so that it sends approximately an equal amount to each of the K hubs. The traffic sent from a given nonhub node to a given hub will include traffic whose final destination is that hub, as well as traffic for other nonhub nodes. The traffic for other nonhub nodes will be switched at the hub and forwarded to its destination. Suppose that each nonhub node can divide its traffic to satisfy the following two conditions.

- i) No more than $H = \lceil (N - 1)r/Kg \rceil$ wavelengths of traffic are sent to each hub from each nonhub node.
- ii) No more than H wavelengths of traffic are received at any nonhub node from any hub.

If the traffic can be divided in this way, each nonhub node will require at most KH ADMs and each hub node will require at most $(N - K)H$ ADMs for sending traffic to nonhub nodes. Thus, all traffic either to or from the nonhub nodes can be supported using at most $2K(N - K)H$ ADMs. Next, we give one construction which shows that the traffic can indeed be divided to satisfy the above two conditions.

Let the nonhub nodes be numbered $1, 2, \dots, N - K$ and the hub nodes be numbered $1, \dots, K$. Recall that the traffic demand between each pair of nodes is r circuits. For $l = 1, \dots, r$, route the l th circuit between nonhub nodes i and j , through hub k , where

$$k = \begin{cases} \lceil ((i - j)r + l) \bmod K \rceil, & \text{if } j > i \\ \lceil ((i - j - 1)r + l) \bmod K \rceil, & \text{if } j < i \end{cases}. \quad (10)$$

This assignment ensures that the hub nodes are evenly loaded and can be thought of as follows: the circuits from any nonhub node to all other nonhub nodes are listed and uniquely labeled with one of $(N - K - 1)r$ consecutive integers. Each circuit is then sent to the hub that corresponds to its label mod K . This results in at most $\lceil (N - K - 1)r/K \rceil$ circuits of nonhub node to nonhub node traffic being sent to each hub from any nonhub node. Each nonhub node will also send the traffic for a given hub

node directly to that hub node; including this traffic we have at most $\lceil (N - K - 1)r/K \rceil + r$ circuits being sent to each hub from each nonhub node. This requires at most

$$\left\lceil \frac{\lceil (N - K - 1)r/K \rceil + r}{g} \right\rceil = \left\lceil \frac{\lceil (N - 1)r/K \rceil}{g} \right\rceil = \left\lceil \frac{(N - 1)r}{Kg} \right\rceil = H \quad (11)$$

wavelengths, where the last equality follows since g is an integer. This shows that condition i) is satisfied by this traffic assignment. Essentially, the same arguments can be used to show that condition ii) is also satisfied by this assignment.

So far, we have only addressed traffic to or from the nonhub nodes. In addition, interhub traffic must also be accommodated. The simplest way to accomplish this is by making one of the K -hub nodes a ‘‘super-hub,’’ to which all hub nodes send their interhub traffic. The super-hub then distributes the interhub traffic to the respective hubs. This requires an additional $2(K - 1) \lceil (K - 1)r/g \rceil$ ADMs for the interhub traffic. Thus, the total number of ADMs required for the above algorithm is given by

$$A = 2K(N - K) \left\lceil \frac{(N - 1)r}{Kg} \right\rceil + 2(K - 1) \left\lceil \frac{(K - 1)r}{g} \right\rceil. \quad (12)$$

While this simple algorithm is generally effective, it should be immediately obvious that when the number of hub nodes is large the algorithm becomes inefficient. This is because the interhub traffic is handled using a single-hub architecture. We know, from our earlier discussion, that when the traffic among nodes exceeds a single wavelength, a single-hub architecture is inefficient. A further improvement can be obtained by using a hierarchical architecture with multiple ‘‘super-hubs’’ that are used for routing the interhub traffic.

The required number of ADMs in such a hierarchical architecture can be calculated recursively. Specifically, let $A(N, K)$ denote the minimum number of ADMs needed for an architecture with N nodes and K hubs, where traffic that originates at the nonhub nodes is routed as above, and interhub traffic is handled using a hierarchical architecture. Let

$$A^*(N) = \min_{K \leq N} \{A(N, K)\}$$

denote the minimum number of ADMs needed when the optimum number of hubs is used. Then assuming that the optimum number of ‘‘super-hubs’’ is used in the above architecture, we have

$$A(N, K) = 2K(N - K) \left\lceil \frac{(N - 1)r}{Kg} \right\rceil + A^*(K). \quad (13)$$

Using (13), the number of ADMs needed for a hierarchical K -hub architecture can be recursively calculated.

