The Role of Switching in Reducing Network Port Counts

Eytan Modiano
MIT Laboratory for Information and Decision Systems
Cambridge, MA 02139
modiano@mit.edu

Randall Berry
Northwestern University
Dept. of Electrical & Comp. Eng.
Evanston, IL 60208
rberry@ece.nwu.edu

Abstract

We consider the role of switching in minimizing the number of electronic ports (e.g., SONET ADMs) in an optical network that carries sub-wavelength traffic. Providing nodes with the ability to switch traffic between wavelengths, such as through the use of SONET cross-connects, can reduce the required number of electronic ports. We show that only small switches at multiple switching nodes are needed for significant reductions in the number of ports. We provide a lower bound on the number of electronic ports that is a function of the number of switching nodes. We show that our lower bound is relatively tight by providing routing and grooming algorithms that nearly achieve the bound. For uniform traffic, we show that the number of electronic ports is reduced when the number of switching nodes used is approximately equal to the number of wavelengths of traffic generated by each node. Finally, we provide a general upper bound on the amount of switching required in the network. For uniform traffic, our bound shows that as the size of the network increases, each traffic stream must be switched at most once (at one of the switching nodes) in order to achieve the minimum port count.

1 Introduction

Optical WDM systems have increasingly been deployed to increase network capacity. Typically these networks have a SONET ring architecture; 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. To reduce the number of electronic ADMs, WDM Add/Drop Multiplexers (WADMs) can be employed; WADMs allow a wavelength to either be dropped at a node or to optically bypass a node. When a wavelength is not dropped at a node, an electronic ADM is not required for that wavelength. The required number of SONET ADMs can be further reduced by *grooming* the lower rate traffic so that the minimum number of wavelengths needs 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 bi-directional rings [4] and unidirectional rings [1]. In much of the work on grooming it is assumed that each low-rate circuit must stay on the same wavelength between the source and destination. This assumption can be relaxed when a node is equipped with SONET digital cross-connect (DCX), which allows for the electronic switching of low rate streams between wavelengths. The added flexibility provided by DCXs can reduce the required number of ADMs in a network. An example of this is given in [1] where it is shown that only equipping a single *hub* node with a DCX can reduce the required number of ADMs over a network with no switching capability.[1] However, there is a non-negligible cost associated with providing this electronic switching. Therefore, it is also desirable to limit the amount of switching in the network.

In this paper we study how much switching is needed for an architecture that attempts to reduce the required number of ports. We consider generalizations of the single hub architecture in [1], where the cross-connect function is distributed among multiple nodes in the ring. Such architectures have three significant advantages. First, the use of multiple cross-connects can reduce the number of SONET ADMs needed. Second, using multiple smaller cross-connects rather than one large cross-connect at the hub reduces the cost of the cross-connects. Finally, the use of cross-connects for grooming adds flexibility to the network over a static solution that does not use a cross-connect. This flexibility allows traffic to be provisioned dynamically thereby reducing the need to know the exact traffic requirements in advance. Another benefit of this flexibility is that the network will be more robust to node failures. Also, we note that a multiple hub solution will often require fewer wavelengths than a single hub solution.

Next we look at some preliminary examples that illustrate the potential benefits of a multiple-hub architecture. We then give a bound on the ADM requirement for a K-hub ring, and provide grooming algorithms for such a ring. Finally we discuss a bound on the average amount of switching needed for a ring that minimizes the required number of ports.

