Traffic Grooming Algorithms for Reducing Electronic Multiplexing Costs in WDM Ring Networks

Angela L. Chiu and Eytan H. Modiano

Abstract—**We develop traffic grooming algorithms for unidirectional SONET/WDM ring networks. The objective is to assign calls to wavelengths in a way that minimizes the total cost of electronic equipment [e.g., the number of SONET add/drop multiplexers (ADM's)]. We show that the general traffic grooming problem is NP-complete. However, for some special cases we obtain algorithms that result in a significant reduction in the number of ADM's. When the traffic from all nodes is destined to a single node, and all traffic rates are the same, we obtain a solution that minimizes the number of ADM's. In the more general case of all-to-all uniform traffic we obtain a lower bound on the number of ADM's required, and provide a heuristic algorithm that performs closely to that bound. To account for more realistic traffic scenarios, we also consider distance dependent traffic, where the traffic load between two nodes is inversely proportional to the distance between them, and again provide a nearly optimal heuristic algorithm that results in substantial ADM savings. Finally, we consider the use of a hub node, where traffic can be switched between different wavelength, and obtain an optimal algorithm which minimizes the number of ADM's by efficiently multiplexing and switching the traffic at the hub. Moreover, we show that any solution not using a hub can be transformed into a solution with a hub using fewer or the same number of ADM's.**

Index Terms—**Add/drop multiplexers (ADM's), optical network design, synchronous optical network (SONET), SONET rings, wavelength-division multiplexing (WDM).**

I. Introduction

MUCH of today's physical layer network infrastructure is built around synchronous optical network (SONET) rings. Typically, a SONET ring is constructed using fiber (one or two fiber pairs are typically used in order to provide protection) to connect SONET add drop multiplexers (ADM's). Each SONET ADM has the ability to separate a high rate SONET signal into lower rate components. For example, four OC-3 circuits can be multiplexed together into an OC-12 circuit and 16 OC-3's can be multiplexed into an OC-48. The recent emergence of wavelength-division multiplexing (WDM) technology has resulted in the ability to support multiple SONET rings on a single fiber pair. Consider, for example, the SONET ring network shown in Fig. 1, where each wavelength is used to form an OC-48 SONET ring. This network is used to provide OC-3 circuits between nodes and SONET ADM's are used to combine up to 16 OC-3 circuits into a single OC-48 that is carried on a wavelength. With WDM technology providing dozens of wavelengths on a fiber, as many OC-48 rings can be supported per fiber pair instead of just one. This tremendous increase in network capacity, of course, comes at the expense of needing additional electronic multiplexing equipment. With the emergence of WDM technology, the dominant cost component in networks is no longer the cost of fiber but rather the cost of electronics.

The SONET/WDM architecture shown in Fig. 1 is potentially wasteful of SONET ADM's because every wavelength (ring) requires a SONET ADM at every node. An alternative architecture, shown in Fig. 2, makes use of WDM add drop multiplexers (WADM's) to reduce the number of required SONET ADM's. A WADM at a given node is capable of dropping and adding any number of wavelengths at that node. In order for a node to transmit or receive traffic on a wavelength, the wavelength must be added or dropped at that node and a SONET ADM must be used. Therefore, with a single WADM at each node it is no longer necessary to have a SONET ADM for every wavelength at every node, but rather only for those wavelengths that are used at that node. Therefore, in order to limit the number of SONET ADM's used, it is better to groom traffic in such a way that all of the traffic, to and from a node, is carried on the minimum number of wavelengths. Notice that this is not the same as minimizing the total number of wavelengths used, a problem that has received much attention recently [1].

Recent work on wavelength assignment in WDM networks considered how to assign wavelength to calls, so that the total number of wavelengths required is minimized [1]. The underlying assumption was that calls required a full wavelength. In practice this is rarely the case. Typically, calls require a small fraction of a wavelength and network providers use electronic multiplexing to allow many users to share a wavelength. For example, SONET multiplexers can be used to aggregate as many as 16 OC-3 circuits onto a single OC-48, which in turn can be carried on a single wavelength.

A large part of the cost in providing network services is in the size and complexity of electronic multiplexing equipment. There is a tradeoff between efficient use of the fiber and the electronic equipment. For example, if wavelengths were not limited, each call can be supported on a dedicated wavelength and no electronic multiplexing would be required. However, in most cases, there is an insufficient number of wavelengths to support all connections with dedicated wavelengths and electronic multiplexing is needed to allow low rate users to share wavelengths. Hence, the network design goal is to minimize overall network

Manuscript received February 11, 1999; revised September 1, 1999. This work was supported by the Defense Advanced Research Project Agency (DARPA) under the Next Generation Internet initiative.

A. L. Chiu was with the Massachusetts Institute of Technology (MIT) Lincoln Laboratory, Cambridge, MA 02139 USA. She is now with the AT&T Laboratories, Middletown, NJ 07748 USA (e-mail alchiu@att.com).

E. H. Modiano was with the Massachusetts Institute of Technology (MIT) Lincoln Laboratory, Cambridge, MA 02139 USA. He is now with the Massachusetts Institute of Technology (MIT) Laboratory for Information and Decision System, Cambridge, MA 02139 USA (e-mail: modiano@mit.edu).

Publisher Item Identifier S 0733-8724(00)00387-X.

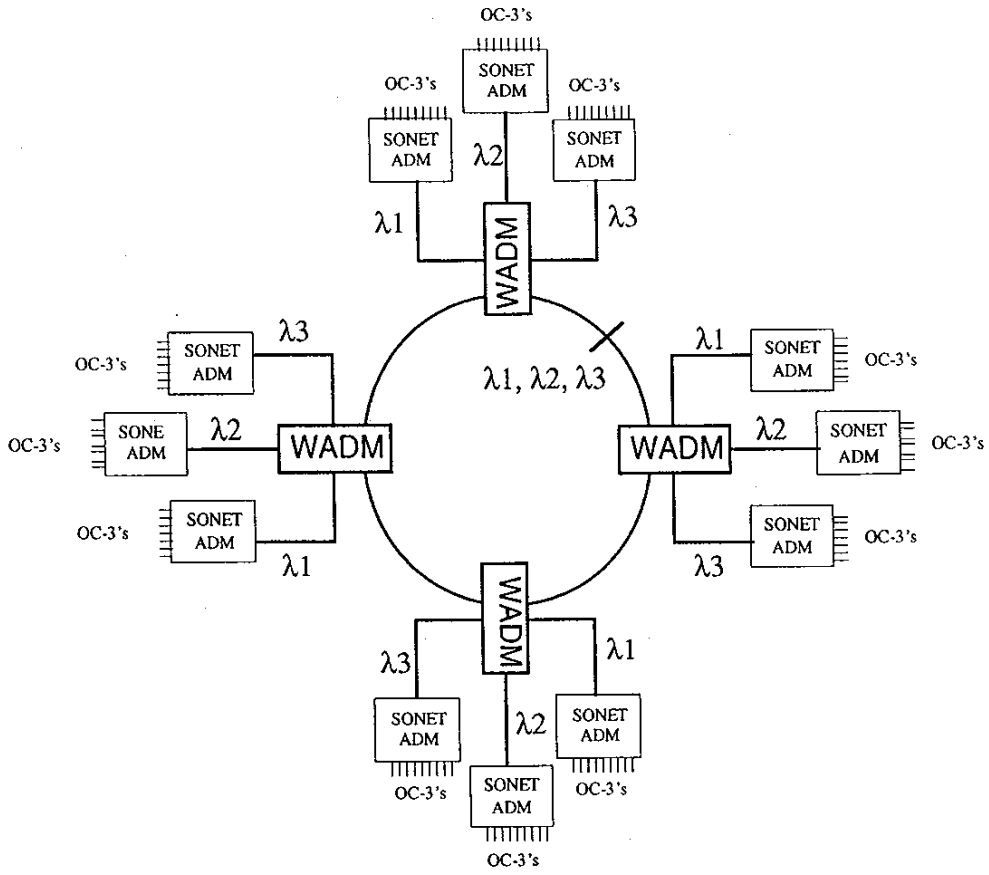

Fig. 1. SONET/WDM rings.

costs not just the number of wavelengths. In this context, the problem is to design traffic grooming algorithms to minimize electronic costs at network edges and to make efficient use of wavelengths.