The results from using the symmetric algorithm are shown in Table I. The five columns on the right show the number of ADMs required when using $K = 1$ to five hubs, respectively. For example, in the case of a 17-node ring, the minimum ADM solution is achieved with four hubs. Highlighted in the table is the solution that achieves the minimum number of ADMs. This

TABLE I
NUMBER OF ADMs NEEDED WITH MULTIPLE HUBS

			Symmetric Algorithm				
N	K^*	Bound	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$
5	1	8	8	14	16	14	8
6	2	12	20	18	22	22	18
7	2	17	24	22	28	30	28
8	2	22	28	26	34	38	38
9	2	29	32	30	40	46	48
10	3	36	54	66	46	54	58
11	3	44	60	74	52	62	68
12	3	53	66	82	58	70	78
13	3	62	72	90	64	78	88
14	4	73	104	98	136	86	98
15	4	84	112	106	148	94	108
16	4	96	120	114	160	102	118
17	4	109	128	122	172	110	128

corresponds exactly to the number of hubs K^* that optimizes the lower bound from Section II-A, which is equal to the number of wavelengths of traffic generated at each node. The value of K^* is given in column 2 and the lower bound when using the optimal number of hubs K^* is given in column 3.

III. DISTRIBUTED HUB ARCHITECTURES

In this section, we relax the assumption that each node is either a hub or a nonhub node, and we allow only a subset of the wavelengths dropped at a node to be switched. In this case, instead of a few hub nodes with complete switching capability, each node may have some partial switching capability, provided by a small cross-connect. Such an architecture has several advantages. The first advantage is that the size of a cross-connect is a significant component of its cost. Using several smaller cross-connects may lead to lower costs. Distributing the switching requirements over all nodes in a ring also allows for more uniform node requirements. Finally, with such an architecture, the ring may be more robust to node failures. The cost of such an approach may be an increase in the complexity of control and management.

To quantify the amount of switching used in different architectures, we assign a *switching cost* of $(ng)^2$ to a DXC that can cross-connect low-rate traffic between n wavelengths. Assuming that the DXC is a crossbar switch, this cost is equal to the number of cross-points in the switch. This is a common metric used in studying switch designs. If multistage switch architectures are used, then this cost could be modified to reflect this. However the above metric will suffice to illustrate our points. The total switching cost for a ring architecture is then the sum of the switching costs of all DXCs in the ring.

A. Example

Consider a unidirectional ring with $N = 9$ nodes, a traffic granularity of $g = 2$ and uniform traffic demand of $r = 1$ circuit between each pair of nodes. In this case, from Proposition 1, we have a lower bound of 48 ADMs. First, we consider supporting this traffic using the symmetric architecture from Section II-B. Each node generates four wavelengths worth of traffic. Thus, from (13), this traffic can be supported with four hub nodes and 50 ADMs. Each hub node receives one wavelength from




Fig. 2. Architecture for supporting traffic between the first group of nodes (1, 2, 3) in the example from Section III-A. The two horizontal lines correspond to two wavelengths. The traffic sent on each lightpath is also indicated in the figure. The cross-connect, indicated by the vertical line is at node 2.

each of the five nonhub nodes and must be able to switch circuits between these wavelengths. This requires a $5g \times 5g$ DXC. Therefore, the switching cost of this architecture is greater than⁶ $4(100) = 400$.

Next, we describe a distributed switching architecture for supporting the same traffic. Consider dividing the nodes into the following groups of three:

- (1, 2, 3) (4, 5, 6) (7, 8, 9) (1, 4, 7)
- (1, 5, 8) (1, 6, 9) (2, 5, 7) (2, 6, 8)
- (2, 4, 9) (3, 6, 7) (3, 5, 9) (3, 4, 8).

Notice that each pair of nodes is in exactly one of these groups. The traffic between all three nodes in each group can be supported by having two of the nodes send all of their traffic to the third node, as depicted in Fig. 2. A $2g \times 2g$ DXC at the third node can be used to switch the incoming traffic, which can then be forwarded to its destination. This requires four ADMs and a switching cost of $(2g)^2 = 16$. Since there are 12 groups, supporting all of the traffic requires 48 ADMs and a total switching cost of 192. Notice that in this case we are using the minimum number of ADMs given by the bound in Proposition 1 and the switching cost is over 50% less than the cost for the symmetric multihub architecture. Also, notice that any node within each group could serve as the “hub” for that group. For example, the switching capability could be spread out among all the nodes in the ring or concentrated at only four nodes.

