Minimizing Electronic Multiplexing Costs for Dynamic Traffic in Unidirectional SONET Ring Networks

Randall Berry
Laboratory for Information and
Decision Systems - MIT
Cambridge, MA 02139
Email: randy1@mit.edu

Eytan Modiano
MIT Lincoln Laboratory
244 Wood St.
Lexington, MA 02420
Email: modiano@ll.mit.edu

Abstract

In this paper we consider circuit assignment algorithms for dynamic traffic in unidirectional WDM/SONET ring networks. Our objective is to minimize the cost of electronic Add/Drop Multiplexers (ADMs) in the network, while being able to support any offered traffic matrix in a rearrangeably non-blocking manner. The only restriction on the offered traffic is a constraint on the number of circuits a node may source at any given time. We provide a lower bound on the number of ADMs required and give conditions that a network must satisfy in order for it to support the desired set of traffic patterns. Circuit assignment and ADM placement algorithms that perform closely to this lower bound are provided. These algorithms are shown to reduce the electronic costs of a network by over 30%. Finally, we discuss extensions of this work for supporting dynamic traffic in a wide-sense or strict sense non-blocking manner as well as the benefits of using a hub node and tunable transceivers.

I. Introduction

Optical WDM systems have increasingly been deployed as a point-to-point technology to increase the number of wavelengths in a network. Typically these networks have a SONET ring architecture and each additional wavelength is used to add an additional SONET ring between the nodes. The nodes in the ring use SONET Add/Drop Multiplexers (ADMs) to electronically combine lower rate streams onto a wavelength, *e.g.* 16 OC-3 circuits can be multiplexed onto one OC-48 stream. The cost of these electronic multiplexers dominates the costs of such a network.

The number of electronic ADMs can be reduced by employing WDM Add/Drop Multiplexers (WADMs) which allow a wavelength to either be dropped at a node or to pass through optically. 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 several papers including [1], [2], [3] and [4]. In [1] it was shown that the general grooming problem is NP-complete. 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 [4] and unidirectional rings [1]. In this paper, instead of considering only a single traffic matrix, we consider minimizing the number of ADMs needed to support any traffic matrix out of a given set. The resulting network can support traffic that dynamically changes within this set.

We consider a unidirectional ring network with N nodes. Each node has one WADM and a SONET ADM for every wavelength dropped at that node. The WADMs are static, *i.e.* the wavelengths that are dropped at each node are fixed. The SONET ADMs multiplex g low rate streams onto a single wavelength. A traffic matrix represents the number of low rate circuits, R_{ij} desired between nodes i and j. We assume that every connection is bi-directional so that $R_{ij} = R_{ji}$. Nodes do not have SONET cross-connects or optical wavelength changers, so a connection occupies a portion of the same wavelength around the entire ring. We limit each node to sourcing at most t bi-directional circuits at a given time, *i.e.*

$$\sum_j R_{ij} \le t \quad \text{for all } i \tag{1}$$

This will be the only constraint on the set of traffic patterns which must be supported. We will call any traffic matrix satisfying (1) *t-allowable*. We will refer to a traffic matrix as *t-maximal* if it is t-allowable and if the addition of any other circuit would make it not t-allowable. Clearly if a network can support every t-maximal traffic matrix it can support every t-allowable one.

For any t-allowable traffic matrix, the maximal load, L_{max}, on any link in a ring is given by

$$L_{\max} = \lfloor Nt/2 \rfloor, \tag{2}$$

and thus the minimum number of wavelengths, $W_{\min}$, that can accommodate any t-allowable traffic matrix is given by

$$W_{\min} = \lceil g^{-1} \lfloor Nt/2 \rfloor \rceil. \tag{3}$$

If we use $W_{\min}$ wavelengths and drop each wavelength at each node, then such a topology can clearly support any t-allowable traffic matrix. Thus $NW_{\min}$ gives an upper bound on the required number of ADMs. We will focus on reducing this number of ADMs while still supporting any t-allowable traffic matrix using $W_{\min}$ wavelengths. In [1] it is shown that the minimum ADM solution often uses the minimum number of wavelengths. Hence, restricting our solutions to those using the minimum number of wavelengths is sensible not only because it makes efficient use of wavelengths but also because it is likely to yield a nearly optimal solution.