In this paper we consider a unidirectional WDM ring network with N nodes numbered $1, 2, \ldots, N$ distributed on the ring in the clockwise direction. Each node, i, has one WADM and D_i SONET ADM's. Each SONET ADM is used to aggregate g low rate circuits onto a single high-rate circuit that is carried on a wavelength. For example, each SONET ADM can be used to multiplex 16 OC-3's ($g = 16$) or four OC-12's ($g = 4$) onto a single OC-48. The traffic requirement is for r_{ij} low rate circuits from node i to node j, for any $i \neq j$. With a WADM at a given node, a wavelength can bypass that node if there is no traffic to be received or transmitted from that node, which results in the saving of a SONET ADM. The objective is to minimize the total number of SONET ADM's used in the network to support all of the traffic by intelligently assigning traffic to wavelengths.

As a simple, illustrative example, consider a ring network with four nodes. Suppose that each wavelength is used to support an OC-48 ring, and that the traffic requirement is for eight OC-3 circuits between each pair of nodes. In this case we have $g = 16$ (16 OC-3's in an OC-48) and $r_{ij} = 8$ for all $i \neq j$. In this example we have six node pairs and the total traffic load is equal to 48 OC-3's or equivalently three OC-48 rings. The question is how to assign the traffic to these three OC-48 rings in a way that minimizes the total number of SONET ADM's required. Consider, for example, the following two circuit assignments of traffic: Assignment 1: $\lambda 1$: $1 \leftrightarrow 2$, $3 \leftrightarrow 4$; $\lambda 2$: $1 \leftrightarrow 3$, $2 \leftrightarrow 4$; $\lambda 3$: $1 \leftrightarrow 4$, $2 \leftrightarrow 3$; and Assignment 2: $\lambda 1$: $1 \leftrightarrow 2$, $1 \leftrightarrow 3$; $\lambda 2$: $2 \leftrightarrow 3$, $2 \leftrightarrow 4$; $\lambda 3$: $1 \leftrightarrow 4$, $3 \leftrightarrow 4$. Since $g = 16$ and $r_{ij} = 8$ each wavelength can support all of the traffic between two pairs of nodes. With the first assignment, each node has some traffic on every wavelength. For example wavelength 1 carries the traffic between nodes 1 and 2 and the traffic between nodes 3 and 4. Therefore, each node would require an ADM on every wavelength for a total of 12 ADM's. With the second assignment each wavelength contains traffic from only three nodes and hence only nine ADM's are needed. Notice that both assignments carry the same amount of total traffic (8 OC-3's between each pair of nodes).

Most previous work in this area has focused on the virtual topology design problem for known and fixed (static) traffic patterns [2], [3]. The general problem of virtual topology design can be formulated as a mixed integer programming problem which is known to be difficult. Heuristic algorithms have been developed to design virtual topologies that minimize the number of wavelengths, delays or blocking probabilities.

While the general topology design problem is known to be intractable, the traffic grooming problem is a special instance of the virtual topology design problem for which, in certain circumstances, a solution can be found. For example, [4] considers

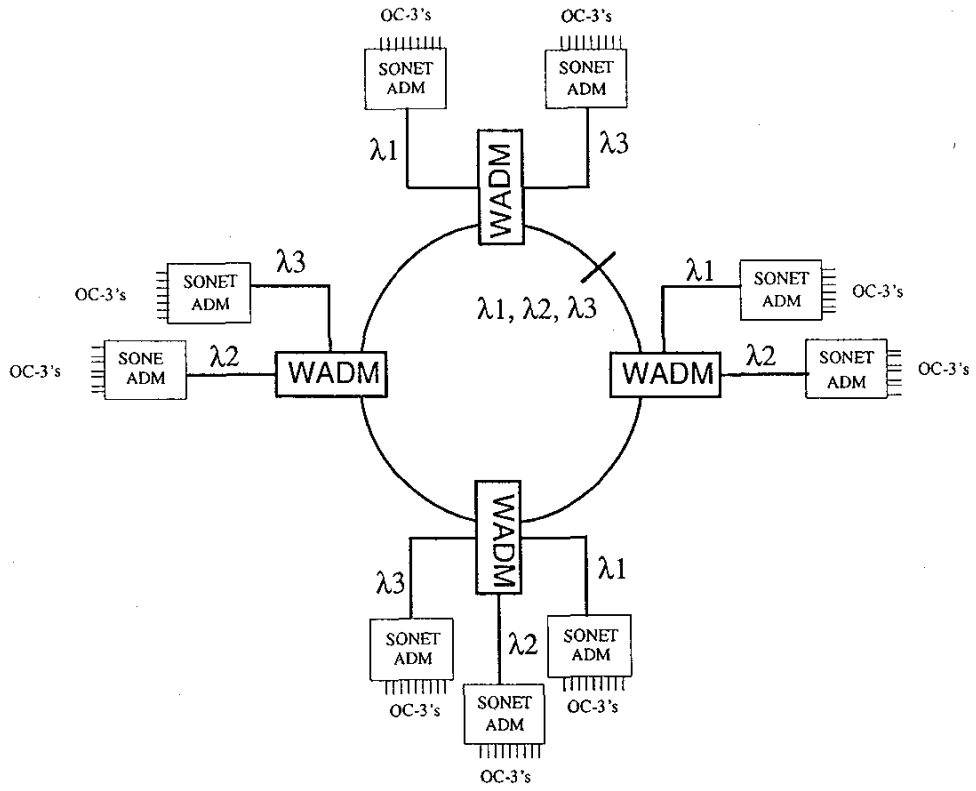

Fig. 2. Using WADM's to reduce the number of SONET ADM's.

traffic grooming for a bidirectional ring with uniform traffic. In this paper we describe solutions for unidirectional rings. In Section II, we consider the simple case of an egress node from which all of the traffic is directed and in Section III, we consider the more general case of all-to-all uniform traffic in a ring network. In Section IV, we extend our traffic model to account for distance dependent traffic and in Section V, we consider the case of a ring network with a hub node, where traffic can be switched between different SONET rings using a SONET cross connect. We summarize the results with remaining issues in Section VI.

II. Egress Node

We start by considering a very simple case of the traffic grooming problem, where all of the traffic on the ring is destined to a single node that we call the egress node. This case is of particular importance in access networks where traffic from the various access nodes on the ring is all destined to the telephone company's central office.[1] In fact, in today's access networks most of the traffic is destined to the central office from where it is routed to more distant locations. Only a small fraction of the traffic travels between nodes on the same access ring. Further, the discussion of the egress node case is also significant because it provides insight to the general traffic grooming problem. For example, we use this case to show that the general traffic grooming problem is NP-complete.

[1]For simplicity of presentation, we discuss the egress node case. However, this discussion also applies to the case of an ingress node where all the traffic comes from one node as well as the case of a single node from and to which all of the traffic is destined.