2 SONET/WDM Ring architectures

Consider three possible ring architectures for the purpose of efficient grooming: a static ring without cross-connects, a single-hub ring, and a ring with multiple-hubs. With the static architecture no cross-connecting is employed, hence each circuit must be assigned to a single wavelength that must be processed (dropped) at both the source and the destination. The static architecture is the traditional SONET ring architecture that has been used in the studies of [1,2,4]. The single hub architecture uses a large cross-connect at one hub node. The cross-connect is able to switch any low rate circuit from any incoming wavelength to any outgoing wavelength. With this architecture, each node sends all of its traffic to the hub node where the traffic is switched, groomed and sent back to the destination nodes. In the multiple hub architecture, K hub nodes are used on the ring. Each hub node has a small cross-connect that can switch traffic among the wavelengths dropped at that node. Each node on the ring sends a fraction of its traffic to one of the hub nodes, where it is properly groomed and relayed to its destination. These three architectures are depicted in Figure 1. Shown in Figure 1a is the static grooming solution where one wavelength is used to support traffic between nodes 1,2 and 3, another for traffic between 2,3 and 4, and a third wavelength for traffic between 1, 3 and 4. The hub architecture shown in Figure 1b has each node send all of its traffic to the hub located at node 3, where the traffic is groomed and relayed back to its destination. Finally shown in Figure 1c is the multiple hub architecture where each node can send its traffic to one or more of the hubs.

To illustrate the potential benefit of the multiple hub architecture, consider a unidirectional ring with N=9 nodes where each wavelength supports an OC-48 and traffic demand is uniform with two OC-12's between each pair. In this case each node generates 16

[1] Moreover, the result in [1] shows that this reduction is possible even if the hub node is required to have an ADM on every wavelength.

OC-12's or four wavelengths of traffic. With the single hub solution, each node can send all four wavelengths worth of traffic to be groomed at the hub at say node 1. In this case, each node would use 4 ADMs, and the hub would use 8x4=32 ADMs for a total of 64 ADMs. In a 2-hub architecture each node would send two wavelengths worth of traffic to each hub (at nodes 1 and 5) and an additional wavelength would be used for traffic between the two hubs, resulting in 58 ADMs. Finally a 4-hub architecture can be used where each node sends one wavelength to each of four hubs and some additional ADMs are used to handle the inter-hub traffic. Using the grooming algorithm given in Sect. 3, a 4-hub architecture can be found that requires only 52 ADMs. In the next section we give a lower bound on the number of ADMs required for with any amount of switching; for this example, this bound is 48 ADMs. Thus, with 4 hubs the bound is nearly met, and any further increase in the amount of switching could at best result in only a moderate additional savings of ADMs.[2] Notice that in this case the number of hubs is equal to the number of wavelengths generated by a node. Also notice that in increasing the number of hubs from 1 to 4 the required number of wavelengths in the ring is reduced from 32 to 26. Thus the 4-hub architecture is more efficient in the use of wavelengths as well as ADMs.