The problem we address can be stated as follows. For given values of N, g and t, we wish to specify a topology, *i.e.* which of the N nodes have ADMs on each of the $W_{\min}$ wavelengths, and this topology must be able to support any

0-7803-5284-X/99/$10.00

t-allowable traffic matrix using the minimum number of ADMs. In the next section we will lower bound the number of ADMs needed and provide necessary and sufficient conditions that a network must satisfy to support such traffic. We then use these conditions to develop algorithms for allocating ADMs to nodes in the network. Finally, in section V, we develop extensions to the basic model that allow a network to be non-blocking in a strict sense; and the use of a hub architecture and tunable lasers in order to achieve further reductions in electronic multiplexing costs is considered.

II. Lower Bound on the number of ADMs.

We start by deriving a lower bound on the required number of ADMs conditioned on using the minimum number of wavelengths. Minimumizing the required number of ADMs is equivalent to starting out with every node having an ADM on each wavelength and maximizing the number of ADMs that can be removed while still supporting every t-allowable traffic matrix. Let $X = \{1,\ldots,N\}$ be the set of nodes and let M_i be the set of nodes with ADMs removed from wavelength i for i=1...$W_{\min}$. The following lemmas will help us to bound the maximum number of ADMs which can be removed from each wavelength.

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

$$|M_i| < (W_{\min} - 1)g/t. \tag{4}$$

Proof: First we show that if we can support every t-allowable traffic set then for all i we must have $|M_i| < N/2$. Assume this is not true, *i.e.* for some i, $|M_i| \geq N/2$ and additionally assume that N is even. In this case we can construct a t-maximal set where every connection involves a node in M_i and thus no circuit in this set can be supported on wavelength i. Furthermore this set will require at least $W_{\min}$ wavelengths and so can not be supported. To see how such a t-maximal set can be formed pair each node not in M_i with a node in M_i. Next, pair the remaining nodes within M_i among each other and assign t circuits between each pair of nodes. Since N is even, every node can be paired up in this way and thus the traffic load for this set will be $L_{\max}$. Clearly, this set requires $W_{\min}$ wavelengths and each connection involves a node from M_i. When N is odd, by a slightly more complicated argument we can construct a similar set.

Now we prove the lemma by contradiction, *i.e.* assume for some i that (4) is not true but we can support every t-allowable traffic set. We have just shown that $|M_i| < N/2$. In this case we can pair up each node in M_i with a distinct node in $X - M_i$. We form a t-allowable traffic set by setting up t circuits between each pair. The resulting traffic matrix consists of $|M_i|t$ circuits, none of which can be placed on wavelength i. This set must be placed on the remaining $W_{\min} - 1$ wavelengths, but these wavelengths can accommodate at most ($W_{\min}$-1)g circuits. Thus this t-allowable set can not be supported, yielding a contradiction. ■

Lemma 2: If a network with $W_{\min}$ wavelengths can support every t-allowable traffic set then

$$\min(|M_i|,|M_j|) \leq (W_{\min} - 2)g/t \text{ for all } i \neq j. \tag{5}$$

Proof: We prove this in a similar manner to Lemma 1. That is we show that we can always construct a t-allowable set with $\min(|M_i|,|M_j|)t$ circuits which can not be carried on either wavelength i or j and thus must be carried on the other $W_{\min}$-2 wavelengths. Since each wavelength can accommodate at most g circuits, (5) must be true for this set to be supported. The proof will be completed once we show how to construct the above set.

Consider two wavelengths i and j and assume that $|M_i| \leq |M_j|$. Let K be the set of nodes removed from both i and j (K may be empty). From the proof of Lemma 1, we can assume that $|M_i|$ and $|M_j|$ are both less than N/2. Thus, we can pair each node in K with a distinct node in $X - M_j$. Likewise, we can pair up every node in $M_i - K$ with a distinct node in $M_j - K$. If we put t circuits between each pair, this gives the required t-allowable set. ■