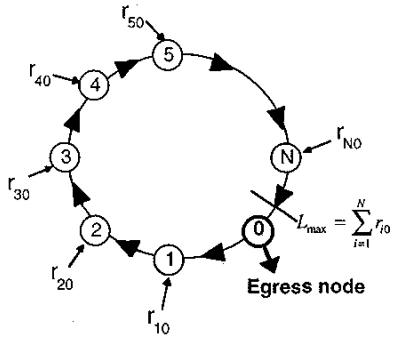

Fig. 3. Unidirectional ring network with an egress node.

Denote the egress node as node 0 and assume that it lies between node N and node 1, as shown in Fig. 3.

The traffic rate between nodes i and j takes on positive values only when $j = 0$ and $i = 1, \ldots, N$. Since the ring is unidirectional, all traffic has to go through the link between node N and node 0. Therefore, link $(N, 0)$ carries the heaviest load, which is given by $L_{\max} = \sum_{i=1}^{N} r_{i0}$.

Hence the minimum number of wavelength required to support this load is, $W_{\min} = \lceil L_{\max}/g \rceil$.

Without loss of generality, we can assume $r_{i0} < g$ for all i for the rest of this section. This is because when $r_{i0} > g$, the portion of the circuits that can fill up a wavelength can be carried on a separate wavelength without the need to be groomed together with traffic from other nodes. That is, $\lceil r_{i0}/g \rceil$ wavelengths can be filled with $\lceil r_{i0}/g \rceil * g$ low rate circuits. The remaining cir-

cuit can be groomed with traffic from other nodes. Hence, the problem is reduced to that of grooming traffic streams of rate $r_{i0} < g$. The objective of the traffic grooming problem is to assign circuits to wavelengths in such a way that the total number of SONET ADM's in the network is minimized.

We start by observing that the general traffic grooming problem, even in this special egress node case, is NP-complete.

Theorem 1: The traffic grooming problem is NP-complete.

Proof: We prove this result by showing that the Bin Packing problem can be transformed into the traffic grooming problem in polynomial time. Since the Bin Packing problem is known to be NP-complete [5], the traffic grooming problem must be NP-complete as well. First we show that there exists an optimal solution such that no traffic from a node is split onto two SONET rings. We prove this claim by showing that for any solution with split traffic, there exist a corresponding solution without split traffic that uses the same number or fewer ADM's. Consider a circuit assignment that has traffic from a node split onto two or more SONET rings. Clearly, each node with traffic on two or more rings must have at least two ADM's. Since $r_{i0} < g$, all of the traffic from that node can be accommodated on a separate SONET ring with two ADM's (one at that node and the other at the egress node). Hence, any circuit assignment can be transformed into one where no node has traffic on more than one ring using the same number or fewer ADM's. Thus, there exists a solution that minimizes the number of ADM's without splitting traffic from a node onto more than one ring.

For any optimal solution with no split traffic, only one ADM is needed for each node while one ADM at each SONET ring is needed for the egress node. Since at least one ADM is needed at every node, the problem is reduced to that of minimizing total number of SONET rings used. This can be achieved by combining the traffic from multiple nodes onto a single ring. This, in fact, is the well known Bin Packing problem [5]. Specifically, the wavelengths correspond to bins, where each bin has finite capacity g; the r_{i0}'s correspond to the sizes of the items to be packed into the bins; and the goal is to minimize the number of bins used. Hence, any solution to the traffic grooming problem can be used to obtain a solution to the bin packing problem in polynomial time. Since it is known that for general g and r_{i0}'s, the bin packing problem is NP-complete, we have shown that the traffic grooming problem must also be NP-complete. Notice, also, that by showing that the traffic grooming problem in the egress node case in NP-complete we have also shown that the general traffic grooming problem with all-to-all traffic is NP-complete. However, in the special case where all of the r_{i0} are equal (i.e., $r_{i0} = r$ for all i) an exact solution for the minimum number of ADM's required and how they should be assigned to circuits can be found.

A. Special Case: $r_{i0} = r$

Using the proof of Theorem 1 we know that there exists an optimal solution that does not require traffic from a node to be split onto multiple wavelengths. Since, without splitting traffic, we can groom the traffic from at most $\lfloor g/r \rfloor$ nodes on one SONET ring, the number of SONET rings needed is, $W = \lceil N/\lfloor g/r \rfloor \rceil$.

Hence, the minimum number of SONET ADM's $M_{\min} = N + W$, because one ADM is needed at every node (since exactly one wavelength is dropped at each node) plus at the egress node one ADM is needed for each wavelength (since all wavelengths are dropped at the egress node). Note that the resulting number of SONET rings may be larger than the minimum number of rings required (taking the case $N = 4$, $g = 7$, and $r = 5$ as an example). Next we provide a solution that uses the minimum number of ADM's required subject to using the minimum number of SONET rings (or wavelengths).

B. Minimizing ADM's Subject to the Minimum Number of Wavelengths

Here we require that the solution will use the minimum number of wavelengths, $W_{\min}$ and find an algorithm that minimizes the number of ADM's subject to this constraint. Hence, the total number of ADM's required will include $W_{\min}$ ADM's at the egress node (one for each ring), plus the total number of ADM's at all the regular nodes. Since we are now restricted to using the minimum number of wavelengths, traffic from a node may have to be split onto multiple rings and each node will have one ADM for each SONET ring used to carry its traffic. We say that a split occurs when some traffic from a node is divided onto two rings. For example, if traffic from a node is divided onto three rings two splits have occurred. Clearly, for each traffic split, a new ADM is needed. Each node needs one ADM plus an additional ADM for each traffic split at that node. Hence the total number of ADM's needed is equal to $W_{\min} + N + S$ where S is the total number of traffic splits over all nodes. Therefore, the minimum number of SONET ADM's subject to minimum number of SONET rings is achieved by minimizing the total number of traffic splits. Obviously, if all of the traffic can fit on the $W_{\min}$ rings with no need for traffic splitting then we have the minimum ADM solution. For each ring with no split traffic, the maximum link load is $L_{ns} = \lfloor g/r \rfloor * r$.

Let W_{ns} be the maximum number of rings containing no split traffic, with L_{ns} circuits each. Since, the remaining $(W_{\min} - W_{ns})$ rings contain at most g circuits, we have

$$W_{ns} * L_{ns} + (W_{\min} - W_{ns}) * g \geq L_{\max} = r * N,$$

where $W_{\min} = \lceil r * N/g \rceil$. Therefore, the maximum number of rings with no split traffic is given by,

$$W_{ns} = \min \left\{ W_{\min}, \left\lfloor \frac{g * W_{\min} - L_{\max}}{g - L_{ns}} \right\rfloor \right\}.$$

If $W_{ns} = W_{\min}$ then all of the traffic can be accommodated without any need to split traffic and the optimal solution is found. Hence, in the following we focus on the case where $W_{ns} < W_{\min}$. In this case not all of the traffic can be accommodated without the need for traffic splitting which implies that there exists at least one traffic-split. The algorithm below assigns circuits to wavelengths in a way that minimizes the number of traffic splits and hence the number of ADM's. The algorithm works for arbitrary positive integer values of g and r and is not restricted to the case of $r < g$. The algorithm is iterative with the following three steps.[2]

[2]The algorithm and its optimality proof were provided by Z. Tang of AT&T Labs, 200 Laurel Ave., Middletown, NJ 07748 USA.