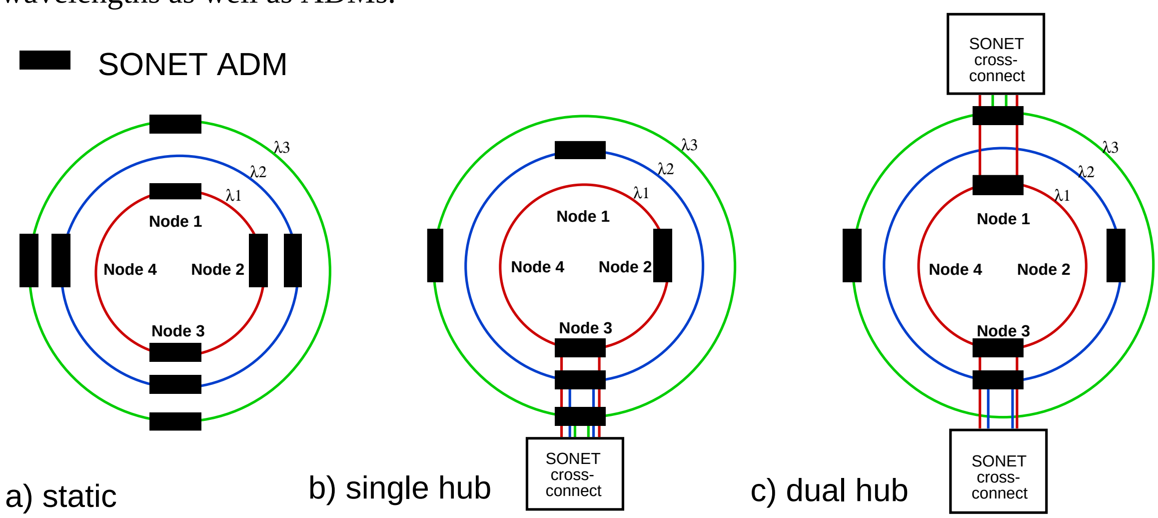

Figure 1. Grooming Architectures

3 Bounds on the required number of ADMs

In the following we develop a lower bound on the required number of ADMs for a K-hub architecture. To simplify the discussion we only consider a unidirectional ring network, *e.g.* a UPSR SONET ring, with N nodes. We consider the case where all traffic has the same granularity of g, and there is a uniform traffic demand of $r \le g$ circuits between each pair of nodes in the ring. The bound we develop is partly based on a lower bound on the number of ADMs needed in a unidirectional ring that was first derived in [5]; we present this bound next. A lightpath in the ring is a single wavelength connection between two nodes, which is not dropped at any intermediate nodes. Thus when a circuit is carried over n lightpaths, it is either dropped and continued or switched n-1 times; in this case each lightpath is said to carry 1/n-th of the "full" circuit. The lower bound is obtained by recognizing that each lightpath in the network must be terminated with exactly two ports. Thus, a lower bound on the number of lightpaths needed to support all of the traffic in the network can be translated into a lower bound on the number of ports (ADMs). Since the direct traffic between two

[2] Indeed, a 4-hub architecture requiring only 49 ADMs can be found by being more clever in the routing of inter-hub traffic, thus the possible savings with more switching capability is at most 1 ADM.

nodes is equal to r low rate circuits, each lightpath can at most carry r "full" circuits entirely from their source to their destination. The remaining capacity of that lightpath (g-r) can only be used to carry circuits that are also carried on *at least* one other lightpath. Hence, each lightpath can carry at most $Q = r + (g-r)/2$ "full" circuits. Since the total traffic demand under the uniform traffic assumption is equal to $L = N(N-1)r$ circuits, the number of lightpaths required is lower-bounded by L/Q. Since each lightpath is terminated at a port, the number of ports needed is at least $2L/Q$. However, SONET ADMs can be used both as a receiving and transmitting port, thus we have the following lower bound[3]

$$\text{ADMs} \geq L/Q = \frac{2N(N-1)r}{g+r}. \quad (1)$$

The bound in (1) will in general not be tight, but it can be achieved in several cases. It is insightful to consider some characteristics of these cases. From the above, it can be seen that for (1) to be tight, each lightpath must be efficiently packed so that it contains Q "full" circuits. This in turn requires the following three conditions to be met: a) Each lightpath must be filled; b) No circuit can travel over more than 2 lightpaths; c) Each lightpath must carry r full circuits.

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

$$\text{LPs} = \text{ADMs} \geq (2(N-1) + 2(N-1)(N-2))\frac{r}{g} = \frac{2(N-1)^2 r}{g}. \quad (2)$$

Note that the difference between the right-hand side of the bound in (2) and the bound in (1) is $(N-1)r - g$. Thus as noted above, when each node generates more than one wavelength worth of traffic, a single hub architecture cannot achieve the bound in (1).

