Traffic Grooming in WDM Networks

Eytan Modiano, MIT
Philip J. Lin, Tellabs

ABSTRACT

The recent emergence of wavelength-division multiplexing technology has led to a tremendous increase in the available transmission capacity in wide area networks. Consequently, these networks may no longer be limited by transmission bandwidth, but rather by the processing capability of electronic switches, routers, and multiplexers in the network. This realization has led to a new wave of research aimed at overcoming the electronic bottleneck by providing optical bypass at the WDM layer. Traffic grooming can be used as a bypass mechanism by which low-rate circuits are assigned to wavelengths in order to minimize the amount of electronic multiplexing equipment. Recently, this topic has received a significant amount of attention in both the research and commercial arenas. In this article we give an overview of the traffic grooming problem and survey some representative work in this area. While most recent work has focused on grooming in SONET rings, grooming traffic in general mesh networks is an important emerging problem.

INTRODUCTION

Wavelength-division multiplexing (WDM) is emerging as a dominant technology for use in backbone networks. WDM significantly increases the capacity of a fiber by allowing simultaneous transmission of multiple wavelengths (channels), each operating at rates up to 10 Gb/s. Systems with over 80 wavelengths are presently being deployed, and capacities that approach 1 Tb/s can be achieved. While such enormous capacity is very exciting, it also places a tremendous burden on the electronic switches and routers at each node that must somehow process all of this information. Fortunately, it is not necessary to electronically process all of the traffic at each node. For example, much of the traffic passing through a node is neither sourced at that node nor destined to it. To reduce the amount of traffic that must be electronically processed at intermediate nodes, future WDM systems will employ WDM add/drop multiplexers (WADMs), which allow each wavelength to either be dropped and electronically processed at the node or optically bypass the node electronics, as shown in Fig. 1.

Much of today's physical layer network infrastructure is built around synchronous optical network (SONET) rings. Typically, a SONET ring is constructed using fiber (one or two fiber pairs are typically used in order to provide protection) to connect SONET add/drop multiplexers (ADMs). Each SONET ADM has the ability to aggregate lower-rate SONET signals into a single high-rate SONET stream. For example, four OC-3 circuits can be multiplexed together into an OC-12 circuit, and 16 OC-3s can be multiplexed into an OC-48. The recent emergence of WDM technology has provided the ability to support multiple SONET rings on a single fiber pair. Consider, for example, the SONET ring network shown in Fig. 2a, where each wavelength is used to form an OC-48 SONET ring. With WDM technology providing dozens of wavelengths on a fiber, dozens of OC-48 rings can be supported per fiber pair using wavelength multiplexers (WDMs) to separate the multiple SONET rings. This tremendous increase in network capacity, of course, comes at the expense of additional electronic multiplexing equipment. With the emergence of WDM technology, the dominant cost component in networks is no longer the cost of optics but rather the cost of electronics.

The SONET/WDM architecture shown in Fig. 2a is potentially wasteful of SONET ADMs because every wavelength (ring) requires a SONET ADM at every node. As mentioned previously, not all traffic needs to be electronically processed at each node. A WADM at a given node is capable of dropping and adding any number of wavelengths at that node. Consequently, it is no longer necessary to have a SONET ADM for every wavelength at every node, but rather only for those wavelengths used at that node. Therefore, in order to limit the number of SONET ADMs required, the traffic should be groomed in such a way that all of the traffic to and from a given node is carried on the minimum number of wavelengths.