1) Algorithm:

Step 1:
Fill each of $W = W_{\min}$ rings with the unsplit traffic from $\lfloor g/r \rfloor$ nodes. The remaining capacity for each of the W rings is $g_1 = g - \lfloor g/r \rfloor r < r$ and the traffic from $N_1 = N - \lfloor g/r \rfloor W$ nodes still needs to be assigned. Notice that N_1 is less than W, since the remaining unassigned traffic, $N_1 r$, needs to fit on the total remaining capacity, $g_1 W$.

Step 2:
Fill the remaining capacity g_1 of each of N_1 rings by the traffic from each of the remaining N_1 nodes. The remaining traffic of each of N_1 nodes becomes $r_1 = r - g_1$.

Step 3:
Now, there are $W_1 = W - N_1$ rings that each has capacity g_1 left, and N_1 nodes that each has traffic r_1 left. Update $W := W_1$, $g := g_1$, $N := N_1$, and $r := r_1$ and repeat Steps 1-3 until the traffic from all nodes has been assigned (i.e., $r_1 = 0$).

2) Example—$N = 5$, $r = 5$, $g = 9$, and $W = 3$: In step 1, we fill each of three rings with $\lfloor 9/5 \rfloor = 1$ node's traffic. A capacity of four remains for each ring and the traffic of two nodes has not yet been assigned. In step 2, we fill up all remaining capacity of two rings by the traffic from each of two nodes left over. There is only one ring left with capacity of four and a single circuit from each of two nodes remains to be assigned. In step 3, we assign those two circuits to the ring with remaining capacity. The total number of traffic-splits resulted is two.

In order to prove the optimality of the above algorithm, we start by showing that there always exists an optimal solution that uses steps 1 and 2 of the algorithm. First, we show that there always exists an optimal solution such that each of W rings is assigned all of the traffic from $\lfloor g/r \rfloor$ nodes without splitting (step 1). To see this consider an optimal solution such that there is one ring, say ring 1, whose unsplit traffic is from fewer than $\lfloor g/r \rfloor$ nodes. By assumption, there exists at least one node whose traffic is split, say node 1. We can swap node 1's traffic assigned to rings other than ring 1 with some split traffic assigned to ring 1 such that all node 1's traffic is on ring 1. Clearly, such swapping will not increase the number of traffic-splits and hence maintain the optimality of the traffic assignment. We can repeat this process until each ring is assigned unsplit traffic from $\lfloor g/r \rfloor$ nodes. Notice that $\lfloor g/r \rfloor$ is also the maximum number of nodes whose traffic can be assigned to a ring without splitting.

Next, we also show that there exists an optimal solution such that all traffic from $\lfloor g/r \rfloor$ nodes is assigned without splitting to each ring, and each of the remaining $N_1 = N - \lfloor g/r \rfloor W$ nodes have their traffic assigned to the remaining capacity of a ring in a one node to one ring fashion with $N_1 < W$. That is, the remaining capacity of each of $N_1 = N - \lfloor g/r \rfloor W$ rings is filled with traffic from exactly one node (step 2). To see this consider an optimal solution where each of the W rings takes all traffic from $\lfloor g/r \rfloor$ nodes (without splitting the traffic) and one ring has its remaining capacity, $g_1 = g - \lfloor g/r \rfloor r < r$, filled with split traffic from more than one node, say nodes 1, 2. We can swap node 2's traffic assigned to this ring with node 1's traffic assigned to other rings, and repeat this process for cases with more than two nodes, until all the split traffic of this ring is from node 1 only. Clearly, this swapping will not increase the number of traffic splits and hence maintain the optimality of the traffic assignment. This process can be repeated until the all the remaining capacity of $N_1 = N - \lfloor g/r \rfloor W$ rings is filled by the traffic from each of the $N_1 = N - \lfloor g/r \rfloor W$ nodes.

Since the algorithm repeatedly uses steps 1 and 2 until all traffic has been assigned, it results in the minimum number of traffic splits. This is because after steps 1 and 2, we are left with $N_1 = N - \lfloor g/r \rfloor W$ nodes each with $r_1 = r - g_1$ traffic to be assigned and $W_1 = W - N_1$ rings each with remaining capacity $g_1 = g - \lfloor g/r \rfloor r$. Assigning the traffic from these N_1 nodes to the W_1 wavelength using steps 1 and 2 of the algorithm will, again, result in the minimum number of traffic splits. Clearly, repeating this process until all traffic has been assigned will result in an optimal solution.

Next, we consider the more general case of a ring network with traffic between all node pairs.

III. All-to-All Uniform Traffic

In this section, we consider the more general case of all-to-all traffic in the ring. Since the solution to the general problem is NP-complete, we consider a more limited case of uniform traffic. That is, $r_{ij} = r$ for all $i \neq j$, where r is some positive integer representing the number of low rate circuits between each pair of nodes. Again, the traffic granularity, g, is equal to the number of low rate circuits that can fit on a single SONET ring (or wavelength). An interesting observation that significantly simplifies the solution for unidirectional rings is that the routing problem is eliminated. All circuits travel in the same direction, and as long as traffic is symmetric (i.e., $r_{ij} = r_{ji}$ for all i, j), all links carry an equal load. This is in contrast to a bidirectional ring where link loads depend on how calls are routed.

We begin with a few definitions that will help our discussion. Let the node load be the number of low rate circuits originating or terminating at a node, then $L_d = (N-1)r$. Let the link load be the number of low rate circuits traversing a link. Clearly, $L = N(N-1)r/2$, because there are $N(N-1)/2$ node pairs each with r circuits between each pair. Now, a lower bound on the number of ADM's needed is given by $M \geq \lceil L_d/g \rceil N$. This number is simply the minimum number of wavelengths required to carry the traffic to and from a node multiplied by the number of nodes (since each wavelength dropped at a node requires an ADM). A tighter lower bound is provided in the next section. The minimum number of wavelengths required to carry all of the traffic in the network is equal to the link load divided by g, i.e., $W_{\min} = \lceil L/g \rceil$. This minimum can be achieved by dropping every wavelengths at every node and would require $W_{\min}$ ADM's at each node yielding an upper bound on the minimum number of ADM's, hence, $M_{\min} \leq W_{\min} * N$.

Similarly, an upper bound on the number of wavelengths used is obtained by using dedicated wavelengths between each pair of nodes. With this approach, $\lceil r/g \rceil$ wavelengths will be needed for each pair of nodes and the total number of wavelengths would be $W_{\max} = \lceil r/g \rceil N(N-1)/2$, and the corresponding number of ADM's, $M = 2W_{\max}$, which is another upperbound on $M_{\min}$. Notice that when $g = 1$ both the upper and lower bounds on the number of ADM's are equal to $rN(N-1)$ and, therefore the solution is optimal.

A. Minimizing the Number of Wavelengths and ADM's Simultaneously

In this section, we show that it is not always possible to minimize both the number of wavelengths and the number of ADM's. Although the cost of electronics dominates network costs, in order to minimize network costs, one would like to find a solution that minimizes both the number of wavelengths used and the number of SONET ADM's. The example below shows that this may not always be possible. Hence, for the remainder of this paper we focus only on minimizing the number of ADM's.

1) Example: In this example, a solution that uses both the minimum number of wavelengths and the minimum number of ADM's does not exist.

This example is of a ring with five nodes ($N = 5$), three circuits between every pair of nodes ($r = 3$) and granularity of four ($g = 4$) (i.e., OC-12's on an OC-48 ring). Note that with this example the link load $L = 30$ and the minimum number of wavelengths $W_{\min} = 8$.