The above considerations lead us to consider a multiple hub architecture, where all traffic is routed to one of K possible hubs. Assume each of the K hubs has a cross-connect capable of switching any circuit from any input wavelength to any output wavelength. Again, consider a unidirectional ring with N nodes, a traffic granularity of g and uniform traffic with r circuits between each pair. With K hubs (and N-K non-hub nodes), a total of $2(N\text{-}K)Kr$ circuits can be routed between the hubs and the non-hubs in one hop. The remaining traffic between the non-hub nodes, of which there are $(N\text{-}K)(N\text{-}K\text{-}1)r$ circuits, will traverse two lightpaths. Therefore, all traffic that is either to or from a non-hub node requires at least

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

lightpaths. Additionally we have to account for the traffic between hub nodes. By the same reasoning as used in deriving (1), this traffic requires at least $2K(K-1)r/(g+r)$ lightpaths. Hence the total number of lightpaths (and ADMs) is bounded by:

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

$$\text{ADMs} \geq \frac{2(N-K)(N-1)r}{g} + \frac{2K(K-1)r}{g+r} \tag{4}$$

Since the bound in (1) does not depend on the number of hubs, (4) can be tightened by combining it with (1). Doing this we have the following bound on the number of ADMs using a K hub architecture:

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

Some insight can be gained from examining the behavior of (5) as K, the number of hubs, varies. Notice that only the first quantity inside the maximization in (5) varies with K, we denote this quantity by $A(K)$. This quantity is the sum of two terms, the first of which is increasing linearly with K while the second term is decreasing quadratically. For

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

$A(K)$ can be shown to be decreasing in K, otherwise it is increasing. When $K=(N-1)r/g$, $A(K)=2N(N-1)r/(g+r)$, *i.e.*, the two quantities in the maximization in (5) are equal. Furthermore for $r \leq g$,

$$\left\lceil (N-1)r/g \right\rceil < 0.5\left[(N-1)(1+r/g)+1\right]. \tag{7}$$

Thus we have that for $r \leq g$, the number of hubs, K^*, that results in the smallest bound on the number of ADMs is given by $K^* = \lceil (N-1)r/g \rceil$; in other words, K^* is equal to the number of wavelengths of traffic generated by each node. When using K^* hubs, the lower bound in (5) is equal to the lower bound in (1). The above example, where $(N\text{-}1)r = g$, provides one case where this bound is tight using $K^* = 1$ hubs. As another example, consider the case where $r = g$, i.e. there is a full wavelength traffic demand between each pair of nodes. Setting up one lightpath between every pair of nodes is clearly the optimal way to route this traffic. This requires $N(N-1)$ ADMs, which meets the bound in (5) with $K^* = N$ hubs, i.e., each node is a hub. We note in this case however that no switching is required at the hubs.

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

4 K-Hub grooming algorithms

We consider several simple architectures and grooming algorithms using K hubs distributed around a ring with $N > K$ nodes. For the purpose of describing these algorithms, the exact location of the hubs is irrelevant. For the algorithms considered, all traffic originating at a non-hub node will be required to pass through a hub where it will be switched. In particular we do not consider architectures where traffic between non-hub nodes can be sent directly (without going through a hub). Our main reason for not allowing traffic to be sent in this way is to focus on architectures that are simple to design, implement and analyze.

4.1 Group algorithms

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

$$\text{ADMs} \leq \left\lceil \frac{\left(N - \lfloor N/K \rfloor\right)\lceil N/K \rceil r}{g} \right\rceil 2K + 2(N-K)\left\lceil \frac{(N-1)r}{g} \right\rceil \tag{8}$$