As a simple illustrative example, consider a unidirectional ring network (e.g., unidirectional path switched ring, UPSR) with four nodes. Suppose that each wavelength is used to support an

OC-48 ring, and that the traffic requirement is for eight OC-3 circuits between each pair of nodes. In this example we have six node pairs, and the total traffic load is equal to 48 OC-3s or equivalently three OC-48 rings. The question is how to assign the traffic to these three OC-48 rings in a way that minimizes the total number of SONET ADMs required. Consider, for example, the following two circuit assignments of traffic:

Assignment #1	*Assignment #2*
λ1: 1↔2, 3↔4	λ1: 1↔2, 1↔3
λ2: 1↔3, 2↔4	λ2: 2↔3, 2↔4
λ3: 1↔4, 2↔3	λ3: 1↔4, 3↔4

With the first assignment, each node has some traffic on every wavelength. For example, wavelength 1 carries the traffic between nodes 1 and 2 and that between nodes 3 and 4. Therefore, each node would require an ADM on every wavelength for a total of 12 ADMs. With the second assignment each wavelength contains traffic from only three nodes; hence, only nine ADMs are needed. Notice that both assignments carry the same amount of total traffic (eight OC-3s between each pair of nodes). The corresponding ADM allocations for both assignments are shown in Figs. 2a and b, respectively.

In a bidirectional ring the amount of electronics can be further reduced if proper routing and wavelength assignment (RWA) is performed on the groomed lightpaths. RWA is important to allow end-to-end lightpaths to share common ADMs [1]. In a SONET bidirectional line-switched ring (BLSR), the ADM is responsible for adding/dropping both the upstream and downstream data (Fig. 3). This is so the data in one direction can be switched to the opposite direction in case of a failure. Consequently, if an ADM has working traffic in the lightpath in one direction (e.g., upstream) and is not supporting traffic on the lightpath in the opposite direction (downstream), the ADM's capability is not fully utilized and the bandwidth in the unused direction is wasted. This is analogous to what is commonly called *stranded bandwidth* in BLSR, except it is occurring at the lightpath level.

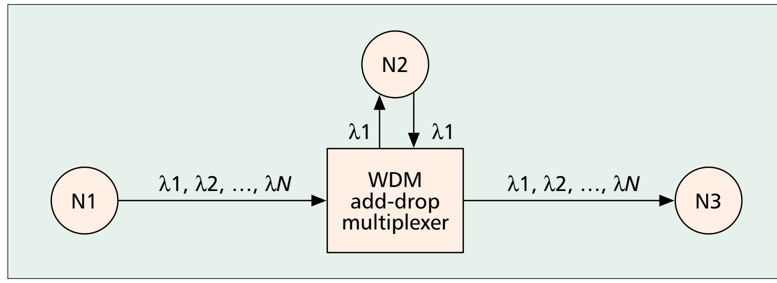

Figure 1. *Using a WADM to provide optical bypass.*

To illustrate the importance of RWA of groomed lightpaths, compare the following two RWAs of the same set of lightpaths, (*i*↔*j* indicates a bidirectional light path between nodes *i* and *j*):

RWA #1
λ1: 1↔2, 2↔3, 4↔5, 5↔6, 7↔8, 8↔9 (9 ADMs)
λ2: 1↔3, 4↔6, 7↔9 (6 ADMs)
Total = 15 ADMs

RWA #2
λ1: 1↔2, 2↔3, 3↔1 (3 ADMs)
λ2: 4↔5, 5↔6, 6↔4 (3 ADMs)
λ3: 7↔8, 8↔9, 9↔7 (3 ADMs)
Total = 9 ADMs

Both RWAs support the same set of traffic demands. The first RWA uses 15 ADMs and two wavelengths. The second RWA uses more wavelengths, but it only requires nine ADMs.

The above example illustrates a few characteristics of the overall problem of network cost minimization. First, the minimum number of ADMs is often not achieved with the minimum capacity usage. In the example, the method that uses the minimum number of ADMs requires an additional wavelength. Standard RWA algorithms that focus on minimizing the number of wavelengths