First we show that there is a solution that uses only 20 ADM's and then we show that any solution that uses only eight wavelengths has to use more than 20 ADM's. With five nodes there are ten node pairs and the traffic between each pair can be supported on a single wavelength requiring two ADM's per wavelength or a total of 20 ADM's. Now, any solution using only eight wavelength would have at least six of them full (contain four circuits). In order to fill a wavelength with four circuits at least three nodes would have to be on that wavelengths. Therefore, each of the six full wavelength requires three ADM's. In addition, the remaining two wavelengths must have at least two nodes and, thus, two ADM's on each wavelength. Hence, the minimum number of ADM's needed for an eight wavelength solution is 22. So we have shown that a minimum ADM solution that uses the minimum number of wavelength does not always exist. However, in many cases a solution using both the minimum number of wavelengths and ADM's can be found. In particular, we have the following conjecture:

2) Conjecture: For $r = 1$ (i.e., one circuit per source/destination pair) and uniform traffic, the minimum number of ADM's can be achieved with the minimum number of wavelengths.

While we are unable to prove this conjecture it appears to hold for all of the cases that we studied. For the remainder of this section we focus on the special case of $r = 1$. We begin with the derivation of a lower bound on the number of ADM's.

B. Lower Bound on Number of ADM's ($r = 1$)

We obtain a lower bound on the number of ADM's by finding the most efficient ways to carry traffic between nodes on the same wavelength. That is, we determine the maximum average number of circuits that can be supported by an ADM, and use that number to lower bound the number of ADM's required in the network. In this section we restrict our discussion to the case of $r = 1$, however our approach can be generalized to other values of r (for example, see Section IV on distance dependent traffic). For a given wavelength, with n nodes (and n ADM's), we classify the traffic into two classes. In the first class, which we call "all-to-all traffic" a circuit is set-up between every pair of nodes. With n nodes on the wavelength, the total number of circuits is $n(n-1)/2$ using n ADM's. Since at most g circuits can be supported on a wavelength, $n(n-1)/2$, must be less than or equal to g. In the second class, which we call "cross traffic," the nodes on the wavelength are divided into two groups of size n_1 and n_2 where $n_1 + n_2 = n$, and a circuit is set-up between every node in one group and every node in the other group. For "cross traffic" the link load is $n_1 * n_2$, and again this load must be less than or equal to g. For a given value of n, the link load is maximized when $n_1 = \lfloor n/2 \rfloor$. Note that with all-to-all traffic among a group of nodes all of the circuits between members of those groups are established. While with cross traffic, only those circuits between members of the two groups are established but circuits between the nodes within each individual group remain unassigned.

Also notice that for a given number of ADM's, "all-to-all traffic" assignments can carry more circuits than "cross traffic" assignments. This is because with cross traffic on average approximately $n/4$ circuits are supported per ADM while with all-to-all traffic on average $(n-1)/2$ circuits are supported per ADM. For particular values of g, this concept can be used to generate a lower bound on the number of ADM's. For example, we demonstrate this approach for obtaining a lower bound when $g = 4$ and $g = 16$.

1) Example—$g = 4$ (e.g., OC-12 Circuits on an OC-48 Ring): It can be shown that with $g = 4$ the most efficient circuit assignment requires 1 ADM per circuit. There are three ways in which circuits can be assigned to wavelengths requiring one ADM per circuit:

1) three nodes per wavelength with all-to-all traffic among the nodes, for a link load of three using three ADM's;
2) four nodes with cross traffic between pairs of nodes, for a link load of four using four ADM's;
3) four nodes with all-to-all traffic among three nodes and cross traffic between the fourth node and one of the three nodes, for a link load of four using four ADM's.

In all three cases, we can support one circuit per ADM on average. Notice, that due to the maximum link load of four circuits per wavelength, many assignments are not possible. For example, all-to-all traffic among four nodes results in a link load of six which cannot be supported on a single wavelength. It can be easily determined that many other assignments can be more efficient. Hence, no matter how circuits are assigned, we need at

least one ADM per circuit, leading to the following lower bound on the number of ADM's:

$$LB(g=4) = (\text{total link load } L \text{ circuits})/(1 \text{ ADM/circuit}) = N(N-1)/2.$$

2) Example—$g = 16$ (e.g., OC-3 Circuits on an OC-48 Ring): It can be shown that the most efficient way to assign nodes and circuits to a ring is six nodes with all-to-all traffic between them. This results in a link load of 15 circuits using six ADM's and a utilization of 2.5 circuits per ADM. Alternatively, eight nodes can be assigned to a ring with cross traffic between two groups of four nodes resulting in a link load of 16 circuits using eight ADM's or an efficiency of two circuits per ADM. Since the maximum efficiency that can be achieved is 2.5 circuits per ADM, an immediate lower bound on the number of ADM's is, $LB(g = 16) = \lceil L/2.5 \rceil$ where L is the total link load and is equal to $N(N-1)/2$. This bound can be made slightly tighter by noticing that there is only a limited number of wavelengths that can be filled with all-to-all traffic and that some wavelengths will have to carry cross traffic. For brevity, the details of the tighter bound are omitted. Similarly, this approach can be extended to obtain lower bounds for other values of g. Next we discuss heuristic algorithms that attempt to assign circuits to rings in order to minimize the number of ADM's required.

C. First Heuristic Algorithm

This algorithm attempts to maximize the number of nodes that only require one ADM, then of the remaining nodes maximize the number of nodes with two ADM's and so on. A node needs k ADM's if it is on k wavelengths. Let, M_k be the number of nodes with k ADM's ($k = 1$ to $W_{\min}$). Then, the algorithms maximizes M_1, then maximizes M_2, ..., maximizes $M_{w\ \min -1}$. Clearly, the motivation of the algorithm is that by maximizing the number of nodes that use fewer ADM's we ultimately reduce the total number of ADM's used. It can be shown that M_i, $i = 1, 2, \ldots, W_{\min}$, is given by

$$M_i = \max\left\{H \text{ s.t. } \sum_{h=0 \text{ to } H} (N-1-h) \le i * g\right\} - \sum_{k=1 \text{ to } i-1} M_k.$$

The algorithm fills each wavelength before assigning traffic to a new wavelength, hence it always uses the minimum number of wavelengths $W_{\min}$ and is optimal for $W_{\min} \le 2$. For cases where $W_{\min} > 2$, the algorithm is not necessarily optimal. This is because by maximizing the number of nodes with only a single ADM, the algorithm forces all other nodes to use their ADM's inefficiently. However, the algorithm results in substantial savings over a system where all wavelengths are dropped at all nodes as would be the case if no WADM's were used. The next algorithm, however, results in much more substantial savings in ADM's.

D. Second Heuristic Algorithm

This algorithm attempts to assign nodes to wavelength by efficiently packing the wavelengths. The algorithm is as follows:

Let $n = \lfloor \sqrt{g} \rfloor$ and divide N into $G = \lceil N/n \rceil$ groups of n nodes, where the last group has only $n_l = (N \bmod n)$ nodes. We assign different pairs of groups to each wavelength with cross traffic between the two groups. By design, the cross traffic between two groups of size $n = \lfloor \sqrt{g} \rfloor$ is less than g circuits and can fit on a wavelength. In order to accommodate all of the cross traffic between the G groups a total of $G(G-1)/2$ wavelengths are needed. The remaining traffic is the all-to-all traffic within each group and is fit on the existing wavelengths if possible, otherwise on additional wavelengths. We illustrate the idea with the following two examples.

