Traffic Grooming in WDM Networks: Past and Future

Rudra Dutta and George N. Rouskas, North Carolina State University

Abstract

Traffic grooming refers to techniques used to combine low-speed traffic streams onto high-speed wavelengths in order to minimize the networkwide cost in terms of line terminating equipment and/or electronic switching. Such techniques become increasingly important for emerging network technologies, including SONET/WDM rings and MPLS/MPλS backbones, for which traffic grooming is essential. In this article we formally define the traffic grooming problem, and we provide a general formulation that captures the features of a wide range of problem variants. We then present a comprehensive comparative survey of the literature that unveils the significant amount of research on this subject (the traffic grooming past). We also offer a broad set of ambitious research directions (the traffic grooming future) that are motivated by the exciting new challenges arising with the advent of MPλS technology.

Over the last few years we have witnessed wide deployment of point-to-point wavelength-division multiplexing (WDM) transmission technology in the Internet infrastructure. The corresponding massive increase in bandwidth due to WDM has heightened the need for faster switching at the core of the network. At the same time, there has been a growing effort to augment the Internet Protocol to support different levels of quality of service (QoS). Label switching routers (LSRs) running multiprotocol label switching (MPLS) are being deployed to address the dual issues of faster switching and QoS support. On one hand, LSRs simplify the forwarding function, thereby making it possible to operate at higher data rates. On the other hand, MPLS enables the Internet architecture, built on the connectionless Internet Protocol, to behave in a connection-oriented fashion that is more conducive to QoS and traffic engineering.

The rapid advancement and evolution of optical technologies makes it possible to move beyond point-to-point WDM transmission systems to an all-optical backbone network consisting of optical crossconnects (OXCs). Each OXC can switch the optical signal coming in on a wavelength of an input fiber link to the same wavelength in an output fiber link. An OXC may also be equipped with converters that permit it to switch the optical signal on an incoming wavelength of an input fiber to some other wavelength on an output fiber link. The main mechanism of transport in such a network is the lightpath, which is a communication channel established between two LSRs (or other edge devices) over the network of OXCs, and may span a number of fiber links (physical hops). A lightpath is a generalization of the MPLS concept of a label-switched path (LSP) with the wavelength color corresponding to the label and the crossconnect matrix corresponding to the label forwarding table.

Optical networks in which OXCs provide the switching functionality are currently the subject of intensive research and development efforts. The Internet Engineering Task Force (IETF) is investigating the use of generalized MPLS (GMPLS) to set up and tear down lightpaths; a number of existing protocols are also being extended to interwork with GMPLS, including OSPF, RSVP, and CR-LDP. GMPLS is an extension of MPLS that supports multiple types of switching, including switching based on wavelengths, usually referred to as multiprotocol lambda switching (MPλS). With GMPLS, the OXC backbone and network of LSRs will share common functionality in the control plane, making it possible to seamlessly integrate OXC backbones within the overall Internet infrastructure. Optical networks have also been the subject of extensive research investigating issues such as virtual topology design, call blocking performance, protection and restoration routing algorithms and wavelength allocation policies, and the effect of wavelength conversion, among others.

With the deployment of commercial WDM systems, it has become apparent that the cost of network components, especially line terminating equipment (LTE), is one of the dominant costs in building optical networks, and is a more meaningful metric to optimize than, say, the number of wavelengths. Furthermore, with currently available optical technology, the data rate of each wavelength is on the order of 2.5–10 Gb/s, while channels operating at 40 Gb/s will be commercially available in the near future. In order to utilize efficiently this capacity, a number of independent lower-rate traffic streams must be multiplexed into a single lightpath. These observations give rise to the concept of *traffic grooming*, which refers to the techniques used to combine lower-speed components onto available wavelengths in order to meet network design goals such as cost minimization. Traffic grooming has received considerable attention recently, and, given the widespread use of SONET networks, it is not surprising that most of the early work has focused on ring topologies. The objective in all the studies is to minimize the network cost either directly, for example, by minimizing the amount of SONET add-drop multiplexers (ADMs), or indirectly, for example, by minimizing electronic switching costs.

This material is based on work supported by the National Science Foundation under grant no. 9701113.

0890-8044/02/$17.00 © 2002 IEEE

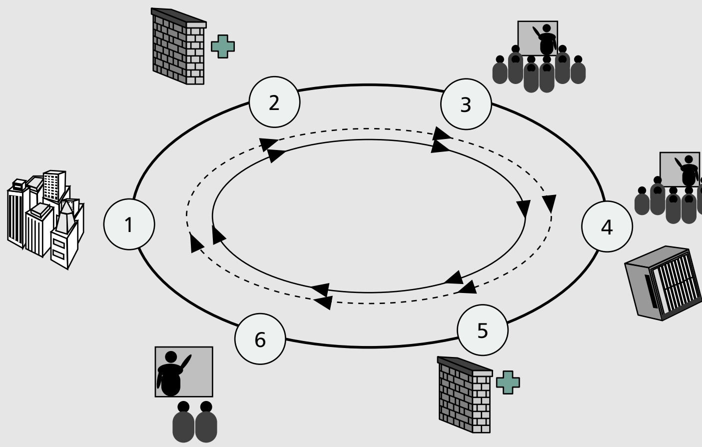

Node 1 — 6

1	0	0	0	250	50	0
	0	0	0	100	50	0
	0	0	0	0	0	100
	800	200	0	0	0	100
	200	600	0	0	0	0
6	0	0	100	100	0	0

■ Figure 1. *A six-node ring with traffic matrix, no optical routing, 5250 Mb/s electronic routing, and 12 ADMs.*

With the emergence of MPLS and MPλS technologies, service providers will eventually build a hierarchy of MPLS networks, with optical MPλS backbones occupying the highest level of the hierarchy. Such a hierarchy is feasible because the GMPLS standard makes it possible to tunnel a set of MPLS label switched paths (LSPs), with the same originating and terminating LSRs, through a network of OXCs by carrying these LSPs within the same lightpath. (Note that this aggregation of LSPs is a form of implicit label stacking with the top label being the color of the wavelength.) Consequently, a new problem arises in which the objective is to *groom a set of MPLS LSPs into lightpaths for transport over an MPλS network of a general topology* so as to minimize the overall network cost.