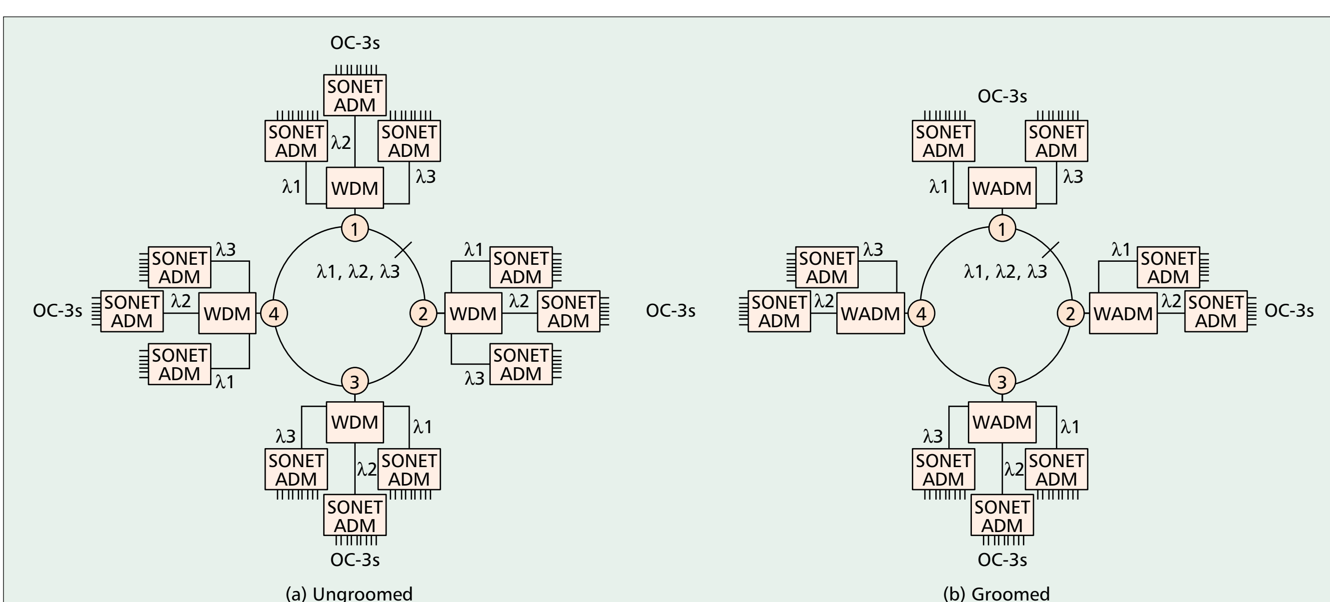

Figure 2. *SONET/WDM rings.*

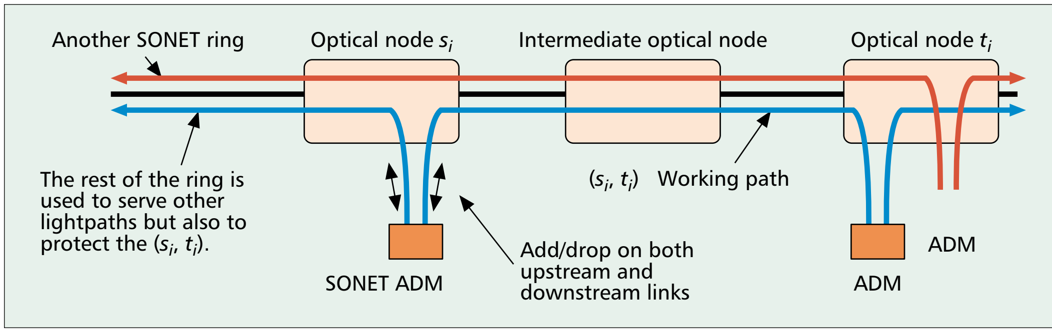

Figure 3. *ADMs in a SONET BLSR.*

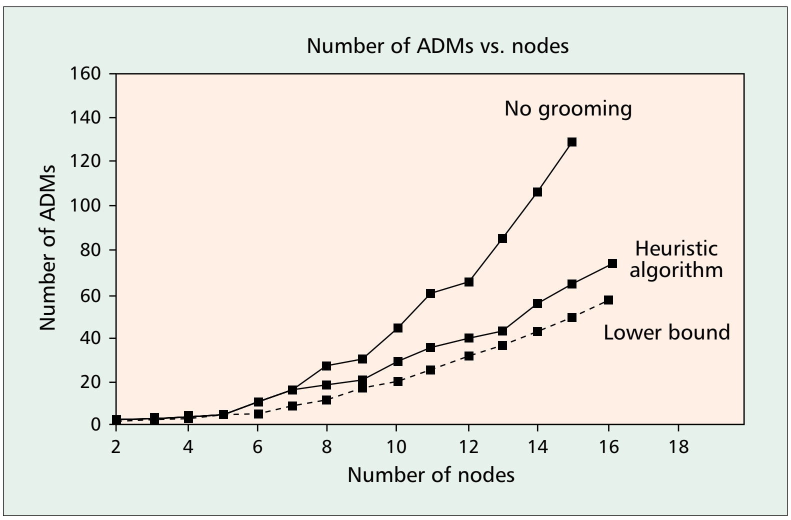

Figure 4. *ADM savings in a unidirectional ring network.*

cannot be directly applied to ADM cost minimization. Instead, algorithms that attempt to jointly optimize the cost of ADMs and wavelengths are needed. Second, the minimum number of ADMs is not achieved with shortest path routing. Since shortest path is desired to reduce network latency, a trade-off exists between network latency and ADM costs. Lastly, the RWA example shows that ADM saving is possible by appropriate RWA without the aid of grooming. This gives us two methods of reducing ADMs: grooming and RWA of groomed light paths. It would be tempting for a network planner to design the network in two steps: first, low-level grooming of tributaries onto lightpaths, and second, RWA of the resulting lightpaths.

Unfortunately, this two-step process will lead to a suboptimal solution. In fact, in [2] it was shown that an improvement of up to 20 percent could be achieved if the two steps are jointly considered in the design process.

Both grooming and RWA have the characteristic of grouping and packing problems. Such problems are often difficult. This intuitively explains why the ADM minimization problem is so complex. In fact, in [3] it was shown that the traffic grooming problem is NP-complete by showing that the bin packing problem can be transformed into the traffic grooming problem in polynomial time. Since the bin packing problem is known to be NP-complete, the traffic grooming problem must be NP-complete as well. As a result, many papers on grooming rely on heuristics and simulations to evaluate the heuristics.

Grooming Static Traffic