1) Example—$g = 4$ (OC-12's on an OC-48 Ring): Since $g = 4$ we divide the N nodes into groups of 2 and have the following two cases.

a) N even => $G = N/2$: $G(G-1)/2$ wavelengths can be filled with cross traffic between different pairs of groups. The all-to-all traffic would require additional $\lceil G/4 \rceil = \lceil N/8 \rceil$ wavelengths with four groups on each wavelength. Hence, each node requires $G = N/2$ ADM's for a total of $N^2/2$ ADM's.

b) N odd => $G = (N+1)/2$: The first $G-1 = (N-1)/2$ groups have two nodes and the last group has only one node. $(G-1)(G-2)/2$ wavelengths can be filled with cross traffic between different pair of groups from the first $G-1$ groups. An additional $\lceil (G-1)/2 \rceil$ wavelengths can be used for cross-traffic with the node from the last group, where each wavelength has two groups (four nodes) from the first $(G-1)$ groups and the node from the last group. If one of the wavelength in the previous step is not full [i.e., $(G-1)/2$ is not an integer], it can be used for the all-to-all traffic within two of the first $G-1$ groups. The remaining all-to-all traffic can be handled by assigning four groups to each wavelength. So the number of ADM's at each node is $G = (N+1)/2$ except for the last node which uses $\lceil (G-1)/2 \rceil = \lceil (N-1)/4 \rceil$ ADM's. Hence, the total number of ADM's used when N is odd equals

$$(N-1)(N+1)/2 + \lceil (N-1)/4 \rceil = (N^2-1)/2 + \lceil (N-1)/4 \rceil.$$

Putting it all together the total number of ADM's required with $g = 4$ is

$$\text{ADM}(g=4) = (N-1 \bmod 2) * N^2/2 + (N \bmod 2) * ((N^2-1)/2 + \lceil (N-1)/4 \rceil).$$

In both cases, since all the wavelengths except the last one are filled with four circuits, the resulting assignment only uses $W_{\min}$ wavelengths. However, for general g, the algorithm may result in number of wavelengths that is slightly larger than $W_{\min}$.

2) Example—$g = 16$ (OC-3's on an OC-48 Ring): Since $g = 16$ we divide the N nodes into groups of four. So we have G= $\lceil N/4 \rceil$ groups where $G_0 = \lceil N/4 \rceil$ of the groups have

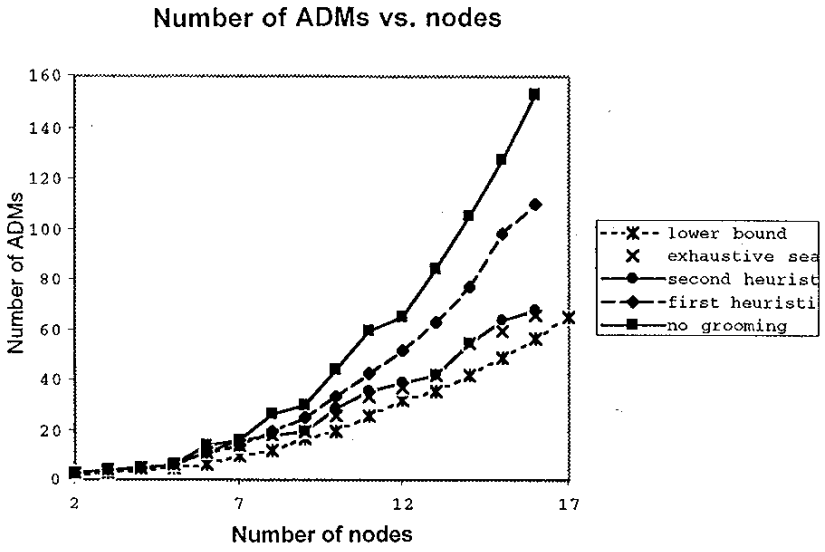

Fig. 4. Comparison of heuristic algorithms ($g = 16$, $r = 1$).

$n = 4$ nodes each and $(G - G_0)$ group has $n_1 = (N \bmod 4)$ nodes. We have the following three cases.

a) $n_1 = (N \bmod 4) = 0 \Rightarrow G = G_0 = N/4$: In this case, $G(G-1)/2$ wavelengths can be filled with cross traffic between different pair of groups. The all-to-all traffic within the groups can be handled by assigning two groups to each wavelength with all-to-all traffic within each group. Hence, each node has exactly G ADM's for a total of $NG = N^2/4$ ADM's.

b) $n_1 = (N \bmod 4) = 1$ *or* $2 \Rightarrow G = \lceil N/4 \rceil$, $G_0 = \lfloor N/4 \rfloor$: In this case, $G_0(G_0-1)/2$ wavelengths can be filled with cross traffic between different pair of groups from the first G_0 groups. An additional G_0 wavelength can be used to handle the cross traffic between each of the first G_0 groups and the last group. Those wavelengths can also be used to carry the all-to-all traffic within each group. This results in link load of $(4 * 3/2 + 2 * 1/2 + 4 * 2 = 15)$ if $n_1 = 2$ or $(4 * 3/2 + 4 * 1 = 10)$ if $n_1 = 1$. So each node uses exactly G_0 ADM's for a total $NG_0 = N\lfloor N/4 \rfloor$ ADM's.

c) $n_1 = (N \bmod 4) = 3 \Rightarrow G = \lceil N/4 \rceil$, $G_0 = \lfloor N/4 \rfloor$: It can be shown, in a manner similar to that of the first two cases, that the total number of ADM's needed in this case is equal to $NG_0 + N - 2 - \lfloor N/4 \rfloor = N\lfloor N/4 \rfloor + N - 2 - \lfloor N/4 \rfloor$. Putting the three cases together we have

$$\mathrm{ADM}(g = 16) = \max\{N,\ N\lfloor N/4 \rfloor + (N - 2 - \lfloor N/4 \rfloor) * \lfloor (N \bmod 4)/3 \rfloor\}.$$

E. Performance Comparison

In Fig. 4 we plot the number of ADM's versus the number of nodes on the WDM ring for $g = 16$ (OC-3 circuits on an OC-48 ring). Plotted in the figure are the lower bound, the number of ADM's used by the first and second heuristic algorithms, the number of ADM's that would be used if all wavelengths were dropped at every node (no grooming) and lastly, the best solution that we have been able to find via exhaustive search. As one can see from the figure, the result of the second heuristic algorithm are very close to the lower bound and almost mirror the best solution. In Fig. 5, we plot the percentage of ADM savings that can be achieved using the second heuristic algorithm over dropping all of the wavelengths at every node. As one can see from the figure, the most savings are achieved when $g = 1$. This, in fact, is a trivial case because each wavelength can only carry the traffic between two nodes and, hence, should only be dropped at those two nodes. It is interesting to note, however, that in general it appears that greater savings can be achieved with smaller values of g. This is due to the fact that when g is small each wavelength can be filled with traffic from just a few nodes while when g is large it takes traffic from many nodes to fill a wavelength.

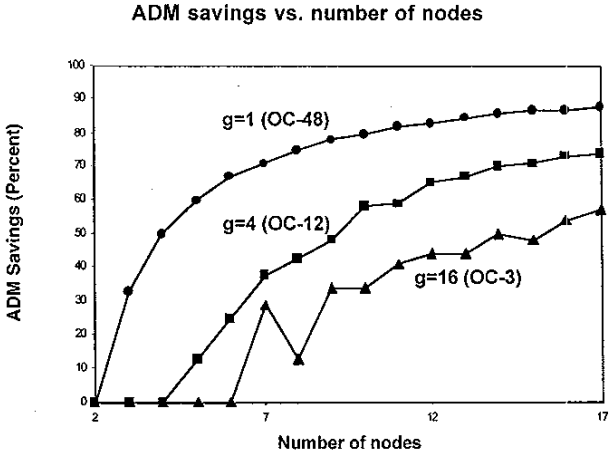

Fig. 5. Percent ADM savings due to grooming.