Notice that with this architecture, traffic between non-hub nodes in different groups needs to be switched at the hub for each group. Such traffic would then be carried over three lightpaths. As discussed in Sect. 3, this precludes such architecture from ever attaining the bound in (1). We consider a variation of this architecture where all traffic travels over at most two lightpaths. Specifically, assume that every node, including the hub nodes, now sends all traffic destined to any node in a group to the respective hub node.[5] The hub nodes once again distribute the traffic to the non-hub nodes in their group. Exact computation of the ADM requirement for this architecture is cumbersome because of the fact that *K* does not always divide *N*. Instead, we proceed with the following approximate, yet insightful, analysis. Assume each node sends 1/*K* of its total traffic to each hub (this assumption would be exact if all groups were of equal size). Hence each node sends $\lceil (N-1)r/Kg \rceil$ wavelengths of traffic to each of the *K* hubs. In addition, each hub node must send the groomed traffic to its subsidiary nodes. Each subsidiary node must receive a total of (*N*-1)*r* circuits using $\lceil (N-1)r/g \rceil$ wavelengths. Hence, each non-hub node generates $K\lceil (N-1)r/Kg \rceil$ wavelengths worth of traffic and receives $\lceil (N-1)r/g \rceil$ wavelengths. This can be accomplished using no more than $K\lceil (N-1)r/Kg \rceil$ ADMs at each non-hub node. Now, each hub node receives $\lceil (N-1)r/Kg \rceil$ wavelengths of traffic from each of (*N*-1) nodes and each hub nodes sends $(K-1)\lceil (N-1)r/Kg \rceil$ to the other hub nodes. Also, each hub node must send $\lceil (N-1)r/g \rceil$ wavelengths of traffic to each of its subsidiary nodes. Hence, the number of wavelengths sourced and terminated at each hub node is approximately the same and equal to $(N-1)\lceil (N-1)r/Kg \rceil$. Summing over all of the nodes, the total number of ADMs required is equal to

$$(N-K)K\lceil (N-1)r/Kg \rceil + K(N-1)\lceil (N-1)r/Kg \rceil = K(2N-K-1)\lceil (N-1)r/Kg \rceil \tag{9}$$

With this algorithm each circuit travels over at most 2 lightpaths, however notice that each non-hub node receives all its traffic from the corresponding hub; thus when every node generates more than 1 wavelength worth of traffic, each lightpath terminated at a non-hub node cannot contain *r* direct circuits, which is another requirement for the bound in (1) to be met. In the next section we consider an algorithm where each non-hub node sends and receives traffic from all of the hub nodes. This approach allows traffic to more closely emulate the characteristics for achieving the bound in (1).

4.2 Symmetric algorithm

In this algorithm, each non-hub node divides its traffic so that it sends approximately an equal amount to each of the *K* hubs. The traffic sent from a given non-hub node to a given hub will include traffic whose final destination is that hub as well as traffic for other non-hub nodes. The traffic for other non-hub nodes will be switched at the hub and forwarded to its

---

[4] The "hierarchical ring" proposed in [3] is similar to this type of architecture.

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

destination. Suppose that each non-hub node can divide its traffic to satisfy the following two conditions: i) No more than $H = \lceil (N-1)r/Kg \rceil$ wavelengths of traffic are sent to each hub from each non-hub node; ii) No more than H wavelengths of traffic are received at any non-hub node from any hub. If the traffic can be divided in this way each non-hub node will require at most KH ADMs and each hub node will require at most $(N-K)H$ ADMs for sending traffic to a non-hub node. Thus all traffic either to or from the non-hub nodes can be supported using at most $2K(N-K)H$ ADMs. Next we give one construction which shows that the traffic can indeed be divided to satisfy the above two conditions.

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

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

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

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

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

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

$$ADMs = 2K(N-K)\lceil (N-1)r/Kg \rceil + 2(K-1)\lceil (K-1)r/g \rceil. \tag{12}$$

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

The required number of ADMs in such a hierarchical architecture can be calculated recursively. Specifically, let $A(N,K)$ denote the minimum number of ADMs needed for an architecture with N nodes and K hubs, where traffic that originates at the non-hub nodes is routed as above, and inter-hub traffic is handled using a hierarchical architecture. Let $A^*(N) = \min\{A(N,K) \mid K \le N\}$ denote the minimum number of ADMs needed when the optimum number of hubs is used. Then assuming that the optimum number of "super-hubs" is used in the above architecture we have,

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

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

The results from using the symmetric algorithm (from (12)) are shown in Table 1 on the right hand side. The four columns on the right show the number of ADMs required when using one, two, three, and four hubs respectively. For example, in the case of a 17 node ring, the minimum ADM solution is achieved with 4 hubs. Highlighted in the table is the solution that achieves the minimum number of ADMs. Notice that this corresponds exactly to the number of wavelengths of traffic generated at each node (column 2); this is consistent with the analysis of the lower bound in Sect. 3.