The static traffic grooming problem is a special instance of the virtual topology design problem. Given a traffic demand of low-rate circuits between pairs of nodes, the problem is to assign traffic to wavelengths in such a way that minimizes the number of ADMs used in the network. Virtual topology design problems can be formulated as a mixed integer programming problem which is known to be difficult. Heuristic algorithms have been developed to design virtual topologies that minimize the number of wavelengths, delays, or blocking probabilities. While the general topology design problem is known to be intractable, the traffic grooming problem is a special instance of the virtual topology design problem for which, in certain circumstances, a solution can be found. For example, [3] considers traffic grooming for a unidirectional ring, and [4] considers the same problem for a bidirectional ring. All show that significant savings in the number of ADMs can be achieved through efficient traffic grooming algorithms. For example, Fig. 4 shows the number of ADMs required when using the traffic grooming algorithm developed in [3] for the unidirectional ring with uniform traffic (a single OC-3 between each pair of nodes groomed onto an OC-48 ring). This number is compared to the number of ADMs required when no grooming is used (i.e., all wavelengths are dropped at all nodes). It is also compared to a lower bound on the number of ADMs. As can be seen from the figure, the algorithms developed in [3] are not far from the lower bound, and achieve significant ADM savings.

The algorithms in [3, 4] consider three different traffic scenarios:

- Uniform traffic in a unidirectional and bidirectional ring
- Distance-dependent traffic where the amount of traffic between node pairs is inversely proportional to the distance separating them
- Hub traffic where all of the traffic is going to one node on the ring

All of those cases yield elegant algorithms that are nearly optimal. The general traffic grooming problem with arbitrary traffic is of course much more challenging. As stated earlier, the general problem can be formulated as an integer program. However, these integer programs are typically very computationally complex and can only be solved for very small problems that are often impractical.