IV. Distance Dependent Traffic

So far in this paper, we considered only uniform and hubbed traffic models. While, admittedly, those traffic models are often unrealistic, they proved very useful in providing analytically tractable and insightful results. In order to account for more realistic traffic, in this section we consider the distance dependent traffic model from [4]. In this model, the amount of traffic between two nodes is inversely related to the physical distance separating them. So that the amount of traffic between the most distant nodes is one unit of traffic and the traffic demand increases by one unit as the distance between nodes decreases by one link. For nodes arranged in a ring topology, we take the internode distance to be that which would result when using shortest path routing (if the ring were bidirectional).[3] Thus, the traffic demand between nodes i and j would be $T_{i,j} = \lceil (N+1)/2 \rceil - \text{distance}(i,j)\ \forall i \neq j$ and 0 if $i = j$. An example of a distance dependent traffic matrix with 4 nodes is shown in Table I.

With this traffic model, the total number of circuits (traffic units) is given by (for N odd), $2N \sum_{i=1}^{i=(N-1)/2} i = N(N^2-1)/4$. Accounting for the fact that traffic between a pair of nodes is symmetric and can be supported on the same circuit, the total number of circuits required, L, is $N(N^2-1)/8$. Similarly, when N is even, the total number of circuits required is $N(N^2+2N-4)/8$.

A. Bounds on the Number of ADM's

Unfortunately, with this traffic model, we are no longer able to obtain closed form expressions for the minimum number of ADM's required. However, we can still obtain some simple bounds and a heuristic algorithm that performs relatively well compared to those bounds. These bounds are a straightforward extension of the bounds derived in Section III. First observe that as before, the minimum number of wavelengths required a is equal to the number of circuits divided by g, and that without any grooming each of these wavelengths would have to be dropped at every node. This remains our upper limit on the number of ADM's needed, i.e., $N\lceil L/g \rceil$. A sometimes tighter upper-bound can be obtained by placing the traffic between nodes on dedicated wavelengths, so that each wavelength carries traffic between two nodes only. This bound becomes tight when the traffic between nodes is relatively large and approaches or exceeds g. With the distance dependent traffic model, this bound becomes particularly useful. Also, an obvious lower bound can be obtained by realizing that each wavelength can at most carry g circuits and has at least 2 ADM's. Hence, at most $g/2$ circuits can be supported per ADM and a lower bound on the number of ADM's is the total number of circuits divided by $g/2$, i.e., $\lceil 2L/g \rceil$.

The lower bound above assumed that all of the ADM's are used in the most efficient way possible. However, not all the circuits between all pairs of nodes can be supported in such way. For example when $g = 4$, if there are only three circuits between a pair of nodes, the associated efficiency is 1.5 circuit per ADM, which is the next most efficient way in which ADM's can be used. So to improve the lower bound, we need to find out the maximum number of circuits that can be supported in the most efficient way with average of two circuits per ADM. Then a tighter lower bound can be found by assuming that the remaining circuits are packed in the next most efficient way with average of 1.5 circuits per ADM. Let $C(N)$ be the maximum number of circuits between two nodes. By definition of the distance dependent traffic, for $N > 3, C(N) = \lfloor N/2 \rfloor$. Next, let $P(k)$ be number of node pairs with k circuits between them, for $k = 1, 2, \ldots, C(N)$. Again based on the definition of the distance dependent traffic, $P(k) = N$ for $k < C(N)$, $P(C(N)) = N/2$ if N is even, and N otherwise. Now, M_2, the maximum number of ADM's that can be used in the most efficient way (supporting two circuits per ADM), is given by $M_2 = \sum_{k=1}^{C(N)} 2 * \lfloor k/4 \rfloor * P(k)$. Finally, the rest of the circuits will be assigned with at most 1.5 circuits per ADM, leading to the improved lower bound for $g = 4$.

$$LB(g=4) = M_2 + \lfloor (L - 2 * M_2)/1.5 \rfloor.$$

TABLE I
Distance Dependent Traffic Matrix with 4 Nodes

N	1	2	3	4
1	**0**	**2**	**1**	**2**
2	**2**	**0**	**2**	**1**
3	**1**	**2**	**0**	**2**
4	**2**	**1**	**2**	**0**

B. Heuristic Algorithm

We now use a simple "greedy" algorithm for assigning traffic to wavelengths. The algorithm starts with wavelength 1 and node 1 and assigns as much traffic between nodes 1 and the other nodes (visited in sequential order). At each step, the algorithm first attempts to assign traffic to the wavelength using nodes that are already assigned to that wavelength. If that is not possible, the algorithm will add a new node to the wavelength if it can assign traffic between the new node and an existing node already on that wavelength. If traffic cannot be assigned by the addition of just one node, then the algorithm will start with a new wavelength.

```
Algorithm
  Start with wavelength 1 and node 1 as-
  signed to it.
  Assign traffic to wavelength i by vis-
  iting nodes in order as follows:
    a) If possible, assign traffic among
  nodes already assigned to wavelength i.
    b) Otherwise, if there exists a node
  that is not already assigned to wave-
  length i and that has traffic to a node
  already on wavelength i, add that node
  to wavelength i (when multiple such
```

[3] Although shortest path routing would not be possible in a unidirectional ring, the traffic demand between two nodes is still correlated to the shortest path distance.

```
nodes exist, they can be visited in
order).
If wavelength i is full or if it is not
possible to assign new traffic by just
adding a single node then start with a
new wavelength.
```

The above algorithm does not attempt to pack the traffic in a particularly efficient way but rather visits the nodes in order and assigns the traffic on a first-fit basis. As an illustrative example consider the traffic matrix of Table 1 with $g = 4$ (up to four circuits per wavelength). Start with the first wavelength and assign to it the traffic between nodes 1 and 2 and nodes 1 and 3 and one circuit between 2 and 3. Assign to wavelength 2 the traffic between 1 and 4 and the traffic between nodes 2–4. Finally, assign the remaining traffic between nodes 2 and 3 and the traffic between 3 and 4 to wavelength 3. The resulting assignment would then be: $\lambda 1$: 1–2, 1–3, 2–3 (1); $\lambda 2$: 1–4 (1), 2–4; $\lambda 3$: 2–3 (1), 3–4. Small improvements to the algorithms can be obtained by packing circuits more cleverly. For example, on $\lambda 1$, instead of putting the traffic between 1–3 and one circuit between 2–3, we could have assigned the circuits between 1–4 to $\lambda 1$ and placed the circuit between 1–3 on a separate wavelength. This minor improvement could result in a savings of one ADM (although in this example it does not result in any savings).

Fig. 6 shows the number of ADM's needed for the distance dependent traffic with $G = 4$ (OC-12 circuits on an OC-48 ring), using the above algorithm. As can be seen from the figure, the algorithm performs somewhere between the upper and lower bounds and results in significant improvement over having to drop all wavelengths at all nodes. As the number of nodes increases the relative amount of ADM savings also increases. This is due to the fact that with this traffic matrix, as the number of nodes increase, so does the amount of traffic between nodes. As the amount of traffic between nodes increases, significant savings are obtained because traffic between pairs of nodes can often be assigned dedicated wavelengths.