The rest of this article is organized as follows. We provide a practical example to illustrate the concept of traffic grooming and the trade-offs involved. We trace the conceptual evolution of the traffic grooming area, and provide an exact formulation of the problem. We present a comprehensive survey of the research related to the traffic grooming problem in WDM networks, and discuss some open problems and offer directions for future research. We then conclude the article.

Practical Example

As a practical illustration of the trade-offs involved in a traffic grooming situation, consider the six-node ring depicted in Fig. 1. For simplicity, the physical topology is assumed to be a unidirectional ring. The matrix indicates the amount of traffic (in megabits per second) that needs to be transferred among the ring nodes. This is a *static traffic* situation, and the matrix elements could represent the aggregate peak rate of traffic demands. For the purpose of the example, we also assume that the data rate of each wavelength channel is 1 Gb/s. Note that, since one node (node 4) originates an amount of traffic greater than the capacity of a single wavelength, at least two wavelengths are required to satisfy the traffic demands.

Figure 1 presents one possible way of carrying the traffic around the ring, in which both wavelengths are terminated at each node. The nodes are equipped only with switches that perform optical-electronic-optical (OEO) conversion of the traffic. This solution requires no optical routing of wavelengths, but it requires significant electronic switching capabilities at the ring nodes to switch traffic en route to its destination. The total amount of electronically switched traffic at all network nodes is 5250 Mb/s; this value is the sum, over all ring nodes, of the amount of traffic that passes through (but does not originate or terminate at) each node. This value is almost twice the total amount of traffic in the matrix. In other words, on average, each traffic component is switched at two intermediate nodes before reaching its destination. Also, this solution requires 12 synchronous optical network (SONET) ADMs, since each ring node must drop and then add both wavelengths. Since grooming usually refers to the appropriate use of hybrid optical and electronic routing, this is a "no grooming" solution.

Figure 2 shows another "no grooming" solution, which is at the other extreme. In this case, every traffic component travels on its *own* lightpath directly from source to destination. Wavelengths are optically routed at each intermediate node (e.g., by employing optical ADMs), and no electronic switching of traffic takes place other than at the source and destination nodes. However, this solution requires a large number of wavelengths, equal to the maximum number of traffic components that must travel over a link (in this case, six). In a realistic situation, these many wavelengths may simply not be available; and even if they are, this approach would still underutilize the network severely. Also, this solution does not reduce the amount of LTE, since it requires the same number of ADMs as the previous one. This result indicates that the elimination of electronic routing does not necessarily translate to a reduction in component cost.

Finally, Fig. 3 shows a solution in between the two extreme scenarios discussed above that uses an appropriate combination of optical routing and electronic switching to

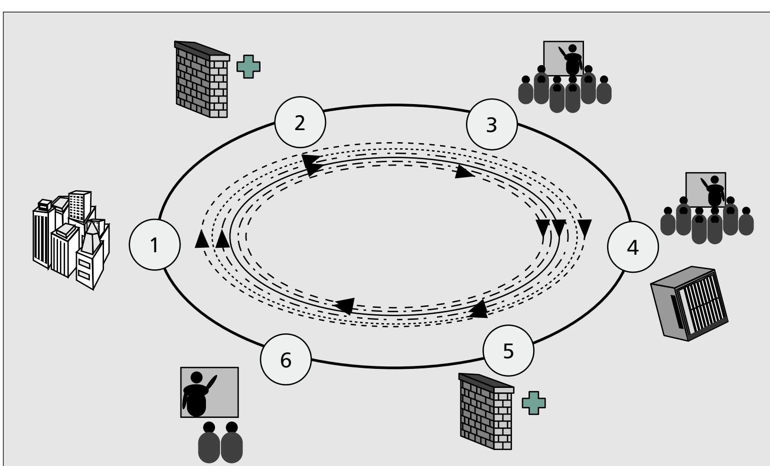

■ Figure 2. *All-optical routing, no electronic routing, and 12 ADMs.*

reduce the overall network cost while utilizing only two wavelengths. One wavelength is electronically terminated at each node, and all traffic from nodes 2, 3, 5, and 6 are carried on it, as well as traffic sourced by node 4 and destined for node 6. The traffic from node 1 to other nodes, and from node 4 to nodes 1 and 2, is put on the second wavelength. This wavelength terminates only at nodes 4 and 1, and optically bypasses all other ring nodes; some of the traffic from node 4 is electronically switched to the first wavelength at node 1. This solution requires only eight ADMs. It also achieves the minimum possible value of electronic switching with no more than two wavelengths, and is in this sense optimal.

This example clearly demonstrates that appropriate "grooming" of traffic components onto the lightpaths can translate into significant benefits in terms of network cost and/or number of required wavelengths, compared to the extreme solutions of point-to-point electronic switching and all-optical routing.

Basic Concepts and Formulation

The traffic grooming problem is essentially a variant of the well-known *virtual topology design* problem [1]. However, most of the virtual topology design studies do not address the traffic grooming or cost minimization aspects that are the subject of this work. In particular, it is either assumed that the set of lightpaths to be established is given or that the average traffic rate for all source-destination pairs is known, while in terms of the objective function, the performance metrics of interest are either the number of wavelengths or the maximum congestion level in the network. Consequently, these approaches cannot be applied directly in the context of SONET rings or MPLS/MPλS backbones where the objective is to groom a set of independent and discrete traffic components in order to minimize network equipment or electronic switching cost.