Zhang and Qiao [5] make an attempt at solving the problem by separating the problem into two parts. In the first part, the heuristic packs the traffic demands (e.g., OC-3s) into *circles*, where each circle has capacity equal to the tributary rate (OC-3) and contains nonoverlapping demands. As many circles as needed are constructed to include all traffic demands. The second part of the heuristic groups circles into wavelengths (e.g., 16 OC-3 circles in one OC-48 ring). Note that this algorithm is different than the two-step process mentioned in the previous section. There, the two steps are first grouping of tributaries into lightpaths, and then RWA of lightpath segments. Here, the two parts are first fitting tributaries onto a circle, and then grouping of the circles.

For this algorithm, the number of ADMs needed for a particular wavelength equal the number of *end nodes* involved. An end node is a node that terminates a connection in the circle. To minimize the number of ADMs, the heuristic attempts to match as many end nodes as possible when grouping the circles. This two-part algorithm can achieve good performance for uniform traffic as long as the grooming factor is reasonably large (e.g., OC-3s onto OC48 wavelengths). Even for nonuniform traffic, this two part algorithm performs reasonably well if a good end node matching algorithm is utilized.

Grooming Dynamic Traffic

Most earlier work on the grooming problem considered static traffic. Static traffic is common for many applications where a service provider designs and provisions network resources based on some estimate of the traffic. In many cases, however, the traffic changes over time. Such changes can be due to slow changes in traffic demands over a long period of time. More recently, such changes can be attributed to the more rapid dynamics of Internet traffic. It is therefore important to design networks that are able to efficiently accommodate changes in traffic.

In [6] the traffic grooming problem is generalized to encompass more general traffic models. The traffic is no longer restricted to be uniform, and nodes are allowed to have dynamically changing connections. The model in [6] defines a new class of traffic called t-allowable which allows each node to source up to t circuits. These t circuits can be destined to any of the nodes in the network without restriction, and the destinations of the circuits can be dynamically changed. The approach taken is to design a network so that it can accommodate any t-allowable traffic matrix in a nonblocking way. The problem is formulated as a bipartite graph matching problem, and algorithms are developed to minimize the number of wavelengths that must be processed at each node. These algorithms provide methods to achieve significant reductions in ADMs under a variety of traffic requirements.

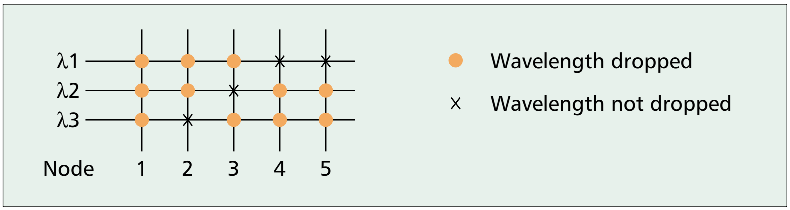

■ **Figure 5.** *Wavelength assignment for 2-allowable traffic.*

For example, suppose we have a ring network with five nodes, two circuits per wavelength, and $t = 2$. It can be shown that any 2-allowable traffic matrix can be supported with three wavelengths, dropped at each of the nodes as shown in Fig. 5 ($\lambda 1$ is dropped at nodes 1, 2, and 3, etc.). Consider now the 2-allowable traffic set, D, consisting of traffic streams {1–2, 1–3, 2–3, 4–5, 4–5}. This set can be supported 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. Notice that for the particular traffic set D, not all of the ADMs shown in the figure are needed. However, these ADMs are there in order to support other potential t-allowable traffic sets.

Grooming with Cross-Connects

Another approach for supporting dynamic traffic is to use a cross-connect at one or more of the nodes in the network. The cross-connect is able to switch traffic from one wavelength onto any other to which it is connected. Not only can the addition of a cross-connect allow for some traffic dynamics, but it can also be used to reduce the number of ADMs required. In [3] it was shown that using a hub node with a cross-connect is optimal in terms of minimizing the number of ADMs required and in [2] it was shown the cost savings can be as much as 37.5 percent. The proof in [3] is obtained by showing that any traffic grooming that does not use a cross-connect can be transformed into one that uses a cross-connect without any additional ADMs.