B. Perfect Architectures

The distributed architecture in the above example meets the lower bound on the required number of ADMs from Proposition 1. As discussed in Section II-A, for any architecture that meets the bound in Proposition 1 with equality, each lightpath must contain r direct circuits and $g - r$ circuits that travel over two lightpaths. Assuming that all the indirect circuits must be switched, then at least a $2g \times 2g$ DXC is required for each pair of indirect circuits; this is exactly what is used in the above example. In this section, we consider a generalization of the above example to other rings, i.e., other values of N , g , and r . We call such an architecture a *perfect architecture*. Specifically, in a perfect architecture the nodes in the ring are divided into groups of $g/r + 1$ nodes such that each pair of nodes is in at most one group. One node in each group serves as a hub node for the group. All other nodes send their traffic to this hub node, where it is switched and forwarded to its destination. For a given set of parameters (N, r, g) , it may

⁶The actual switching requirements will be larger than this because we have not accounted for the switching required for interhub traffic.

not be possible to divide the traffic in the above manner; i.e., a perfect architecture may not exist (e.g., if g/r is not an integer). However, when it does exist it uses the minimum number of ADMs for any architecture.

Proposition 3: If a perfect architecture exists for a unidirectional ring with parameters (N, r, g) , then it uses the minimum number of ADMs of any architecture that can support this traffic.

Proof: Each group of $g/r + 1$ nodes requires $2(g/r)$ ADMs and supports $(g/r + 1)(g/r)r/2$ circuits. Since there are $N(N - 1)r/2$ total circuits, the total number of groups is

$$\frac{N(N - 1) \left(\frac{r}{2}\right)}{\left(\frac{g}{r} + 1\right) \left(\frac{g}{r}\right) \left(\frac{r}{2}\right)} = \frac{N(N - 1)r^2}{(g + r)g}. \quad (14)$$

Hence, the total number of ADMs is

$$\left(\frac{N(N - 1)r^2}{(g + r)g}\right) \left(\frac{2g}{r}\right) = \frac{2N(N - 1)r}{g + r}. \quad (15)$$

This exactly meets the lower bound in Proposition 1 and, therefore, is the minimum number of ADMs needed for any architecture. ■

Let $M = g/r + 1$ and assume that this is an integer. The problem of finding groups of M nodes for which a perfect architecture exists can be described in graph theoretic terms. Consider a fully connected graph with N nodes; denote this graph by K_N . Assume each node in this graph represents a node in the ring; a pair of nodes is represented by a link in this graph. Each group in the above construction can be viewed as a fully connected subgraph with M nodes. The above construction gives a family of subgraphs that are edge disjoint and cover the graph, K_N . Such a family is referred to as a *decomposition* of the original graph. In this case, each subgraph in the decomposition is isomorphic to K_M (a fully connected graph with M nodes). This is referred to as a K_M -decomposition of K_N . In these terms, a perfect architecture can be found if there exists a K_M -decomposition of K_N , where $M = g/r + 1$ is an integer.

The problem of graph decompositions has been well studied in the graph theory literature and is related to combinatoric problems such as finding a block orthogonal designs or Steiner triple systems of a given order [18]. The next proposition provides a necessary condition for the existence of a K_M -decomposition of K_N .

Proposition 4 ([18]): If there exists a K_M -decomposition of K_N , then the following hold:

$$M - 1 | N - 1 \text{ and } M(M - 1) | N(N - 1).$$

Here, we use the notation $a|b$ to denote that a is a divisor of b . Furthermore, the above conditions can be shown to be sufficient for all but a finite number of values of M and N [18]. By combining the above arguments, we have that unless $M - 1 | N - 1$ and $M(M - 1) | N(N - 1)$, where $M = g/r + 1$, a perfect architecture cannot be found. Also, except for a finite number of values of M and N the above conditions are sufficient. Notice that for the example in Section III-A, the above conditions are met.