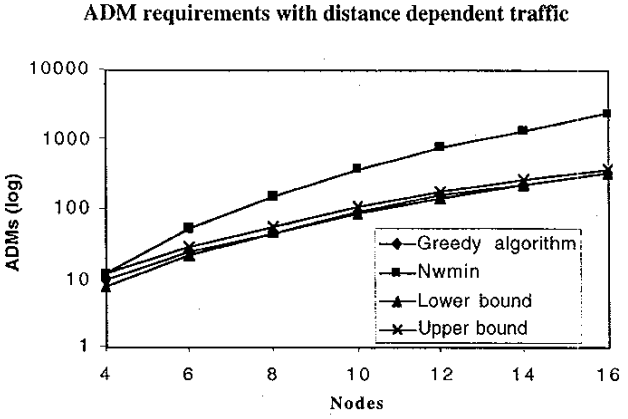

Fig. 6. ADM requirement for distance dependent traffic ($g = 4$).

V. Using a Hub with a SONET Cross Connect

Here we allow one node to have a SONET cross-connect, say node N, and we require that the cross-connect be present on every wavelength, so that all of the traffic can be routed through it if needed. We denote this node as a hub. The hub can take a circuit from one SONET ring and switch it to another ring. Again, we focus on the case with all-to-all uniform traffic where $r_{ij} = r$ for all $i \neq j$, and we assume that $L_d = r(N-1) \leq g$ (i.e., all of the traffic to and from a node can be carried on one wavelength).

Theorem 2: The optimal solution with one hub is either as good as or better than the optimal solution not using a hub in terms of minimizing the total number of ADM's.

Theorem 2 states that if we require the hub to be present on every wavelength, the resulting number of ADM's is not increased. While it may appear that the addition of a hub node should not result in any additional ADM's, it is not at all obvious that forcing the hub to be on every wavelength does not require additional ADM's. The following proof shows by construction that this is indeed the case, and in fact, the use of a hub often results in significant ADM savings.

Proof: Let node N be the hub node and consider any solution where node N is not on every wavelength, there exists a corresponding solution with the hub on every wavelength using the same or fewer ADM's. Since the hub node is also a regular node, every node has traffic going from and to the hub. Therefore each of the nodes on those wavelength(s) without the hub must have at least two ADM's (one on some wavelength without the hub and one on some wavelength with the hub). Since the traffic from each node can be carried on a separate wavelength through the hub using just two ADM's, any assignment not using the hub can be transformed into an assignment with the hub present on every wavelength using no additional ADM's. Of course, this solution may not use the minimum number of wavelengths. A further reduction in the number of ADM's can be obtained by packing the wavelengths optimally as we show next for the case of $r = 1$.

A. Optimal Algorithm when $r = 1$ and $L_d = N - 1 \leq g$

With the same argument as used for the egress node case, it can be shown that there exists a minimum ADM solution such that no traffic to and from a node is split onto two rings. This means that only one ADM is needed for every node except the hub, which has W ADM's, where W is the number of wavelengths used. This reduces the problem to minimizing W, which is equivalent to maximizing the number of nodes carried on a wavelength. Let K be the maximum number of nodes on a wavelength (including the hub node), then each wavelength with K nodes needs to carry two types of traffic. All-to-all traffic within the K nodes that does not need to go through the hub, of which there are $K(K-1)/2$ circuits; and cross traffic between the $K-1$ (excluding the hub) nodes and the remaining $N-K$ nodes not on the same wavelengths of which there are $(K-1)(N-K)$ circuits. This combined traffic load must be less than or equal to g, hence, $K(K-1)/2 + (K-1)(N-K)$ must be less than or equal to g. Expanding this expression and using the quadratic formula we obtain

$$K = \left\lfloor N + \frac{1}{2} - \frac{\sqrt{4N^2 - 4N - 8g + 1}}{2} \right\rfloor.$$

It can be shown that as long as K is less than N ($K = N$ corresponds to the case of $W = 1$ where all the traffic can be carried on one wavelength), the above expression yields a real value for K. The corresponding number of wavelength is $W = \lceil (N-1)/(K-1) \rceil$ and the corresponding number of ADM's $M = \lceil (N-1)/(K-1) \rceil + N - 1$, which is optimal.

VI. Conclusion

This paper studies the problem of assigning circuits to wavelengths with the objective of minimizing the cost of electronic multiplexing equipment. In particular, we consider the special case of SONET/WDM unidirectional ring networks, and attempt to minimize the number of SONET ADM's. While we show that the general problem is NP-complete, we are able to obtain encouraging results for some special cases where circuit rates are the same. In particular, in the case of an egress node we obtain the solution that minimizes the number of ADM's as well as a solution that minimizes the number of ADM's subject to using the minimum number of wavelengths. For all-to-all uniform traffic, and distance dependent traffic, we obtain a lower bounds on the number of ADM's and simple heuristic algorithms that perform close to the bounds.

We were also able to make a number of interesting observations and conjectures that provide insight into the traffic grooming problem. For example, we were able to show that it is not always possible to minimize both the number of ADM's and the number of wavelengths simultaneously. We conjecture, however, that when $r = 1$ such a solution can be found.

Finally, we consider the use of a hub node where traffic can be switched between SONET rings and show that, for the case where all of the traffic to and from a node can be carried on a single wavelength, a solution using a hub node always requires fewer or the same number of ADM's compared to a solution not using a hub node. We also obtain the optimal solution using a hub node and the corresponding minimum number of ADM's.

Yet, the work of this paper is preliminary and considers only a select number of special cases. Many interesting problems remain to be solved. For example, we still need to find the optimal solution and the optimal algorithm in the all-to-all uniform traffic case. Also, the benefits of using one or more hubs with a cross-connect require further study.

References

[1] O. Gerstel, G. Sasaki, and R. Ramaswami, "Dynamic wavelength allocation in WDM ring networks with little or no wavelength conversion," in *Proc. 1996 Allerton Conf.*, Oct. 1996.

[2] J. Bannister, L. Fratta, and M. Gerla, "Topological design of WDM networks," in *Proc. INFOCOM'90*.

[3] 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.

[4] J. M. Simmons, E. L. Goldstein, and A. A. M. Saleh, "On the value of wavelength-add/drop in WDM rings with uniform traffic," in *Proc. OFC'98*, San Jose, CA, Feb. 1998.

[5] M. R. Garey and D. S. Johnson, *Computers and Intractability*, New York: W. H. Freeman, 1979.

[6] X. Zhang and C. Qiao, "An effective and comprehensive solution to traffic glooming and wavelength assignment in SONET/WDM rings," in *Conf. On All-Optical Networking, SPIE Proc.*, vol. 3531, Boston, MA, Sept. 1998.

Angela L. Chiu received the B.S. degree in electrical engineering from the University of Connecticut, Storrs, in 1991 and the S.M. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1993 and 1997, respectively.

She was a Research Staff Member in the Optical Communication Technology Group of MIT Lincoln Laboratory, Lexington, MA. She is currently a Senior Technician Staff Member with the IP Data Services Integration Division of AT&T Labs, Middletown, NJ, where she is now responsible for the overall architecture and design in quality of service (QoS) and multiprotocol label switching (MPLS) for AT&T Worldnet IP backbone. Her research interests include architecture, design, and analysis of IP QoS, MPLS, WDM networks, and high-speed networks.

Eytan H. Modiano 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 Postdoctoral Fellow during 1992–1993, where he was conducting research on security and performance issues in distributed network protocols. Between 1993 and 1999, he was with the Communications Division at the MIT Lincoln Laboratory, Cambridge, where he worked on communication protocols for satellite, wireless, and optical networks and was the Project Leader for the MIT Lincoln Laboratory's Next Generation Internet (NGI) project. In 1999, he joined 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.