In [7] various network architectures with different amounts of cross-connect capabilities are compared. In one extreme, all traffic is routed to a hub node which has a cross-connect. The single-hub design has the lowest electronics cost since all the non-hub nodes need only terminate their own local traffic and optically bypass all other traffic. However, the single-hub network uses many wavelengths because all traffic is transported to the hub. On the other extreme, one can place cross-connects at all of the nodes and allow traffic to be switched and groomed at each node. This point-to-point WDM (PPWDM) design uses the minimum number of wavelengths but is the most expensive in terms of electronic processing cost. The cross-connect provides the flexibility for both the single-hub and PPWDM to accommodate dynamic traffic. At the expense of using more wavelengths, a hierarchical ring design that uses less electronics can be used to

In future IP networks, SONET ADMs may no longer be needed to multiplex traffic onto wavelengths. Instead, future IP networks will involve routers that are connected via wavelengths using WDM cross-connects.

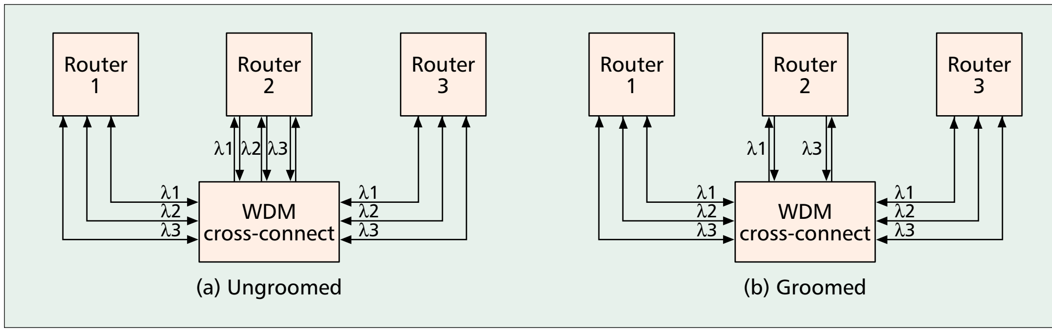

Figure 6. *Grooming in an IP/WDM architecture.*

simulate PPWDM. The idea is to designate some nodes as backbone nodes where all wavelengths are processed and other nodes as local nodes where only the local wavelengths are processed. With careful design, one can trade off electronic processing cost with the number of wavelengths.

Further cost savings are possible at the expense of less flexibility. For example, if the traffic demand is known (i.e., static traffic), [7] showed that a double-hub ring reduces the number of wavelengths required to half of that used by a single-hub ring. This is true for both uniform and nonuniform traffic. Also, if the traffic is incremental (i.e., there's only new arrivals but no departures), the number of wavelengths in the hierarchical ring can be reduced.

Furthermore, it was shown in [8] that additional savings could be obtained by distributing the cross-connect function among multiple nodes. These savings are significant in two ways. 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.

The above papers all conclude that 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.

Bounds on Grooming

One can derive a simple lower bound on the number of ADMs by noting that the number of ADMs required at each node must be large enough to take care of inserting and extracting the traffic into and out of the network. However, this rather obvious lower bound is very loose because typically traffic cannot be arranged in such a way that all of the traffic coming to a node can be carried on the minimum number of wavelengths. This is especially true if the demands have very fine granularity and a large number of demands are required to fill up one wavelength. An improved lower bound is obtained in [3] by considering the most efficient way in which traffic can be groomed onto a wavelength. This leads to an upper bound on the number of circuits that can be supported by an ADM, and consequently a lower bound on the number of ADMs required in order to support all of the traffic. The bounds developed in [3] are shown to be tight when traffic is uniform and no switches are employed (i.e., traffic cannot be switched from one wavelength onto another). To quantify the benefits of switching, [2] developed an improved lower bound that allows for cross-connects at each node to switch traffic between wavelengths. This improved lower bound is achieved by looking at the transit demands at each ADM. A transit demand visits an ADM, and therefore does not use ADMs as efficiently as those demands that optically bypass all intermediate nodes. The improved lower bound is achieved by counting the number of ADMs required to accommodate transit traffic. This improved lower bound is loose because it assumes that each transit demand visits at most one intermediate ADM. The authors in [9] describe a tighter bound for uniform traffic in bidirectional rings. The derivation is based on careful calculation of the bandwidth occupancy on each ring. The bound is shown to be tight by an α-approximation algorithm where the performance of the algorithm is within a factor of α from the lower bound. Unfortunately, the bound in [9] is better than the bound in [2] only if the network is large. Furthermore, the bound in [9] only applies for uniform traffic, whereas the bound in [2] applies for general traffic.