When a perfect architecture can be found, it will have $N(N - 1)/K(K - 1)$ DXCs, and each DXC will have a switching cost of $((K - 1)g)^2$. Thus, the total switching cost is

$$N(N - 1) \left(1 - \frac{1}{K}\right) = N(N - 1) \left(\frac{g^3}{g + r}\right). \quad (16)$$

C. Grooming Algorithm

From the preceding section, for an arbitrary N , g , and r , a perfect architecture may not exist. In this section, we give a heuristic algorithm for routing and grooming traffic for an arbitrary ring that attempts to mimic a perfect architecture. The basic idea of this algorithm is to first find subsets of the total traffic requirement that are similar to the subsets used in a perfect architecture. Each subset will then be supported using a single small cross-connect at one hub node. In a perfect architecture, these subsets of traffic correspond to all-to-all traffic among a group of M nodes, where each node in the group generates a full wavelength worth of traffic. In the general case, these subsets will not necessarily correspond to all-to-all traffic between the nodes in a group. In particular, a pair of nodes may appear in multiple groups, but the traffic between the pair will only be assigned to one of the groups. In addition, each node in a group may not generate a full wavelength worth of traffic.

The algorithm sequentially forms groups of nodes and, for each group, a corresponding subset of the offered traffic. Each group is formed by adding nodes, in a greedy fashion, in an attempt to form perfect subsets. The corresponding traffic subset is simultaneously formed by adding all the remaining traffic between each node added to the group and the nodes already in the group. By remaining traffic, we mean circuits that have not yet been assigned to another group.

We give a more precise description of this algorithm next. To simplify the discussion, we only describe the case where $r = 1$. We maintain a list of the circuits C_i originating at each node i in the ring that have not yet been assigned to a subset. A list of the nodes in each group and the corresponding traffic subsets are also maintained. When we say a node i is *added to a group*, this implies that all the remaining traffic between that node and any other node in the group is assigned to the traffic subset. A *bidirectional circuit* between two nodes i and j includes both a circuit from i to j and a circuit from j to i .

Grouping Algorithm

- 1) Set $n = 1$.
- 2) Choose as the first node in group G_n , a node with the maximal remaining circuits to be assigned.
- 3) Add a node to group G_n that will result in the largest increase in the number of circuits in the corresponding traffic subset.
- 4) If more than one node in group G_n has g bidirectional circuits in the traffic subset, or there are no circuits that are not yet assigned to the nodes in group G_n (i.e., C_i is empty for all i in G_n) continue to step 5). Otherwise, go to 3).
- 5) If all circuits have been assigned, stop. Otherwise, set $n = n + 1$, go to 1.

In steps 2) and 3), any ties can be broken arbitrarily; for example, by choosing the node with the smallest label. Step 4)

ensures that no nonhub node in the resulting architecture will generate more than one wavelength worth of traffic.

As an example, consider applying the algorithm to the ring from Section III-A, with $N = 9$, $g = 2$, and $r = 1$. Assume that all ties are broken by choosing the node with the smallest label. Initially, all nodes have all eight circuits to be assigned, so any node can be chosen to start the group G_1 . Using our tie breaking rule, we choose node 1. Next, adding any other node to G_1 will result in adding one bidirectional circuit to the traffic subset, so node 2 will be chosen. At this point, both nodes in the group have $1 < g$ bidirectional circuits assigned, so we may add another node. Again, the choice of node does not matter; so node 3 is chosen. Now, $G_1 = \{1, 2, 3\}$, and each node has g bidirectional circuits in the traffic subset, thus we begin forming G_2 . Every node except 1, 2, or 3 have all eight circuits yet to be assigned, and thus can be chosen to start G_2 . Continuing in this manner, it can be shown that the algorithm will form the groups given in the example.

After forming groups using the above algorithm, the traffic for each group can then be supported using a single DXC at one “hub” node for the group. This “hub” node will be chosen from the nodes that have the maximal number of circuits in the traffic subset. Each “nonhub” node in the group will generate no more than one wavelength worth of traffic and send all of the traffic to the hub node. If there are K nodes in a group and each nonhub node uses a different wavelength, the traffic can be supported using $2(K - 1)$ ADMs and a switching cost of $(Kg)^2$. In cases where each node in the group does not generate a full wavelength of traffic, the number of ADMs and the switching cost can often be reduced by allowing nodes to share a wavelength. If all traffic must go through the DXC, then assigning traffic to wavelengths to minimize the needed number of ADMs is equivalent to the egress grooming problem studied in [1]. This problem can be reduced to the well-known bin packing problem [1]; any heuristic for the bin packing problem can then be used to assign the traffic to wavelengths.