An immediate corollary of lemma 2 is that for every wavelength except one we must have

$$|M_i| \leq (W_{\min} - 2)g/t \tag{6}$$

Lemma 1 gives a bound on the ADMs that can be removed on the remaining wavelength. Thus, we have the following upper bound:

$$\text{ADMs removed} \leq \left\lfloor \frac{g}{t}(W_{\min} - 2) \right\rfloor (W_{\min} - 1) + \left\lfloor \frac{g}{t}(W_{\min} - 1) \right\rfloor \tag{7}$$

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

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

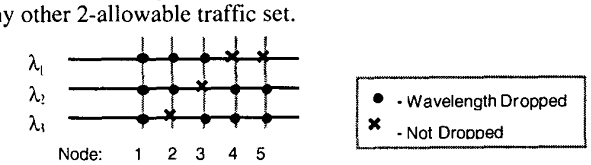

Figure 1: Provisioning of ring in Example 1.

In the next section we will establish a connection between this problem and bipartite matching problems. Exploiting this connection we will come up with necessary and sufficient conditions for an allocation of ADMs to be able to support any t-allowable traffic matrix. This will then be used to develop several heuristic algorithms for removing ADMs from wavelengths.

III. Bipartite Matching formulation

A bipartite graph, (C,D,E), is a graph with two disjoint sets of nodes, C and D, and a set of edges, E, where each edge is between a node in C and a node in D. For a given ring network, we want to construct a bipartite graph which represents the possible placements for each call from a given t-allowable traffic matrix. We will denote one set of nodes in the graph by $D=\{\lambda_{1,1},\ldots,\lambda_{1,g},\lambda_{2,1},\ldots,\lambda_{W_{\min},g}\}$. This set contains g elements for each of the $W_{\min}$ wavelengths corresponding to possible circuit assignments on that wavelength. The other set of nodes will correspond to the set of requested circuits, we will denote this set by C. There is an edge in the graph between $\lambda_{i,j}$ and a circuit involving nodes k and l if both nodes k and l have an ADM on wavelength i.

Example 2: Consider the allocation of ADMs and the traffic set from example 1. The corresponding bipartite graph is shown in Figure 2.

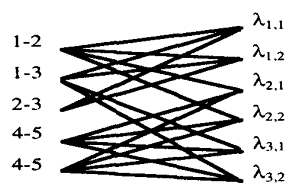

Figure 2: A Traffic matrix and the corresponding bipartite graph.

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

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

A proof of this theorem can be found in many texts on combinatorics such as [6]. As stated, this theorem is useful to check that a single traffic matrix can be supported, but for the problem at hand we are interested in supporting every t-allowable traffic matrix. The following theorem provides a necessary and sufficient condition for accommodating every t-allowable matrix.

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

1) For every pair of nodes i and j there exists a wavelength on which both i and j have an ADM.

2) For any group of m wavelengths, there exists at most gm circuits, out of any t-allowable set, which must be routed on this group. We say a circuit between nodes i and j must be routed on a set of wavelengths if for any wavelength not in this set, either i or j does not have an ADM on that wavelength.

Proof: We first show that the above conditions are necessary. Clearly, if 1 is not satisfied then any t-allowable set containing a circuit between i and j can not be accommodated. If 2 is not satisfied then there exists a set of m wavelengths on which we must route more than gm circuits in some t-allowable set, C. Consider the bipartite graph corresponding to C. Let S be the subset of C containing the above circuits, then we have $|N(S)|=gm$ and $|S|>gm$. Thus by Hall's theorem there exists no C-saturated matching, and this traffic matrix can not be accommodated.

Next we show that these conditions are sufficient. Assume that they are not sufficient, so that there exists an assignment of ADMs to $W_{\min}$ wavelengths which satisfies 1 and 2 above, but which can not support some t-allowable traffic set C. Since C cannot be supported, by Hall's theorem there exists a subset S of C such that $|N(S)|<|S|$. Let k be a nonnegative integer such that $(k-1)g<|S|\leq kg$. For a bipartite graph corresponding to an allocation of ADMs, $|N(S)|$ will always be a multiple of g. Thus, $|N(S)|<|S|$ implies that $|N(S)|\leq(k-1)g$. Therefore this set of more than $(k-1)g$ calls must be routed on a set of $k-1$ or fewer wavelengths, which contradicts condition 2, completing the proof. ■

