Switching and Traffic Grooming in WDM Networks

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

Eytan Modiano
MIT Laboratory for Information
and Decision Systems
Cambridge, MA 02139
modiano@mit.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 distributed among the nodes in the ring are needed for significant reductions in the number of ports. We characterize a class of network architectures that use both the minimal amount ADMs and the minimal amount of switching. We also give an algorithm for designing a topology that is efficient in both the required number of electronic ports and the amount of switching used.

1 Introduction

Wavelength Division Multiplexing (WDM) systems have increasingly been deployed to increase the capacity of optical networks. These networks typically have a SONET ring architecture. Nodes in such a 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. With WDM, multiple SONET rings can be supported on a single fiber pair; however, each additional ring will require additional ADMs. 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 a number of recent papers including [1-7]. 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]. Heuristic algorithms for general (non-uniform) traffic have also been presented in [6-7]. In much of the work on grooming, such as [1,2,6,7], 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 SONET rings (i.e. 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, even when the hub node is required to have an ADM on every wavelength. In [5] it was shown that the cost savings, in term of ADMs, with a single-hub architecture can be as high as 37.5 percent. In other work, such as [3], it is assumed that every node can cross-connect every wavelength that is dropped at that node. Clearly, more switching capability will not increase the required number of ADMs. However, there is a non-negligible cost associated with providing this electronic switching. Therefore, in addition to minimizing the required number of ADMs, it is also desirable to limit the amount of switching in the network.

In this paper, we consider architecures that are efficient both in terms of the number of ADMs used, as well as the amount of switching provided. In [8], it is shown that it is often possible to minimize the required number of ADMs in a network while only providing a limited amount of switching. In other words, additional switching capability does not lead to any further reduction in the number of ADMs. In [8], multiple-hub architectures were considered, where the nodes in the ring are divided into hub nodes and non-hub nodes. Each hub node can cross-connect every wavelength dropped at the hub while non-hub nodes have no DCXs. For such an architecture, with uniform traffic, it is shown in [8] that the optimal number of hub nodes is generally equal to the number of wavelengths of traffic generated by a node.

In this paper, we relax the assumption that each node is either a hub or a non-hub node, and we allow only a subset of the wavelengths dropped at node to be cross-connected. In this case, instead of a few hub nodes with complete switching capability, each node may have some partial switching capability. We give examples to show that such an architecture can result in both an efficient use of ADMs as well as a small switching cost. Next, we describe the basic ring model to be considered and discuss quantifying the switching cost of a ring. Following this we give an example to illustrate the approach we are considering. We then characterize cases where an "optimal" architecture can be found. Finally we give a heuristic algorithm for switching and grooming in a ring.

2 Ring Model

In this paper, we consider unidirectional ring networks, such as a UPSR SONET ring. This is done primarily to simplify our description; much of the following can be easily generalized to bi-directional rings and, in some cases, to arbitrary mesh networks. Let N denote the number of nodes in the ring, and assume these are numbered 1, 2,. . . , N. We assume that all traffic has the same granularity of g, *i.e.*, g low-rate circuits can be combined on each wavelength, and that there is a uniform traffic demand of $r \leq g$ low-rate circuits between each pair of nodes in the ring. We also assume that sufficient wavelengths are available so that any wavelength limitations can be ignored. For the above situation, a lower bound on the number of ADMs needed, regardless of the amount of switching, is given by the following expression [5]:

$$ADMs \geq \frac{2N(N-1)r}{g+r} \qquad (1)$$

In general, this bound is not tight, but it can be achieved in several cases. Some conditions needed for this bound to be tight are given in [8] and used to motivate the architectures presented there.

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

3 Example

Consider a unidirectional ring with $N=9$ nodes, a traffic granularity of $g=2$ and uniform traffic demand of $r=1$ circuit between each pair of nodes. In this case, from (1) we have a lower bound of 48 ADMs. First we consider supporting this traffic using a multi-hub architecture as in [8]. Each node generates 4 wavelengths worth of traffic. Using the symmetric hub architecture from [8], this traffic can be supported with 4 hub nodes and 50 ADMs. Each hub node receives one wavelength from each of the 5 non-hub nodes and must be able to switch circuits between these wavelengths. This requires a $5g \times 5g$ DCX. Therefore, the switching cost of this architecture is greater than[1] 4(100)= 400.

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

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

Notice that each pair of nodes is in exactly one of these groups. The traffic between all three nodes in each group can be supported by having two of the nodes send all of their traffic to the third node. A $2g \times 2g$ DCX at the third node can be used to switch the incoming traffic, which can then be forwarded to its destination. This requires 4 ADMs and a switching cost of $(2g)^2$= 16. Since there are 12 groups, supporting all of the traffic requires 48 ADMs and a total switching cost of 192. Notice that in this case we are using the minimum number of ADMs given by the bound in (1) and the switching cost is over 50 percent less than the cost for the symmetric multi-hub architecture. Also notice that any node within each group could serve as the "hub" for that group. For example, the switching capability could be spread out among all the nodes in the ring or concentrated at only 4 nodes.