In this section we review the known results on the complexity of the *routing and wavelength assignment* (*RWA*) problem, the fundamental problem underlying the virtual topology design, and we formally define the traffic grooming problem. Before we proceed, we introduce some terminology. A *transparent node* (respectively, *opaque node*) in a WDM network is one in which all switching in the data plane takes place in the optical (respectively, electronic) domain. A translucent node is transparent with respect to some of the optical data channels (i.e., these channels are optically switched at the node) and opaque with respect to the others (i.e., these channels terminate at the node and undergo electronic processing and switching). We assume that the set of transparent/opaque channels is configurable. An OXC (respectively, OEO switch) is an example of a transparent (respectively, opaque) node, while a translucent node can be thought of as a combination of an OXC and an OEO switch. A *transparent network* (respectively, *opaque network*) consists exclusively of transparent (respectively, opaque) nodes, while a *translucent network* consists of a mix of transparent, opaque, and translucent nodes. Note that a transparent or opaque node is a special case of a translucent node, and a similar relationship exists between transparent, opaque, and translucent networks.

The RWA Problem

From a graph theory perspective, an optical network is viewed as a graph $G = (V, E)$ where each edge represents an optical fiber link between its endpoints. A request $r = [u, v]$, for a connection from node u to node v, is *satisfied* by
- Assigning to r a path p_r of links in G from u to v
- Assigning to p_r a wavelength to carry the information over the links of p_r (the path p_r and associated wavelength constitute a *lightpath*)

Consider a multiset R of connection requests (a multiset is a set in which the same element can occur two or more times). Then the RWA problem is to satisfy all requests in R in such a way that if two requests r and r' have paths p_r and $p_{r'}$, respectively, that share a common edge, they are assigned different wavelengths. The goal is to satisfy all requests in R using the minimum number of wavelengths. The RWA problem is defined on transparent networks and, in its pure form, assumes no wavelength conversion (i.e., the *same* wavelength is assigned on all links along the path p_r of request r).

First note that if the network G is a tree, every pair of nodes is joined by a unique path, so part (i) of the RWA problem is solved. It is known then that wavelength assignment to minimize the number of wavelengths can be solved in polynomial time in paths, stars, and spiders (a spider network is a tree with exactly one vertex of degree greater than two), but that it is NP-hard in general trees. If G is a path, part (ii) is equivalent to the interval graph coloring problem, which can be solved in linear time by a greedy algorithm [2]. If G is a star, part (ii) is equivalent to finding a minimum edge coloring in a bipartite graph, which is solvable in polynomial time by combining theorems of Hall and Konig [3]. Minimizing the number of wavelengths is also polynomial in spider networks [4]. If G is a general tree, the problem of minimizing the number of wavelengths is known to be NP-hard, even in special cases such as binary trees [5]. However, given any set of requests in a tree network, there is a polynomial time algorithm to produce a wavelength assignment that is no more than 5/2 optimal [6]. In ring networks, a request r can be routed in two different ways, so both parts (i) and (ii) must be addressed. This problem is known to be NP-hard [4]. On the other hand, given any set of requests in a ring network, it is possible in polynomial time to find a set of paths and corresponding wavelengths to route the requests, so the number of wavelengths used is no more than twice optimal [4].

The RWA problem is NP-hard in general topologies, and remains so even in special cases,

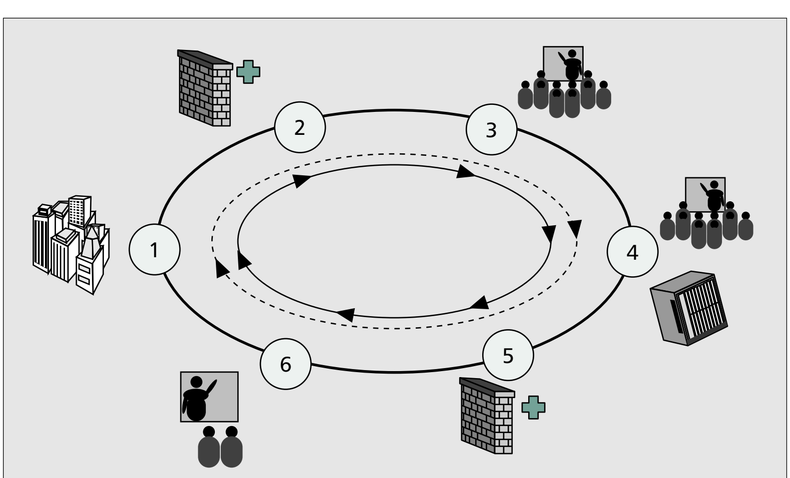

■ Figure 3. *Best grooming solution: 2450 Mb/s electronic routing and 8 ADMs.*

for example, when each lightpath traverses no more than two (physical) hops [7]. However, the RWA problem becomes easier when some degree of wavelength conversion is allowed at (some of) the OXCs. For instance, it is known that even when a single node in a ring network is capable of full wavelength conversion, the number of wavelengths required is equal to the congestion (i.e., the maximum number of lightpaths traversing any link) [8]. In fact, it is shown in [8] that very limited conversion may provide significant improvements in the traffic-carrying capacity of ring, star, and tree networks, as well as general topologies. Finally, recent results demonstrate similar improvements in multifiber networks [9, 10]; note that a multi-fiber network is equivalent (from the RWA point of view) to one in which each OXC is capable of *limited* conversion, similar to the model in [8].

The Traffic Grooming Problem

In order to utilize bandwidth more effectively, new models of optical networks allow several independent traffic streams to *share the bandwidth* of a lightpath. If the multiplexing and demultiplexing of lower-rate traffic components is performed at the boundaries of the network only (i.e., at LSRs), and the aggregate traffic transparently traverses the MPλS network, this problem is equivalent to RWA. However, it is in general impossible to set up lightpaths between every pair of LSRs (e.g., due to wavelength constraints or constraints on the number of optical transceivers at each LSR). Therefore, it is natural to consider optical networks with translucent or opaque nodes at which traffic on terminating lightpaths is electronically switched (groomed) onto new lightpaths toward the destination node. Introducing some amount of electronic switching within the MPλS network has two advantages: it significantly enhances the degree of *virtual connectivity* among the edge LSRs, while at the same time it may drastically reduce the wavelength requirements within the optical network for a given traffic demand. The trade-off is an increase in network cost due to the introduction of expensive active components (i.e., optical transceivers and electronic switches). These observations motivate us to define the following traffic grooming problem.