IV. Algorithms for removing ADMs

We now use the results of the previous section to develop algorithms for removing ADMs from wavelengths. First note that if $W_{\min}=1$, we cannot remove any ADMs and still support all t-allowable traffic matrices. If $W_{\min}=2$, we have to leave every node on one wavelength and can remove at most $\lfloor g/t\rfloor$ nodes from the other wavelength. This follows directly from Lemmas 1 and 2. Furthermore, we can always remove $\lfloor g/t\rfloor$ nodes from one wavelength. To see this note that the most circuits that will be forced on 1 wavelength is $\lfloor g/t\rfloor t<g$. So, by the Theorem 1 we can accommodate all t-allowable circuits. Thus for $W_{\min}\leq 2$ we have a trivial algorithm which yields the minimum number of ADMs. Therefore in the following we shall only consider the case where $W_{\min}\geq 3$.

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

For a circuit to be forced on a set of i wavelengths, one of the two nodes in the circuit must be removed from each of the remaining $W_{\min} - i$ wavelengths. When $W_{\min} - i > 2$ then at least one of the two nodes must have an ADM removed from more than one of the wavelengths. So if we remove at most one ADM for each node, we only have to check condition 2 for sets of $W_{\min}$ -1 and $W_{\min}$-2 wavelengths. Clearly, we can remove $\lfloor N / W_{\min} \rfloor$ nodes from each wavelength so that no node will be removed from more than one wavelength. Also if we remove $\lfloor (W_{\min} - 2)g / t \rfloor$ or fewer nodes per wavelength, then no more than $(W_{\min}$-$2)g$ circuits will be forced on any set of $W_{\min}$ - 2 or $W_{\min}$ - 1 wavelengths. Thus if we remove

$$M_1 = \min\left(\lfloor (W_{\min} - 2)g / t \rfloor, \lfloor N / W_{\min} \rfloor\right)$$

ADMs per wavelength and no more than one ADM per node, condition 2 of theorem 1 is satisfied. Condition 1 is also easily satisfied in this case. Thus we have proved the following lemma which immediately yields an algorithm for allocating ADMs.

Lemma 3: For $W_{\min} > 2$, one can always remove M_1 ADMs from each of $W_{\min}$ wavelengths such that no node has more than one ADM removed and any t-allowable traffic matrix can be supported.

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

$$\lfloor (W_{\min} - 2)g / t \rfloor < \lfloor N / W_{\min} \rfloor \tag{8}$$

the above algorithm removes the most nodes possible from every wavelength except possibly one. When $W_{\min}$ becomes large for a given N, then the inequality in (8) is reversed and the procedure in lemma 3 results in removing only a small number of ADMs from each wavelength. In such cases, we have to consider removing nodes from more than one wavelength.

We consider a procedure that allows nodes to be removed from multiple wavelengths. This procedure is a generalization of the one in Lemma 3. For now we assume that $N \geq W_{\min}$. For given integers x and k, suppose we remove nodes $(i-1)\lfloor N / W_{\min} \rfloor + 1$ to $\left((i + x - 2)\lfloor N / W_{\min} \rfloor + k\right) \bmod N$ from wavelength i, where $0 \leq k \leq \lfloor N / W_{\min} \rfloor$. Thus we remove $(x-1)\lfloor N / W_{\min} \rfloor + k$ nodes from each wavelength, and a node is removed from at most x wavelengths. Traffic is then only forced onto groups of $W_{\min}$-$2x$ or more wavelengths. If we set x=1, then each node has at most one ADM removed. In this case, by Lemma 3 if $k<M_1$ we can support all t-allowable traffic. For an arbitrary value of x, as long as x is less than $W_{\min}/2$ then condition 1 of Theorem 1 is satisfied. We only need to check condition 2 for sets of $W_{\min}$-$2x$ or more wavelengths. The most circuits that can be forced on a set of $W_{\min}$ -$2x$ wavelengths is kt. This occurs when we consider $2x$ adjacent wavelengths and note that there are k nodes without an ADM on any of the first x wavelengths and k other nodes without an ADM on the next x wavelengths. If we pair up each node in the first set with a distinct node in the second set and establish t circuits between these nodes, then this set of calls can not be supported. By similar reasoning we can find the most circuits that can be forced on sets of $W_{\min}$ -($2x$ -1) to $W_{\min}$ -1 wavelengths. In this manner we get the following set of inequalities which must be satisfied for condition 2 to hold.