As an example of this algorithm consider a ring with $N = 6$, $g = 4$, and $r = 1$. In this case, $g/r + 1 = 5$, and 4 is not a divisor of 5, so a perfect architecture cannot be found. Using the above algorithm results in the following subsets of traffic.

Subset 1: all-to-all traffic between $\{1, 2, 3, 4, 5\}$.

Subset 2: traffic between 6 and $\{1, 2, 3, 4, 5\}$.

The first subset of traffic requires eight ADMs and a switching cost of $(4g)^2 = 256$. The second subset requires seven ADMs and a switching cost of $(2g)^2 = 64$. Therefore, this architecture requires 15 ADMs and a total switching cost of 320. For comparison, the best symmetric hub architecture from Section II-B will require 18 ADMs and a switching cost of 512.

A more extensive comparison is shown in Figs. 3 and 4. In Fig. 3, the required number of ADMs for a distributed hub architecture based on the above algorithm is plotted for a ring with $g = 16$ and $r = 1$ as the number of nodes varies from $N = 6$ to 20. The number of ADMs needed for the symmetric hub architecture and the lower bound from Proposition 1 are also shown. For comparison, we also give a lower bound from [1] on the number of ADMs required without switching. In general, this lower bound is overly optimistic, i.e., the actual number of ADMs required without switching is typically greater than this




Fig. 3. ADMs needed with symmetric hub architecture, distributed hub architecture, and lower bound for ring with $r = 1$ and $g = 16$.

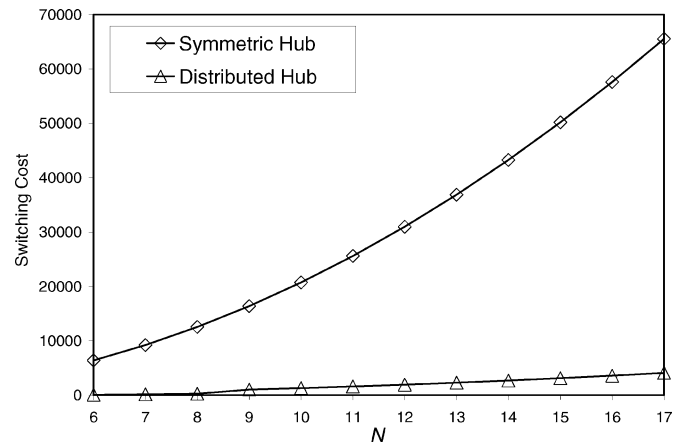


Fig. 4. Switching cost with symmetric hub architecture and distributed hub architecture for ring with $r = 1$ and $g = 16$.

bound. Note that both the symmetric hub architecture and the distributed hub use nearly the same number of ADMs for the range of values shown. Both approaches meet the lower bound from Proposition 1, for specific values of N (e.g., $N = 17$). Also, note that the curves for the symmetric and distributed hub cases both sharply increase when $N = 18$. This is the value of N , where the number of wavelengths generated by each node increases from 1 to 2. The savings relative to the bound on a network with no switching is greatest just before these jumps, i.e., when each node is generating enough circuits to fill up an integer number of lightpaths. The largest gain shown is for the case when $N = 17$, where using switching reduces the number of ADMs by more than 40%. For certain values of N , the symmetric hub architecture uses more ADMs than the lower bound from [1]. This is partly due to the fact the lower bound in [1] is not always obtainable; also, for these cases, we can further reduce the needed number of ADMs in the symmetric hub architecture by multiplexing the traffic from several nonhub nodes onto a single wavelength. Fig. 4 compares the switching cost between the symmetric hub architecture and the distributed hub architecture for the same ring. Though these two approaches required a similar number of ADMs, the distributed hub architecture has a significantly smaller switching cost. As N increases,

the difference in switching requirements increases; for $N = 17$ the distributed hub architecture has a switching cost that is more than 97% less than the symmetric hub case. We have observed similar trends for other parameter settings. In general, the benefits of switching appear to be greater for larger values of g . For example, when $g = 4$ and $r = 1$, over the same range of N , switching reduces the needed number of ADMs by at most 20% compared with the lower bound in [1].