Let C be the capacity of each wavelength. A traffic pattern is represented by a demand matrix $T = [t^{(sd)}]$ where $t^{(sd)}$ denotes the number of traffic streams (each of unit demand) from node i to node j. Given matrix T on network G, the traffic grooming problem involves the following conceptual subproblems (SPs):

- *Virtual topology SP*: find a set R of lightpath requests
- *Lightpath routing and wavelength assignment SP*: solve the RWA
- *Traffic routing SP*: route each traffic stream through the lightpaths in R

The goal is to minimize a cost function defined on the set R of lightpaths. In this context, the number of wavelengths per fiber link is taken into consideration as a constraint rather than as a parameter to be minimized.

Cost Functions — The following three cost functions are relevant in a traffic grooming scenario.

Total number of lightpaths. In this model, the cost to be minimized is the cardinality of the set R of lightpaths. This cost model is motivated by the observation that each lightpath requires a transceiver at its origin and terminal node, and that terminating a lightpath requires sufficient capacity for switching its constituent traffic streams. An alternative performance measure is the total amount of LTE, that is, the number of SONET ADMs for a ring network, or the number of OXC ports for an arbitrary topology. This cost function captures both the LTE and electronic switching costs involved in satisfying a given demand matrix.

Total amount of electronic switching at all network nodes. Rather than counting a unit of cost each time a lightpath is terminated at a translucent or opaque node, this model counts instead the number of traffic streams carried by such a lightpath. Consequently, it is a more accurate measure of the amount of electronic switching inside the network, but it only indirectly captures the LTE (transceiver) cost.

Maximum number of lightpaths terminating/originating at a translucent node. This cost model is appropriate for minimizing the electronic switching or LTE cost not over all network nodes, but at the node where it is maximum. Such a min-max objective is of great practical interest because it sets a limit on the equipment (hence cost) required at a single node.

As defined, the traffic grooming problem is a generalization of a number of variants that have been studied mainly in the context of ring networks. Our definition assumes no wavelength conversion. However, wavelength conversion can be modeled as a special case of electronic switching whereby *all* traffic streams on a given lightpath (and *only* those) are switched onto a new lightpath (i.e., *no* grooming). Therefore, models that capture the cost of wavelength conversion can be easily incorporated into the formulation. In what follows, we assume that it is not necessary to explicitly capture the cost of wavelength conversion unless specifically mentioned otherwise.

Formulation

We now provide an exact integer linear program (ILP) formulation of the grooming problem following our own work [1, 11]. The formulation is provided for the sake of completeness, but also because we feel it allows a better understanding both of the subproblems we identified above as well as their coupled nature. As such, we have concentrated not on the most efficient formulation but rather on the illustrative value of it.

We make the assumption that the physical links and the lightpaths are directed, and each link is composed of a single fiber. We use the following notation: b_{ij} is the integer *lightpath count* from node i to j, $b_{ij}(l, m)$ is the number of such lightpaths that traverse the physical link from node l to m. The *link lightpath wavelength indicator* $c_{ij}^{(k)}(l, m)$ is 1 if such a lightpath uses the wavelength k over the physical link from l to m, 0 otherwise. The *link indicator* p_{lm} indicates whether a physical link exists in the physical topology from node l to m. The other variables relate to traffic; they are all integers since we assume that traffic demands and allocation are all quantized, and can be expressed as multiples of some basic rate. The bandwidth of a single wavelength channel is denoted by C in terms of that basic rate. The demand on the network are embodied in the input parameters $t^{(sd)}$, which is the amount of traffic that needs to flow from node s to d. Collectively, these form the traffic matrix $T = [t^{(sd)}]$. The lightpath(s) from node i to j carry aggregate traffic t_{ij}, and $t_{ij}^{(sd)}$ denotes the amount of this traffic that is due to the demand $t^{(sd)}$. The domain for each of i, j, s, d, l, m is $\{0, \ldots, N-1\}$, where N is the number of nodes in the physical topology. We also assume that the number of wavelength channels W supported by each directional physical link is given as a parameter. The domain for k is $\{0, \ldots, W-1\}$. Accordingly, we formulate the traffic grooming problem as follows:

Minimize: (one of the following functions)

Total number of lightpaths: $\Sigma_{i,j}\, b_{ij}$ (1)

Total amount of electronic switching:

$$\Sigma_{s,d,i,j}\, t_{ij}^{(sd)} - \Sigma_{s,d}\, t^{(sd)} \quad (2)$$

Maximum number of lightpaths at a node:

$$\max_i(\max(\Sigma_j\, b_{ji}, \Sigma_j b_{ij})) \quad (3)$$

Subject to:

- **Physical Topology Constraints**:

$$b_{ij}(l, m) \le b_{ij} p_{lm}, \forall i, j, l, m \quad (4)$$

$$c_{ij}^{(k)}(l, m) \le p_{lm}, \forall i, j, k, l, m \quad (5)$$

- **Lightpath Routing SP Constraints**:

$$\sum_{l=0}^{N-1} b_{ij}(m,l) - \sum_{l=0}^{N-1} b_{ij}(l,m) = \begin{Bmatrix} b_{ij}, & m=i \\ -b_{ij}, & m=j \\ 0, & m \ne i, m \ne j \end{Bmatrix} \forall m,i,j \quad (6)$$

$$\sum_{i,j} b_{ij}(l,m) \le W, \forall l,m \quad (7)$$

- **Lightpath Wavelength Assignment SP Constraints**:

$$\sum_{k=0}^{W-1} c_{ij}^{(k)}(l,m) = b_{ij}(l,m) \quad \forall i,j,l,m \quad (8)$$

$$\sum_{i,j} c_{ij}^{(k)}(l,m) \le 1 \quad \forall k,l,m \quad (9)$$

$$\sum_{l=0}^{N-1} c_{ij}^{(k)}(m,l) - \sum_{l=0}^{N-1} c_{ij}^{(k)}(l,m) \begin{Bmatrix} \le b_{ij} & m=i \\ \ge -b_{ij} & m=j \\ =0, & m \ne i, m \ne j \end{Bmatrix} \quad \forall i,j,k,m \quad (10)$$