$$\left(W_{\min} - 2x\right)g \geq kt$$

$$(W_{\min} - (x+i))g \geq \left((x-1-i)\lfloor N / W_{\min} \rfloor + 2k\right)t, \quad \forall i = 1,.., x-1$$

$$(W_{\min} - i)g \geq \left((x-1)\lfloor N / W_{\min} \rfloor + k\right)t, \quad \forall i = 1,.., x$$

Out of this set of $2x$ inequalities, it can be shown via algebraic manipulations that if the following three inequalities are satisfied then the entire set of $2x$ must also be.

$$(W_{\min} - (2x-1))g \geq 2kt \tag{i}$$

$$(W_{\min} - (x+1))g \geq \left((x-2)\lfloor N / W_{\min} \rfloor + 2k\right)t \tag{ii}$$

$$(W_{\min} - x)g \geq \left((x-1)\lfloor N / W_{\min} \rfloor + k\right)t \tag{iii}$$

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

$$\max \ (x-1)\lfloor N / W_{\min} \rfloor + k$$

subject to:

$$\begin{aligned} &(W_{\min} - (2x-1))g \geq 2kt \\ &(W_{\min} - (x+1))g \geq \left((x-2)\lfloor N / W_{\min} \rfloor + 2k\right)t \\ &(W_{\min} - x)g \geq \left((x-1)\lfloor N / W_{\min} \rfloor + k\right)t \\ &0 \leq k \leq \lfloor N / W_{\min} \rfloor \\ &1 \leq x \leq \lfloor W_{\min} / 2 \rfloor. \end{aligned} \tag{P}$$

Where x and k are constrained to be integers. This optimization problem can be solved in the following manner: First set k=0 and find the largest value of x which satisfies the constraints. Next, fix x at this value and find the largest value of k satisfying the constraints. Again we summarize these results in the following lemma.

Lemma 4: Consider a ring with $W_{\min}$ >2. Then we can remove $(x-1)\lfloor N / W_{\min} \rfloor + k$ ADMs per wavelength in the above manner where x and k are solutions to the integer program (P) and still support every t-allowable traffic matrix.

This lemma immediately yields an algorithm for removing ADMs. The following provides an example of this algorithm along with the algorithm in Lemma 3.

Example 3: Consider a ring with 15 nodes, g = 16, and t = 10. For this ring, $W_{\min}$ = 5 and $\lfloor N / W_{\min} \rfloor$ = 3. Thus using the algorithm from Lemma 3 we can remove 15 ADMs. Using the algorithm in Lemma 4, one finds that x = 2 and k = 1, and thus one can remove 4 nodes per wavelength, for a total of 20 ADMs removed. For comparison, the upper bound on the number of ADMs removed from (6) is 22. The resulting allocation of ADMs for both algorithms is shown in Figure 3.

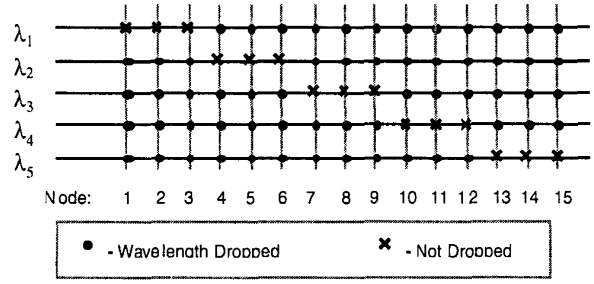

Figure 3a: Allocation corresponding to Lemma 3.

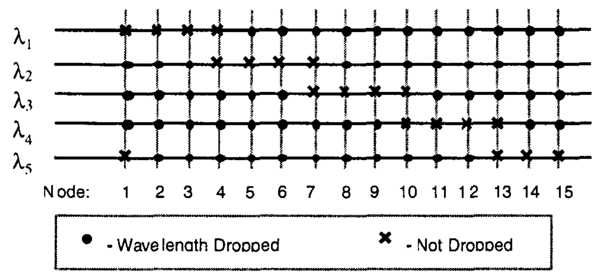

Figure 3b: Allocation corresponding to Lemma 4.