4 Perfect Architectures

The distributed architecture in the above example meets the lower bound on the required number of ADMs from (1). This architecture can also be shown to have the smallest switching cost of any architecture that uses this number of ADMs. The proof of this follows from the characterization of architectures that achieve (1) given in [8]. We refer to such an architecture as *perfect*, i.e., a perfect architecture both meets the bound on ADMs in

---

[1] The actual switching requirements will be larger than this because we have not accounted for the switching required for inter-hub traffic.

(1) and has the smallest switching cost of all architectures that meet this bound. In this section we consider when the above example can be generalized to other cases, i.e., other values of *N*, *g*, and *r*. A sufficient condition for this is given in the following proposition.

Proposition 1: *A perfect architecture for a unidirectional ring with parameters N, g, and r can be found if the nodes in the network can be divided into groups of g/r + 1 nodes such that each pair of nodes is in at most one group.*

The resulting perfect architecture is a natural generalization of that in the preceding section. For the case of r=1, the above condition can be shown to be necessary as well as sufficient.

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

The problem of graph decompositions has been well studied in the graph theoretic literature and is related to combinatoric problems such as finding a block orthogonal designs or Steiner triple systems of a given order [9]. The next proposition provides a necessary condition for the existence of a K_M-decomposition of K_N.

Proposition 2: *If there exists a K_M-decomposition of K_N, then the following hold:*

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

Here we use the notation $a|b$ to denote that *b* is divisible by *a*. Furthermore, the above conditions can be shown to be sufficient for all but a finite number of values of *M* and *N* [9]. By combining the above arguments we have that unless $M-1|N-1$ and $M(M-1)|N(N-1)$, where $M=g/r +1$, a perfect architecture can not be found. Also, except for a finite number of values of *M* and *N* the above conditions are sufficient. Notice that for the example in Sect. 3 the above conditions are met.

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

$$N(N-1)(1-1/K) = N(N-1)\left(\frac{g^3}{g+r}\right). \qquad (2)$$

5 Algorithm

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

The algorithm forms groups of nodes and, for each group, a corresponding subset of the offered traffic. We try to form sets of all-to-all traffic between the pairs of nodes in a group. When this is not possible, we form a set that contains all the remaining traffic between each pair of nodes in the group, such that no more than one node in the group generates more than a wavelength worth of traffic. Furthermore, we try to form the largest groups that satisfy these properties. These sets are formed by adding nodes to a group one at a time and adding the corresponding traffic to the subset. We give more details of this algorithm next. To simplify the discussion we only describe the case where r=1. We maintain a list of the circuits originating at each node in the ring that have not yet been assigned to a subset. A list of the nodes in each group and the corresponding traffic subsets are also maintained. When a node is added to a group, all remaining traffic from that node to any other node in the group is added to the traffic subset.

Grouping Algorithm:

1. Choose as the first node in a group, a node with the maximal remaining circuits to be assigned.
2. Add nodes to the group until more than one node in the group has *g* circuits in the traffic subset or there are no other nodes with any traffic to send to the nodes in the group. Add nodes to the group sequentially; at each time adding the node that will result in the largest increase in the number of circuits in the corresponding traffic subset.
3. If all circuits have been assigned, stop. Otherwise go to 1.

We note that for the case when a perfect architecture exists, the above algorithm will divide the nodes into the groups given in Proposition 1.

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

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

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

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

The first subset of traffic requires 8 ADMs and a switching cost of $(4g)^2=256$. The second subset requires 7 ADMs and a switching cost of $(2g)^2=64$. Therefore, this architecture requires 15 ADMs and a total switching cost of 320. For comparison, the best multi-hub architecture from [8] will require 18 ADMs and a switching cost of 512. Some other examples for a ring with $g = 4$ and $r = 1$ are given in Table 1 below.

N	ADMs	Switching
6	15	320
7	22	656
8	26	912
9	34	1152
10	43	1504
11	55	1968
12	60	2592
13	64	2368
14	86	3040
15	98	3280

Table 1: Results of algorithm for ring with g=4, r=1 and N nodes.

6 Conclusions

In this paper, we discussed the role of switching in reducing ADM or port counts in WDM ring networks. We studied architectures where multiple nodes have some partial switching capability Such an architecture was shown to both reduce the needed number of ADMs and have a small switching cost. A heuristic algorithm for designing such architectures was also presented.

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

References

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

[6] H. Ghafouri-Shiraz, G. Zhu, and Y. Fei, "Effective Wavelength Assignment Algorithms for Optimizing Design Costs in SONET/WDM Rings," *IEEE Journal of Lightwave Technology*, vol. 19, no. 10, October, 2001.

[7] X. Zhang and C. Qiao, "An Effective and Comprehensive Approach for Traffic Grooming and Wavelength Assignment in SONET/WDM Rings," *IEEE/ACM Trans. on Networking*, vol. 8, no. 5, October, 2000.

[8] E. Modiano and R. Berry, "The Role of Switching in Reducing Network Port Counts," *Proc. 2001 Allerton Conf. on Communication, Control, and Computing*, Oct 3-5, 2001.

[9] J. Bosák, *Decompositions of Graphs*, Kluwer Academic, Boston, MA, 1990.