- **Traffic Routing SP Constraints**:

$$t_{ij} = \sum_{s,d} t_{ij}^{(sd)}, \forall i,j \quad (11)$$

$$t_{ij} \le b_{ij} C, \quad \forall_{ij} \quad (12)$$

$$\sum_{j=0}^{N-1} t_{ij}^{(sd)} - \sum_{j=0}^{N-1} t_{ji}^{(sd)} = \begin{Bmatrix} t^{(sd)} & i=s \\ -t^{(sd)} & i=d \\ 0, & i \ne s, i \ne d \end{Bmatrix} \quad \forall m,i,j \quad (13)$$

Most of the above constraints are self-explanatory. We note that constraint (7) is implicit in constraints (6), (8), and (9), but we chose to include it for illustrative purposes. Constraint (7) would also be needed if full wavelength conversion capability were available at each node of the network, in which case the entire wavelength assignment SP would be eliminated. An important feature of the formulation is that the wavelength continuity constraint (10) can be relaxed on a per-node basis, making it easy to model the more realistic situation whereby some nodes possess wavelength conversion capability while others do not.

A Survey of Traffic Grooming Research in WDM Networks

The importance of traffic grooming in optical WDM networks is underscored by the recent surge in interest within the research community, and the emergence of network technologies (especially GMPLS) for which grooming techniques are essential. A recent special issue of the *Optical Networks Magazine* [12] investigated a number of aspects related to traffic grooming, from fairness considerations to the role of wavelength conversion, and from reconfiguration issues to the digital format mechanisms of carrying groomed traffic in higher-speed channels. In this section we present a comprehensive survey of the traffic grooming literature, focusing on research that addresses the problem of minimizing component cost in optical WDM networks.

The vast majority of the traffic grooming research deals with ring networks, with only few recent studies focusing on non-ring topologies. Accordingly, we present the literature on ring networks in the following section, while the research on other topologies is discussed in the (significantly shorter) next section. In Table 1, we classify the context of the ring network studies along two dimensions. Along one dimension we categorize the traffic scenarios considered, and along the other we indicate whether full, partial, or no interchange of traffic is allowed between the wavelengths around the ring; that is, whether SONET crossconnects are employed everywhere in the ring, at certain nodes only, or not employed at all. The detailed discussion below clarifies these concepts. There is necessarily some approximation involved in the classification; our purpose is to indicate general characteristics of the studies in helping a reader locate the appropriate ones.

Traffic Grooming in Ring Networks

Static Grooming Model, Uniform Traffic Pattern

No Traffic Exchange Among Wavelengths — The objective of [13, 14] is to minimize the number of ADMs or OXC ports in the whole network; this objective is related to the first cost function we specified above. In [14], this is posed as a reduction in router sizes, both in terms of the number of ports as well as physical size. These objectives are expressed as a *through-to-total* ratio of connections at each node in the network: the larger the number of connections that can be routed optically at each node, the higher the through-to-total fraction, and the smaller the number of ADMs required. While [13] only considers ring topologies, [14] also considers a Manhattan street network physical topology. Two traffic patterns are considered: the uniform all-to-all pattern, and a distance-dependent pattern similar to the "falling" pattern of [11]. In this work, as well as in other studies discussed later, the wavelength assignment subproblem of the RWA problem is addressed implicitly rather than explicitly: because of the approach taken in placing ADMs on each wavelength used, there is no chance to form lightpaths that cannot be colored.

A number of strategies are presented in [13] for the case $C = 1$, that is, when each traffic component occupies a whole wavelength, as well as for the special cases of $C = 2$, 4, 8, 16, under a uniform traffic scenario. The key concept introduced is that of the *super node*, which is used when the capacity C is some integer greater than 1. A super node is an aggregation of nodes, and is chosen so that traffic between super nodes becomes equal to the bandwidth that fills one wavelength. This is a heuristic approach to the grooming problem, and the benefits that can be obtained be using this approach are approximately quantified. In [14], the results of using the super node approach for the mesh network are also presented.

As the authors remark, it should be possible to improve this approach by grooming traffic within super nodes as well. This becomes more important as the number of nodes increases without a corresponding increase in the number of wave-

Grooming model	Traffic pattern	Wavelengths may exchange traffic		
		No	At hub node only	At all nodes
Static	Uniform all-to-all	[13, 14, 15]	[16]	[17, 18]
	Other general	[19–25]	[24, 26]	[11]
Dynamic	General	[27]		[18]

Table 1. *Classification of traffic grooming studies for ring networks.*

lengths, and *C* increases; in this case, more and more nodes must be grouped into super nodes. Thus, it is not clear how this approach will scale. In a sense, this approach is somewhat complementary to the approach taken in [11], which is described later. As we mention later in this section, [16] provides another approach related to grouping the nodes in larger units.