In Figure 4 we have plotted the number of ADMs resulting from both of these algorithms as t ranges from 1 to 14. The number of ADMs with no grooming is also plotted along with the lower bound from (6). With the grooming algorithm from Lemma 4, the number of ADMs is reduced by up to 27%. In [1] it was found that approximately 60 ADMs were needed to support uniform all-to-all traffic in this network. Uniform all-to-all traffic is (N-1)-maximal. In this example, supporting all (N-1)-allowable traffic requires 77 ADMs, but this is a much less restrictive set of traffic.

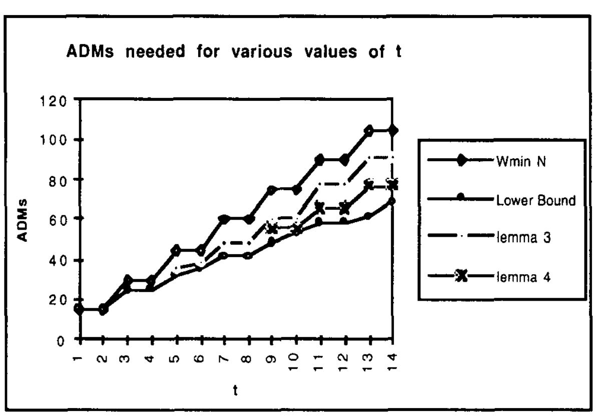

Figure 4: Required ADMs for heuristic algorithms, N=15, g=16.

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

$$kN + \lceil \lfloor Nt'/2 \rfloor / g \rceil = \lceil \lfloor Nt/2 \rfloor / g \rceil$$

i.e. decomposing traffic in this way requires no more wavelengths. Now since the number of wavelengths needed for each set in this decomposition is less than or equal to N, we can apply the above algorithms to remove ADMs from each set to get an allocation of ADMs which will support all t-allowable traffic sets.

Example 4: Consider a ring with N = 5, g = 2, and t = 6 so that $W_{\min}$ =8. Applying the above procedure we get one set of 5 wavelengths which must support 4-allowable traffic and one set of 3 wavelengths which must support 2-allowable traffic. Applying Lemma 4 to both of these sets, we find we can remove 1 ADM from each wavelength and thus eliminate a total of 8 ADMs.

V. Extensions to the basic model.

In this section we will describe a number of extensions to our basic model. First, we discuss the application of our approach to a system with a hot-spot node which can source an unlimited number of circuits. We then discuss the use of a strict sense non-blocking network to support rapidly changing traffic and finally we discuss the benefits of using a hub node and tunable lasers.

A. Hot-spot node

Suppose there is one node in the network which has no restriction on the number of circuits it can source. We will refer to this node as a "hot spot". Such a node can be used to model a ring with a central office node. In this case the minimum number of wavelengths required to support all t-allowable traffic is given by:

$$W_{\min} = \lceil (N-1)t/g \rceil \qquad (9)$$

Clearly the hot spot node needs an ADM on each wavelength. Consider applying the grooming algorithms to the set of nodes not including the hot spot with the above number of wavelengths. The resulting provisioning of the network will then handle all t-allowable traffic between these nodes. This allocation is also sufficient to handle all t-allowable traffic including the hot spot node. To see this note that by including the hot spot node we are forcing no additional calls onto any group of wavelengths, and thus by Theorem 1 we can still support any t-allowable traffic matrix. This procedure applies with an arbitrary number of hot spots.

B. Supporting dynamic traffic without rearranging calls

In the above sections we examined finding minimal allocations of ADMs needed to support any t-allowable traffic matrix. Given such an allocation of ADMs , one then needs to find how to assign calls from a given traffic matrix to the proper ADMs. When t is only one, there are $N!(2^{-N/2})/((N/2)!)$ different maximal traffic matrices and an assignment is needed for each of these. For a given traffic matrix one could set up a bipartite graph as in section III, and solve the maximum matching problem. Polynomial algorithms for solving this problem are known and can be found, for example, in [7]. In many cases an assignment can be found by inspection and the above matching problem would not need to be solved. These assignments can be all computed off-line and stored in a look-up table. Also, in some cases, the assignments can be stored in a more compact form than simply listing every possible assignment.