Grooming in IP Networks

In future IP networks, SONET ADMs may no longer be needed to multiplex traffic onto wavelengths. Instead, future IP networks will involve routers that are connected via wavelengths using WDM cross-connects, as shown in Fig. 6a. Since the SONET multiplexers have been eliminated, the function of multiplexing traffic onto wavelengths has now been passed onto the IP routers. Unless optical bypass is intelligently employed, with the new architecture, all of the traffic on all fiber and on all wavelengths (which amounts to multiple terabits) will now have to be processed at every IP router. Routers of this size and capacity far exceed any near-term prospects; and even when such routers could be built, they are likely to be very costly.

This situation can be alleviated through the use of a WDM cross-connect to provide optical bypass as shown in Fig. 6b. In order to achieve maximum efficiencies, one would need to bundle traffic onto wavelengths so the number of wavelengths that have to be processed at each router is minimized. This objective results in reducing both the number of ports needed on the routers (one per wavelength add/dropped at the router) as well as the total switching capacity of the router.

This problem is similar to that of grooming SONET streams described earlier. However, a number of important differences arise when considering the grooming of router traffic. First, unlike SONET networks, which are typically arranged in a ring topology, IP networks are arranged in a more general topology; hence, the earlier grooming results cannot be applied directly. Second, SONET circuits are typically provisioned well in advance and remain for very long periods of time. As a result, in the case of a SONET network, the traffic grooming problem can be solved in advance, and network equipment laid out accordingly. Most previous work on grooming for SONET rings considered particular traffic patterns (typically uniform traffic) for which a solution to the grooming problem was obtained. In the case of an IP network, not only is a uniform traffic pattern inappropriate, but the traffic patterns are highly dynamic; hence, a static solution would not be of much use.

Grooming in Mesh Networks

Most of the research on grooming has focused on ring networks. This is appropriate because today's backbone networks are organized in rings. However, due to the growth in Internet traffic, an increasing number of networks are being arranged in a general mesh topology. Therefore, there is a need to extend the grooming work to general mesh networks. In [10] an attempt is made to solve the general grooming problem by formulating it as a multicommodity network flow problem. Given a traffic matrix, each entry in the traffic matrix is viewed as a different commodity. A logical complete graph is created where the nodes represent the actual physical nodes, but the links in the logical graph represent potential lightpaths connecting the nodes. The flow problem is solved over the logical complete graph with the goal of minimizing the number of links used. Minimizing link usage is equivalent to minimizing electronic cost because each link represents a lightpath, and each lightpath requires the appropriate electronics for terminations and processing of the terminated traffic. Therefore, the new problem of traffic grooming is mapped into a well studied problem of multicommodity flow. Any known algorithm for solving multicommodity flow can now be applied to traffic grooming on the mesh. The solution to the flow problem gives the number of required lightpaths, and routing of these lightpaths (e.g., via shortest path) completes the solution.

Conclusions

Traffic grooming is an effective mechanism for providing optical bypass in WDM-based networks to save electronic processing cost. In this article we attempt to provide an overview of this important problem. Inevitably, due to space limitation, this overview is by no means complete. It is our hope that we provide the reader with sufficient background on this important problem that further work can be pursued.