Another benefit of the distributed hub architecture is that the size of the cross-connect used is reduced when compared to a single hub architecture. With only a single hub, each node would send $W = \lceil (N-1)r/g \rceil$ wavelengths to the hub and the hub would need a cross-connect of size (N-1)W x (N-1)W. However, when multiple hubs are used each node would only send a single wavelength to each hub and hence each hub would only need a cross-connect of size (N-1) x (N-1). This reduction in cross-connect size is significant because the size of the cross-connect has a dramatic impact on its cost.

	a) Lower bound						b) Algorithm			
N	W	K=1	K=2	K=3	K=4		K=1	K=2	K=3	K=4
5	1	8	8	8	8		**8**	14	16	14
6	2	13	12	12	12		20	**18**	22	22
7	2	18	17	17	17		24	**22**	28	30
8	2	25	23	23	23		28	**26**	34	38
9	2	32	29	29	29		32	**30**	40	46
10	3	41	37	36	36		54	66	**46**	54
11	3	50	46	44	44		60	74	**52**	62
12	3	61	56	53	53		66	82	**58**	70
13	3	72	67	63	63		72	90	**64**	78
14	4	85	79	73	73		104	98	136	**86**
15	4	98	92	86	84		112	106	148	**94**
16	4	113	106	99	96		120	114	160	**102**
17	4	128	121	114	109		128	122	172	**110**

Table 1. The number of ADMs needed with multiple hubs.

5 A bound on the average amount of switching needed

So far we have considered lower bounds on the number of ADMs needed and we have argued that if these bounds can be achieved then it will require that all circuits be switched at most once. Based on this we developed algorithms in Sect. 4 which can significantly reduce the required number of ADMs. However, in general there is no assurance that the bounds in Sect. 3 are tight; in which case it might be possible to further reduce the ADM requirement by allowing more switching in the network. In this section we provide some insight into this situation by developing an upper bound on the amount of switching needed in a network, in term if the number of ports in the network. Using this bound we show that reducing the number of ports inherently requires that traffic be only switched a small number of times.

Again we consider a WDM network with N nodes, where g low rate traffic streams are multiplexed onto each wavelength. Let T be the total number of ports in the network, where each lightpath is terminated by 2 ports (i.e., there are $T/2$ ADMs). Let C be the total number of low rate (unidirectional) circuits in the network. Now, let $f = T/C$, this is the average number of ports per circuit. Note that by setting up a point-to-point lightpath for each circuit, f can always be made equal to 2. Of course, more efficient grooming algorithms would lead to f being less than 2. For i=1,…,C, assume the ith circuit uses L_i ports, i.e., this circuit is sent over $L_i - 1$ lightpaths. Define $K = (1/C)\sum_i L_i$, so that K denotes the average number of ports used by a circuit. Finally, let S be the average number of times that a circuit is switched. We want to show that in a network architecture that minimizes the overall number of needed ports, T, each circuit needs to be switched on average only a small number of times. Note that for a given C, minimizing T is equivalent to minimizing f. Also, note that the number of times a circuit is switched is upper-bounded by $L_i/2$-1. (This is an upper bound, since the wavelength a circuit is on may be dropped at an intermediate node only to add/drop another circuit sharing that wavelength, but not switched.) Thus S is upper bounded by $K/2$-1. Since each port is shared by at most g circuits, we have that $CK/g \le T$, which implies that $K \le fg$. Hence,

$$S \le fg/2 - 1. \tag{14}$$

Thus for a given topology and traffic demand, any upper bound on f (or equivalently T) can be converted into an upper bound on S. This suggests that a topology that is efficient in the use of ports (small f) will not use much switching (small S). We consider some specific examples of this bound next.