In this section we want to consider traffic which is dynamically changing, but is always t-allowable. When discussing provisioning rings for dynamic traffic we will use some standard definitions from switching theory which we repeat here. A ring is *strict sense* or *strictly non-blocking* if any t-allowable circuit between nodes can be established without interference from any other existing allowable circuits. A ring is *wide sense non-blocking* if any t-allowable circuit between nodes can be established without interference from any other existing allowable circuits, provided that the existing circuits have been established according to some algorithm. A ring is *rearrangeably non-blocking* if any t-allowable circuit can be established by possibly re-routing any existing circuits. Clearly the following relationship holds:

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

The converse implications do not in general hold.

A ring provisioned according to the algorithms in section IV is rearrangeably non-blocking but not necessarily strictly or wide-sense non-blocking. If traffic changes frequently then the control overhead associated with re-arranging existing circuits may not be acceptable. In such a case, one may prefer a ring that is either wide-sense non-blocking or strictly non-blocking. If every node has an ADM on each of $W_{\min}$ wavelengths then the resulting ring is strictly non-blocking. The following result shows that for a ring with $t = 1$ and $W_{\min}$ wavelengths to be strictly non-blocking then each node must have an ADM on every wavelength, in other words one cannot save on the cost of ADMs by grooming.

Theorem 2: A strictly non-blocking ring with $t = 1$ and $W_{\min}$ wavelengths must have an ADM for each node on each wavelength.

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

For $W_{\min} > 2$, we will proceed by induction. First note that there must be at least $(N/2)+1$ nodes on each wavelength for the ring to be even rearrangeably non-blocking; this follows from the proof of Lemma 1. When $W_{\min} > 2$ and $t = 1$, it follows from the definition of $W_{\min}$ that $N/2 \geq 2g+1$. Thus there must be more than $2g+2$ nodes on each wavelength. Now assume that the theorem is true for $W_{\min} = k$ wavelengths, and consider the case when $W_{\min} = k+1$. Without loss of generality we can assume that nodes $1,\ldots,2g+2$ are on wavelength 1. Thus we can consider any 1-allowable set of g circuits between $2g$ of these nodes and place these circuits on wavelength 1. Then any other 1-allowable set of calls between the remaining N-$2g$ nodes must be placed on the remaining k wavelengths. If we consider a ring with these N-$2g$ nodes, then the minimum number of wavelengths for this ring is k. Therefore, by the induction hypothesis, we can't remove any of these N-$2g$ nodes from the remaining k wavelengths. The original $2g$ nodes were picked arbitrarily from the set of $2g+2$ nodes that must be on wavelength 1, and by choosing different sets and repeating this argument we have that every node must be on the remaining k wavelengths. Likewise by repeating this argument but starting with a different initial wavelength we see that every node must be on every wavelength. Thus the theorem is true for $W_{\min} = k+1$, and by induction is true for any ring with $t = 1$. ■

Though we have only proved this for $t = 1$, the proof can be modified for an arbitrary value of t. As a consequence of this theorem, if we want to save on ADMs for dynamic traffic without re-arranging, we must consider wide-sense non-blocking networks. Analyzing wide-sense non-blocking rings is more difficult than the other cases due to the fact that a routing algorithm must also be considered. The following gives an upper bound on the ADMs that can be removed for a wide-sense non-blocking ring.

Lemma 5: Consider a unidirectional ring with $W_{\min}$ wavelengths. Let M_i be the set of nodes removed from wavelength i. For the ring to be wide-sense non-blocking for t-allowable traffic, where t is even, we must have for all i:

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

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

First we prove if the ring is wide-sense non-blocking, then:

$$|X| \geq |M_i| / 2 \tag{10}$$

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

If (10) is true, beginning with a t-maximal set as above, assume that the $|M_i|t/2$ circuits involving the nodes in M_i are disconnected along with $|M_i|t/2$ circuits involving nodes from X. Then we can form $|M_i|t$ circuits as above, where each circuit is between one node from M_i and one node that was in a circuit in X. These additional circuits must be routed on the remaining $W_{\min}$ -1 wavelengths without rearranging the active calls. This means that at most $(W_{\min} - 1)g - |M_i|t$ calls not involving the nodes in M_i can be routed on these wavelengths. There are $(N - |M_i|)t/2$ circuits in the original maximal set not involving nodes in M_i, thus we must have