The objective of the study in [15] is to create not a single, but multiple virtual topologies in a WDM bidirectional ring, which together satisfy the traffic demands of node pairs. These topologies are then implemented in a sequence, determined by a schedule of topologies. The objective is to obtain as short a schedule as possible. The scheduling offers a way to aggregate traffic streams over time into the available wavelengths, and thus has an applicability to grooming problems. From the grooming point of view, the main contribution of the study is a method of constructing *circles*, that is, a set of lightpaths which occupy the whole ring at a single wavelength and take care of all the internode traffic of a set of nodes. These are the topologies that are then combined by scheduling. The circle construction problem is simplified somewhat due to the assumption of the uniform all-to-all traffic assumption. Two methods of constructing circles are presented, one for a ring consisting of an even number of nodes and the other for a ring of an odd number of nodes. The strategy is to systematically and exhaustively group nodes to form circles by defining a variable called *stride* that counts the number of links between nodes in the group, and iterating over this variable. For rings with an even number of nodes, each group has either two or four connections, and for rings with an odd number of nodes, each group has three or four connections.

In [25], an approach of forming circles similar to that in [15] is taken, but now the focus is explicitly on grooming, and a scheduling approach is not taken. The objective is to minimize either the number of wavelengths used or the number of SONET ADMs. Both uniform (all-to-all) and nonuniform traffic is considered. For uniform traffic, the algorithms of [15] are used. A heuristic algorithm is proposed to form circles in the case where the traffic matrix is general. This algorithm is based on maintaining all desired connections that cannot be groomed into complete circles at any given time in a "Gap Maker" list, and attempting to fit elements from this list into existing partial circles such that the number of ADMs (or wavelengths) is minimized. Algorithms to groom these circles onto wavelengths, and to find lower bounds on the number of ADMs needed are also provided; for the uniform traffic case, these bounds are tight.

Wavelengths May Exchange Traffic —The study in [17] concentrates specifically on the SONET architecture. One of the important features of this study is a description of the relevant SONET components and a detailed cost model that accounts for these network components. The authors then abstract a mathematical model of the ring network they consider, in which they focus on the number of SONET ADMs as the objective to be minimized. The uniform all-to-all traffic model is considered, and a single-hub design is proposed. Unlike the studies discussed earlier in this section, the authors allow SONET crossconnects. They demonstrate that a solution involving digital crossconnection that allows traffic components to change wavelengths at the nodes where they are electronically switched, will outperform one in which this capability is absent. They also provide an example to show that an integrated design approach performs better than the two-step approach in which the virtual topology subproblem and the traffic routing subproblem are decoupled. An extension to the case where the same network contains SONET rings of different rates is considered. In such a network, some of the wavelengths are implemented with SONET equipment capable of lower rates than others. If the additional bandwidth is not needed, this approach reduces cost. Networks with uniform (but different) SONET rates are compared to networks with a mix of slow and fast SONET equipment.

The focus of the study in [16] is not in using realistic traffic models, but rather in gaining insights by using simplistic ones. Since traffic is allowed to change wavelengths only at the hub node, wavelength assignment is implicit (similar to [13, 14]). This study proves the interesting result that even in the case of the very restrictive traffic pattern of a single egress node (all nodes source traffic destined to only one destination node), the grooming problem remains NP-complete under the assumption that no SONET crossconnects are used. An all-to-all uniform traffic pattern and some cross-traffic patterns are also considered, and lower bounds on the number of ADMs are derived. It is demonstrated that the number of wavelengths and the number of ADMs cannot always be simultaneously minimized. There is some treatment of using a SONET crossconnect, but only a single node with such a crossconnect is considered. Two heuristic algorithms are proposed. One of the algorithms attempts to group the nodes of the ring such that each group has one or more wavelengths dedicated to it. The problem then becomes one of carrying the intergroup and intragroup traffic, and intelligent group formation leads to whole wavelengths packed tightly. This approach is interesting because it may have wider applications beyond the context in which it is presented. The grouping of nodes is a technique also encountered in [13, 14] (which we have already discussed), but the approach in this study is more general since the nodes forming the groups do not need to be contiguous on the ring.

Static Grooming Model, General Traffic Patterns

No Traffic Exchange Among Wavelengths — The study in [20] does not strictly address a grooming problem, since it addresses only the wavelength assignment subproblem; it is assumed that lightpaths have already been formed and routed as well. However, wavelength assignment is a part of the overall grooming problem, and this study also focuses on how the number of ADMs (hence, network cost) can be reduced by proper combining of lightpaths. It thus complements some other work in the grooming area, especially the work by some of the same authors in [18] (which we discuss below). The basic model is again that of a SONET ring in each wavelength, with the objective being to minimize SONET ADMs. A static traffic model is assumed, but otherwise the traffic model is not detailed; this is reasonable since the topology problem is already assumed solved. The lightpaths are not assumed to have any particular forms, which can be seen as a lack of assumption regarding specific traffic characteristics.

The authors make the important distinction between minimizing the number of lightpaths and the number of ADMs in the context of a ring. The two are similar if the resulting ring is fully loaded with lightpaths for all the wavelengths used, but otherwise there could be different solutions using the same number of lightpaths that use very different number of ADMs. The authors describe the concept of *lightpath splitting*. This is the term for terminating a lightpath at an intermediate node before it reaches its destination, then creating a new lightpath starting at that node to reach the final destination of the original lightpath. Naturally, this increases the number of required ADMs directly, but if this is carried out at strategic points, the resulting new lightpaths may be able to share ADMs, thus reducing the total number. Thus, the distinction between number of lightpaths and number of ADMs is directly used. However, this technique is likely to be useful only when the ring is lightly loaded, or (equivalently) the number of wavelengths available can be increased at will; this reduces the usefulness of this approach to some degree. Examples are provided to show that it may not be possible to minimize the number of wavelengths utilized and the number of ADMs simultaneously. Two heuristics are then developed, Cut-first and Assign-first. The Cut-first heuristic is the extension of a previously known heuristic (in a different context) that splits all the lightpaths passing through some single node, and then assigns wavelengths. The authors enhance this heuristic by choosing the unique node and the means of assigning wavelengths to the lightpaths so that the number of ADMs is reduced. The Assign-first heuristic picks a node to start at and colors the lightpaths greedily. A comparison of the two heuristics is provided. Three transformations are provided that can be used as post-assignment fine tuning (regardless of the method used to assign the wavelengths), but it is not clear that they can always be used for significant improvement in a computationally tractable manner.