For any topology and traffic demand, as noted earlier, $f \le 2$, substituting this into (14) we have

$$S \le g - 1. \tag{15}$$

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

$$S \le \frac{4\lceil (N-1)r/g \rceil (N-1)g}{2N(N-1)r} - 1 = 2\left\lceil \frac{(N-1)r}{g} \right\rceil \frac{g}{Nr} - 1 \le 2\left(\frac{N-1}{N}\right) + 2\frac{g}{Nr} - 1. \tag{16}$$

We emphasize that while (16) was developed by considering a single hub architecture, the bound applies to an architecture with an arbitrary number of hubs. Note that for any fixed g/r as N gets large, the upper bound on the average amount of switching in (16) approaches 1. Also note when g/r <1, then the right-hand side of (16) is less than one for any N. In other words, when each node generates more than a wavelength of traffic for each other node, the average amount of switching per circuit in an architecture that efficiently uses ports will be less than one. Of course, for the hub architecture S is less than one by design. However, the above tells us that any architecture that sought to further reduce the number of ADMs would not require more switching than the bound on S given in (16). Furthermore, instead of using a single hub architecture to bound the number of ports needed, a better bound on the number of ports can be found by using the bound from (13) for the symmetric multi-hub architecture. Figure 2 shows the resulting bound on S as a function of N, for a ring with g=4, r = 1. In finding this bound the optimal number of hubs were chosen for each N. Note that for all but 3 values of N this bound is less than one, indicating that each circuit needs to be switched at most once.

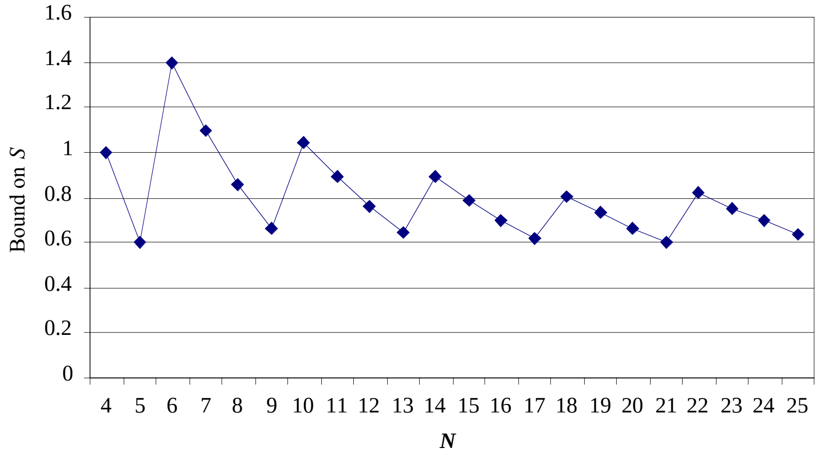

Figure 2. Plot of bound on S in (12) when f bounded using the symmetric grooming algorithm.

6 Conclusions

We show that using multiple hubs with cross-connects can reduce the number of ADMs needed in a SONET/WDM ring network. This architecture also reduces the size of the cross-connects used. Perhaps the greatest benefit of using this cross-connect approach is that it allows for the bandwidth to be provisioned dynamically.

The approach discussed in this paper can be extended in a number of ways. First, the symmetric hub architecture requires that each node send a portion of its traffic to one of the hub nodes and that all traffic goes through some hub. Both of these assumptions can be relaxed. Relaxing these assumptions can reduce the required number of ADMs but results in more complicated architectures. For example, for the 9 node ring example in Sect. 2 we can find an architecture with 5 hubs that requires only 48 ADMs (which is optimal), however each node only sends traffic to four of the hubs. Another generalization of the above model is to consider the case where a node only cross-connects a subset of a wavelengths that are dropped at that node. This generalization allows one to examine further trade-offs regarding where and how much cross-connect functionality is needed in the network.

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. Extended version appeared in *IEEE Journal of Lightwave Technology*, January, 2000.

[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, R. Ramaswami, and G. Sasaki, "Cost Effective Traffic 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, "Combined WDM and SONET Network Design," *Infocom '99,* New York, NY, March, 1999.