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

We also must have $|X| \leq g$; combining these and performing some algebra yields the desired result. ■

Note we assumed that t was even in this lemma just to simplify the proof, a similar bound could be found for t odd. Consider our previous example with N =15, g = 16, and t = 10. In this case the above bound is $|M_i| \leq 1$ so for the network to be wide-sense non-blocking the most ADMs that could be removed is 5. Compare this with 20 ADMs that can be removed for a rearrangeably non-blocking ring. In fact for this example with t taking on any even value between 2 and 14, the size of M_i is always bounded to be less than or equal to 1, resulting in at most an 8% reduction in ADMs.

These results suggest that to get great benefits from grooming with dynamic traffic, some rerouting of existing traffic is needed, at least within the unidirectional ring model considered here.

C. Using a hub node and tunable lasers

By investing in a more sophisticated components else where in the network, one can gain further reductions in the cost of the electronic layer multiplexing. Two examples of this which we will consider is the use of a hub node and the use of tunable lasers. First we consider a hub node. By a hub node we mean a node which has ADMs on every wavelength and has a SONET cross-connect. By similar arguments to those used in [2] we can show that making one node in the ring such a hub node will not require any more ADMs than were required without the hub. Assuming that $t \leq g$ we can then show that the minimum number of ADMs needed to support all t-allowable traffic with a single hub node is given by

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

For example suppose that we have a ring with N = 7, g = 2, and t = 1. Using the algorithm from Lemma 4 we need 12 ADMs to support all t-allowable traffic. By making one node a hub node, we can support this traffic using only 10 ADMs. We can also reduce the required number of ADMs if instead of having fixed tuned lasers, each node is equipped with tunable lasers. For example, again consider the ring with N = 7, g = 2, and t = 1. If nodes are equipped with tunable lasers then each node only needs one ADM, and thus we need only 7 ADMs for the entire ring. In this case using tunable lasers reduced the required number of ADMs by 58%.

Clearly with tunability a node needs no more than t ADMs, thus when t is small there is a clear advantage to tunability. On the other hand for larger values of t the gain from tunability is not as obvious and is an open issue.

VI. Conclusion

In this paper we examine the problem of designing a WDM ring network to support dynamic SONET traffic. The goal of our design is to minimize the number of electronic multiplexers (e.g., SONET ADMs) used in the network. We developed a number of algorithms for assigning ADMs to wavelengths in a way that supports every t-allowable traffic matrix in a rearrangeably non-blocking manner. These algorithms are shown to reduce the number of ADMs needed by up to 33%. We also derive a lower bound on the number of ADMs required to support all t-allowable traffic and show that in some cases our algorithms perform close to this bound. Finally, we discuss extensions of our model to include supporting dynamic traffic in a strictly non-blocking manner and to hot-spot traffic where some nodes have a heavier traffic load.

References:

[1] E. Modiano and A. Chiu, "Traffic Grooming Algorithms for Minimizing Electronic Multiplexing Costs in Unidirectional SONET/WDM Ring Networks," CISS 98, Princeton, NJ, February, 1998.

[2] A. Chiu and E. Modiano,"Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming," Globecom 98, Sydney, Australia, Nov. 1998.

[3] O Gerstel and R. Ramaswami, "Cost effective Grooming in WDM Rings," Infocom 98, San Francisco, CA, April, 1998.

[4] J. Simmons, E Goldstein, and A. Saleh, "On the value of Wavelength Add/Drop in WDM Rings With Uniform Traffic," OFC '98, San Jose, CA, February, 1998.

[5] O. Gerstel, P. Lin and G. Sasaki, "Wavelength Assignment in a WDM Ring to Minimize the Cost of Embedded Sonet Rings," Infocom 98, San Francisco, CA, April, 1998.

[6] F. Roberts, *Applied Combinatorics*, Prentice Hall, 1984.

[7] D. Bertsimas and J. Tsitsiklis, "Introduction to Linear Optimization," Athena Scientific, 1997.