As we mentioned above, [20] does not consider lightpath routing as part of the problem. In [21], a problem similar to that in [20] is considered, but the problem of routing the lightpaths is considered as well as the wavelength assignment problem. An Euler-trail decomposition of the ring network is presented, and a heuristic that uses one of the heuristics of [20] and performs rerouting of lightpaths as well, based on the Euler-trail decomposition, is presented. Several theoretical results are proved in designing the algorithm. The time complexity of the heuristic is not discussed; in particular, the time complexity of a crucial step (the partitioning of the graph into nondecomposable trails) is not clear. The authors show that the simultaneous consideration of lightpath routing can significantly improve the performance of the heuristic over the cases where shortest path routing or random routing is used. This is a useful and striking result in the context of rings, where there are only two routing options.

In [22], the primary concern is wavelength assignment in a WDM ring and grooming is not considered; however, we consider this part of the area literature for the same reasons as [20]. A useful result is a proof of the NP-completeness of the problem of minimizing ADMs by transformation from an arc-colorability problem. The authors formulate the problem as a multicommodity flow ILP, and thereby obtain an approximation algorithm. A refinement of one of the heuristics of [20] is included. Finally, three greedy heuristics are provided to minimize ADMs by assigning wavelengths to a set of lightpath requests, and their approximation characteristics are given without proof. One of the heuristics is similar to the one proposed in [21].

The study in [23] explicitly distinguishes between the two cases of ring grooming when lightpath routing is or is not a part of the problem, and considers both. It follows the same approach of decomposing the problem into the formation of circles (called *primitive rings* here) and then grooming them as in [25]. The first phase of the problem then becomes similar to the wavelength assignment problem considered in [22] by the same authors. The heuristics of that study as well as others are then presented as primitive ring formation algorithms for the cases including and excluding routing, together with proofs of the corresponding approximation ratios. In the second phase, some theoretical results are provided regarding the number of groups (wavelengths) that are required to groom the primitive rings. A heuristic is provided for the general case, and it is shown that this can be done optimally for the special case $C = 2$. The authors remark on the NP-completeness of the subproblems they address in the two phases.

The approach of [19] is from a theoretical standpoint. The authors consider a family of graph problems that arises out of the second phase (grooming of primitive rings) of the problem considered in [23] and related works by the same authors. The formation of primitive rings is not considered part of the problem in this work. The authors effectively assume that SONET crossconnects are not used, and in addition make an assumption that SONET TDM traffic streams carried from one lightpath to another must occupy the same TDM slot. This would simplify implementation but constrains the solution significantly. The goal is to minimize the number of ADMs. The case of $C = 4$ is examined to illustrate the combinatorial formulation and approach. The design procedure can be likened to the canonical grooming methods used in [26] (discussed next), but the approach is that of decomposing the graph, not the traffic matrix, and the results have applicability to more general traffic scenarios. The case of $C = 8$ is then examined in an exhaustive manner. The algorithms can be thought of as grooming heuristics because it is not known that they can decompose every graph arising out of a traffic grooming problem, but some theoretical results are presented which indicate that the algorithms are good in many realistic scenarios.

Wavelengths May Exchange Traffic — The study in [26] also focuses on SONET rings. In addition, only the single-hub architecture is considered. The objective is to minimize SONET ADMs under the assumption that only the hub node has SONET crossconnect ability. This study introduces the interesting concept of "canonical grooming," that is, considering the static traffic matrix as being composed of a small set of traffic patterns and associated groomings. This is a form of decomposition applied to the traffic pattern. The study discusses such canonical grooming only with respect to the single hub architecture, but the concept may have wider applicability. Similar approaches of defining small tractable subsets of traffic have been considered by other authors [16, 25], but they do not make the concept of decomposition explicit. The authors go on to prove an NP-completeness result following the same approach as [16], and use the relationship of the grooming problem with the bin-packing problem to make a remark on the near-optimality of a grooming solution if an approximation scheme for the bin-packing problem is used. More specific results are presented for the special cases when the wavelength capacity $C = 2$, 4, and 8. The study also considers SONET rings of different speeds, similar to [17].

In [24], a heuristic based on a simulated annealing approach is applied to the problem. A WDM ring is considered, and the objective is to minimize the number of ADMs. The architectural assumptions, however, are somewhat restrictive. Two cases are considered: in one (*single-hop*), each traffic component has to be carried on just one lightpath, in a single logical hop from source to destination. In the other (*multihop*), each traffic component is terminated exactly once between source and destination, at the hub node (which is assumed to have full SONET crossconnect capabilities). ILP formulations are provided that can be extended to bidirectional or unidirectional rings. However, the formulations do not exactly implement the single-hop lightpaths; rather, a "single-hop" connection may be optically terminated at several intermediate nodes that are source or destination nodes for other traffic components sharing the same wavelength. This approach is similar to the "combine" ADMs approach of [27], discussed in the next subsection. Such sharing is realistic in terms of the number of wavelengths likely to be available or (conversely) bandwidth utilization. Since single-hop connections can be electronically routed at several SONET ADMs, the two architectures are equivalent to the assumption of no SONET crossconnects (the single-hop case) or full SONET crossconnection at one special hub node (the multihop case). Since these problems are NP-complete, a heuristic approach is proposed. The problem is again considered in two phases as in [25], and the heuristic proposed in [25] is used for the first (circle formation) phase. The emphasis in this work is on obtaining some benefit in the second (circle grooming) phase by using a simulated annealing heuristic rather than a greedy approach. The perturbation is to exchange circles or (if possible) partial circles between two wavelengths in a groomed solution. Numerical results show that simulated annealing outperforms the greedy heuristic.