IV. BOUND ON THE AVERAGE AMOUNT OF SWITCHING NEEDED

So far, we have considered lower bounds on the number of ADMs needed, and we have argued that if these bounds can be achieved then it will require that all circuits be switched at most once. Based on this, we developed algorithms that can significantly reduce the required number of ADMs. However, in general, there is no assurance that the bounds in Section II-A are tight; in which case it might be possible to further reduce the ADM requirement by allowing more switching in the network. In this section, we provide some insight into this situation by developing an upper bound on the amount of switching needed in a network, in terms of the number of ports in the network. Using this bound, we show that reducing the number of ports inherently requires that traffic be only switched a small number of times.

Again, we consider a WDM network with N nodes, where g low-rate traffic streams are multiplexed onto each wavelength. Let T be the total number of ports in the network, where each lightpath is terminated by two ports (i.e., there are $T/2$ ADMs). Let C be the total number of low-rate (unidirectional) circuits in the network. Now, $f = T/C$ is the average number of ports per circuit. Note that by setting up a point-to-point lightpath for each circuit, f can always be made equal to 2. Of course, more efficient grooming algorithms would lead to f being less than 2. For $i = 1, \dots, C$, assume the i th circuit uses L_i ports, i.e., this circuit is sent over $L_i - 1$ lightpaths. Define

$$K = \frac{1}{C} \sum_i L_i$$

so that K denotes the average number of ports used by a circuit. Finally, let S be the average number of times that a circuit is switched. We want to show that in a network architecture that minimizes the overall number of needed ports, T , each circuit needs to be switched on average only a small number of times. Note that for a given C , minimizing T is equivalent to minimizing f . Also, note that the number of times a circuit is switched is upper-bounded by $L_i/2 - 1$. (This is an upper bound, since the wavelength a circuit is on may be dropped at an intermediate node only to add/drop another circuit sharing that wavelength, but not switched.) Thus, S is upper bounded by $K/2 - 1$. Since each port is shared by at most g circuits, we have that $CK/g \leq T$, which implies that $K \leq fg$. Hence

$$S \leq \frac{fg}{2} - 1. \quad (17)$$

Thus, for a given topology and traffic demand, any upper bound on f (or equivalently T) can be converted into an upper bound on S . This suggests that a topology that is efficient in the use of




Fig. 5. Plot of the bound on S in (12) when f is bounded using the symmetric hub grooming algorithm.

ports (small f) will not use much switching (small S). Moreover, the most port efficient topology will yield the tightest bound on S in (17). We consider some specific examples of this bound next.

For any topology and traffic demand, as noted earlier, $f \leq 2$, substituting this into (17), we have

$$S \leq g - 1. \quad (18)$$

When $g = 1$, this implies that no switching is required, as one would expect (since point-to-point circuits are most efficient). For $g > 1$, the above bound is very loose because establishing point-to-point circuits is inefficient in terms of the number of ports. Next, consider a unidirectional ring with uniform traffic of r circuits between all N nodes so that $C = N(N - 1)r$. The number of ports needed with an arbitrary number of switches is upper-bounded by the ports required in a single-hub architecture. The total number of ports for a single-hub architecture is given by $T = 4 \lceil (N - 1)r/g \rceil (N - 1)$. Thus, we have

$$\begin{aligned} S &\leq \frac{4 \lceil (N - 1)r/g \rceil (N - 1)g}{2N(N - 1)r} - 1 \\ &= 2 \left\lceil \frac{(N - 1)r}{g} \right\rceil \frac{g}{Nr} - 1 \\ &\leq 2 \left(\frac{N - 1}{N} \right) + 2 \frac{g}{Nr} - 1. \end{aligned} \quad (19)$$

We emphasize that while (19) was developed by considering a single-hub architecture, the bound applies to an architecture with an arbitrary number of hubs and general topology. Notice that for any fixed g/r as N gets large, the upper bound on the average amount of switching in (19) approaches 1. Also, notice that when $g/r < 1$, then the right-hand side of (19) is less than one for any N . In other words, when each node generates more than a wavelength of traffic for each other node, the average amount of switching per circuit in an architecture that efficiently uses ports will be less than one. Of course, for the hub architecture S is less than one by design. However, the above tells us that any architecture that sought to further reduce the number of ADMs would not require more switching than the bound on S given in (19). Furthermore, instead of using a single-hub architecture to bound the number of ports needed, a better bound on the number of ports can be found by using a more efficient architecture. As an example, we consider the symmetric multihub

architecture from Section II-B. In this case, the number of ports can be found using (13). Fig. 5 shows the resulting bound on S as a function of N , for a ring with $g = 4$, $r = 1$. In finding this bound, the optimal number of hubs were chosen for each N . Notice that for all but three values of N this bound is less than one, suggesting that each circuit needs to be switched at most once.