Most early work on traffic grooming was focused on SONET rings, where traffic is often static and known in advance. Grooming in SONET rings is of critical importance to both equipment vendors and service providers. However, as networks evolve to become more IP focused, grooming for IP traffic will become an important area for future work. In the IP environment, however, traffic is typically neither static nor known in advance. Furthermore, as network architectures transition from rings to meshes, grooming in mesh networks will become an important extension to current ring-based algorithms.

Grooming in SONET rings is of critical importance to both equipment vendors and service providers. However, as networks evolve to become more IP focused, grooming for IP traffic will become an important area for future work.

References

[1] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength Assignment in a WDM Ring to Minimize the Cost of Embedded SONET Rings," *Proc. INFOCOM*, San Francisco, CA, Apr. 1998, pp. 94–101.

[2] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET Network Design," *Proc. INFOCOM*, New York, NY, Mar. 1999, pp. 734–43.

[3] E. Modiano and A. Chiu, "Traffic Grooming Algorithms for Minimizing Electronic Multiplexing Costs in Unidirectional SONET/WDM Ring Networks," CISS '98, Princeton, NJ, Feb., 1998; extended version appeared in *IEEE J. Lightwave Tech.*, Jan. 2000.

[4] J. Simmons and A. Saleh, "Quantifying the Benefit of Wavelength Add-Drop in WDM Rings with Distance-Independent and Dependent Traffic," *IEEE J. Lightwave Tech.*, Jan. 1999, vol. 17, pp. 48–57.

[5] X. Zhang and C. Qiao, "An Effective and Comprehensive Approach to Traffic Grooming and Wavelength Assignment in SONET/WDM Rings," *SPIE Proc. Conf. All-Opt. Networking*, vol. 3531, Boston, MA, Sept. 1998.

[6] R. Berry and E. Modiano, " Reducing Electronic Multiplexing Costs in SONET/WDM Rings with Dynamically Changing Traffic," *IEEE JSAC*, Oct. 2000.

[7] O Gerstel and R. Ramaswami, "Cost Effective Grooming in WDM Rings," *INFOCOM '98*, San Francisco, CA, Apr. 1998.

[8] E. Modiano and R. Berry, " Using Grooming Cross-Connects to Reduce ADM Costs in SONET/WDM Ring Networks," *Proc. OFC 2001*, Anaheim, CA, Mar. 2001.

[9] T. Y. Chow and P. J. Lin "The Ring Grooming Problem," submitted to *SIAM J. Discrete Math.*, available at Los Alamos ArXiv, http://www.arxiv.org/abs/math.OC/0101091

[10] V. R. Konda and T. Y. Chow "Algorithm for Traffic Grooming in Optical Networks to Minimize the Number of Transceivers," *IEEE 2001 Wkshp. High Perf. Switching and Routing*, to be published.

Biographies

EYTAN MODIANO [SM] (modiano@mit.edu) received his B.S. degree in electrical engineering and computer science from the University of Connecticut at Storrs in 1986, and his M.S. and Ph.D. degrees, both in electrical engineering, from the University of Maryland, College Park, in 1989 and 1992, respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992, and a National Research Council Post Doctoral Fellow during 1992–1993 while he was conducting research on security and performance issues in distributed network protocols. Between 1993 and 1999 he was with the Communications Division at Massachusetts Institute of Technology (MIT) Lincoln Laboratory where he designed communication protocols for satellite, wireless, and optical networks, and was the project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999 he has been a member of the faculty of the Aeronautics and Astronautics Department and the Laboratory for Information and Decision Systems (LIDS) at MIT, where he conducts research on communication networks and protocols with emphasis on satellite and hybrid networks, and high-speed networks.

PHILIP J. LIN (pjlin@ieee.org) is a staff research engineer at the Tellabs Research Center and a visiting scientist at MIT. He received a B.S. (1989) in engineering from the California Institute of Technology and an S.M. (1991) and a Ph.D. (1996) in electrical engineering from MIT. His Ph.D. thesis work is on all-optical backbone networks. His current research areas include optical network architecture, optical system design, and optical access networks.