V. CONCLUSION

We have shown that providing limited switching ability can aid in reducing the number of ADMs needed in a SONET/WDM ring network. We have considered providing this switching ability in two types of architectures—multiple-hub architectures and distributed hub architectures. In both cases, significant savings in ADMs are possible. We introduced a notion of the switching cost needed in a ring and showed that the distributed hub architecture incurred a significantly smaller switching cost than a multiple-hub architecture. We also presented an upper bound on the amount of switching needed in a network; this bound suggests that in a network that efficiently utilizes ADMs, circuits need to be switched at most once, as they are in the architectures presented here.

In addition to reducing the number of ADMs, other advantages of switching include the ability to better support dynamic traffic and to improve a network's robustness to node failures. In addressing such issues, the placement of switches within a ring will likely be an important consideration.

REFERENCES

- [1] E. Modiano and A. Chiu, "Traffic grooming algorithms for minimizing electronic multiplexing costs in unidirectional SONET/WDM ring networks," in *Proc. CISS*, Princeton, NJ, Feb. 1998.
- [2] J. Simmons and A. Saleh, "Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic," *J. Lightwave Technol.*, vol. 17, pp. 48–57, Jan. 1999.
- [3] R. Berry and E. Modiano, "Reducing electronic multiplexing costs in SONET/WDM rings with dynamically changing traffic," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1961–1971, Oct. 2000.
- [4] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost effective traffic grooming in WDM rings," in *Proc. IEEE INFOCOM*, San Francisco, CA, Apr. 1998, pp. 69–77.
- [5] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET network design," in *Proc. IEEE INFOCOM*, New York, Mar. 1999, pp. 734–743.
- [6] —, "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] E. Modiano and P. Lin, "Traffic grooming in WDM networks," *IEEE Commun. Mag.*, vol. 39, pp. 124–129, July 2001.
- [8] H. Ghafouri-Shiraz, G. Zhu, and Y. Fei, "Effective wavelength assignment algorithms for optimizing design costs in SONET/WDM rings," *J. Lightwave Technol.*, vol. 19, pp. 1427–1439, Oct. 2001.
- [9] J. Wang, W. Cho, V. R. Vemuri, and B. Mukherjee, "Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multihop connections," *J. Lightwave Technol.*, vol. 19, pp. 1645–1653, Nov. 2001.

- [10] X. Zhang and C. Qiao, "An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings," *IEEE/ACM Trans. Networking*, vol. 8, pp. 608–617, Oct. 2000.
- [11] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 122–133, Jan. 2002.
- [12] R. Dutta and G. Rouskas, "On optimal traffic grooming in WDM rings," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 110–121, Jan. 2002.
- [13] T. Y. Chow and P. J. Lin, "The ring grooming problem," *Networks*, vol. 44, pp. 194–202, Aug. 2004.
- [14] J. Q. Hu, "Traffic grooming in WDM ring networks: A linear programming solution," *J. Opt. Networks*, vol. 1, pp. 397–408, 2002.
- [15] G. Calinescu, O. Frieder, and P. J. Wan, "Minimizing electronic line terminals for automatic ring protection in general WDM optical networks," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 183–189, Jan. 2002.
- [16] H. Zhu, H. Zang, K. Zhu, and B. Mukherjee, "A novel generic graph model for traffic grooming in heterogeneous WDM mesh networks," *IEEE/ACM Trans. Networking*, vol. 11, pp. 285–299, Apr. 2003.
- [17] J. Q. Hu and B. Leida, "Traffic grooming, routing, and wavelength assignment in optical WDM mesh networks," in *Proc. IEEE INFOCOM*, Hong Kong, Mar. 2004.
- [18] J. Bosák, *Decompositions of Graphs*. Boston, MA: Kluwer, 1990.



Randall A. 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 communication, data networks, and information theory.

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 Postdoctoral Fellow from 1992 to 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, Massachusetts Institute of Technology (MIT) Lincoln Laboratory, Cambridge, 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.