We have studied the traffic grooming problem in WDM rings [11] under general traffic patterns and the second cost function. We have developed a novel dynamic programming approach for optimally placing a set of opaque nodes in the network. Having some of the nodes be opaque effectively *decomposes* the ring into a set of *independent path segments*. Each segment is solved optimally by solving the corresponding ILP exactly. The ILP for a path segment is identical to the ILP for the ring network, except that the constraints and variables related to wavelength assignment are removed. Removing these constraints/variables does not affect optimality in the path case, but reduces the computation time for a path segment by several orders of magnitude from that of a ring with the same number of nodes. The optimal path solutions are then appropriately combined into *lower* and *upper* bounds for the original ring network. As the length of path segments increases (i.e., the number of opaque nodes decreases), the quality of the bounds also increases. Our experiments have shown that the upper and lower bounds converge quickly within a tight range that depends on the traffic matrix, and the upper bounds represent feasible solutions to the traffic grooming problem for the original ring. As a result, we obtain provably near-optimal solutions to large ring networks that are impossible to solve directly. Since the problem solved for the path segments is identical to that for the ring network (except for the wavelength assignment variables), this approach to obtaining bounds is valid for any variant of the traffic grooming problem on rings with the first or second cost function.

Dynamic Grooming Model — In [27], the grooming problem is studied in the context of a WDM ring in which "circuits" (traffic components) are requested between node pairs, each circuit requiring $1/C$ of the bandwidth of a wavelength. The objective is to reduce the network cost by eliminating as many SONET ADMs from the "no grooming" scenario as possible; this is a variation of the first cost function we defined earlier. Additional assumptions are made that simplify the problem to some degree. Specifically, only symmetric traffic requirements are considered, that is, for any two nodes x and y, the traffic flowing from x to y is the same as the traffic from y to x. Also, no circuit is allowed to switch wavelengths at an intermediate node; that is, no SONET crossconnects are available at the ring nodes.

This work considers the "t-allowable" traffic matrix model. Instead of focusing on a specific traffic matrix as a single problem instance, an entire family of traffic matrices, having some unifying characteristics, is seen as a single problem instance. A traffic demand consisting of a static set of circuits is considered t-allowable if no node sources/sinks more than t duplex circuits. Several lower bounds for the number of ADMs that are needed to support all t-allowable traffic matrices are presented; they are derived by constructing "t-maximal" traffic matrices, which are effectively the hardest among all t-allowable problem instances. (The approach of considering a traffic family by defining allowable traffic matrices is similar to "Traffic Assumption A" of [18] which is discussed next, but in this case a single parameter is used to define the entire family, whereas in [18] the parameters are more related to capability limits of individual nodes.) The problem of designing good virtual topologies (deciding which ADMs may be removed) is converted to a matching problem in bipartite graphs. Virtual topologies that can support t-allowable matrices are characterized, and two different heuristic algorithms are proposed. Finally, it is discussed how some of the results must change if the circuits are dynamically added to the ring, and the ring must be wide-sense or strictly nonblocking.

In [18], three cost functions (number of wavelengths, number of ADMs, and maximum number of hops) are considered for each ring design proposed, but the authors remark that they consider the ADM cost most appropriate in reflecting the overall network cost. A distinction is made between static, dynamic, and incremental traffic; the last is a special case of the dynamic traffic where very long holding times are modeled by calls arriving but never departing. Three different traffic scenarios are considered. One is a static traffic model, with the uniform all-to-all pattern. The other two are dynamic scenarios. Grooming is made an integral part of the problem, in that some integer $C > 1$ is assumed to be the number of traffic components that can fit in a wavelength. Several ring designs are then proposed, each appropriate for some of the traffic scenarios considered. These designs can be viewed as heuristic solutions. For each design, algorithms are provided to obtain the virtual topology. Wavelength assignment and traffic routing are implicit. This study provides a very useful comparison of the different cost functions for the different ring designs proposed, and their suitability to different traffic scenarios. Blocking characteristics of the designs are also derived.

Traffic Grooming in Other Network Topologies

We have studied the traffic grooming problem in star networks [28] under the second cost model, and we believe it to be NP-hard. We have devised a strategy for obtaining a strong sequence of lower and upper bounds on the optimal solution, and we have shown that a greedy approach yields an approximation algorithm. In a star network only two types of lightpaths are possible: single-hop lightpaths originating or terminating at the hub node, and two-hop lightpaths originating and terminating at non-hub nodes. The problem has a polynomial time solution in some special cases, for instance, when only one non-hub node originates any traffic. In the general case, however, there are some indications that the traffic grooming problem is related to Knapsack (or, more generally, 0–1 integer programming) problems.

Traffic grooming in trees is NP-hard since part (ii), the RWA, is NP-hard. We studied this problem [28] and devised two strategies to obtain lower and upper bounds by decomposing the tree into star networks that are then optimally or approximately solved in isolation. For each internal tree node i define the star s_i as the part of the tree that includes node i, its parent node, and its children. In the first approach, all internal tree nodes at an even (or odd) depth from the root are assumed to be opaque. Then the star s_i of each nonopaque internal node i becomes independent of any other such star since all nodes of degree 1 are either leaves or opaque. Each star s_i is solved in isolation (either exactly or approximately), and the individual results are combined to obtain a feasible solution for the tree. Note that this solution is similar to the decomposition of rings into path segments separated by opaque nodes. In the second approach, we use translucent rather than opaque nodes. Specifically, all nodes i at an even (or odd) depth from the root are assumed to be opaque with respect to traffic to or from nodes outside the star s_i, but can be transparent with respect to traffic originating and terminating at non-hub nodes of s_i. In either approach, the solution to the tree problem contains no lightpaths longer than two hops. We have also developed two greedy heuristics that create longer lightpaths and have good average case performance.

The study in [29] is one of the few that expressly consider traffic grooming in an arbitrary network topology. A dynamic traffic scenario is considered, with the assumption of incremental traffic as in [18]. The individual traffic streams are of equal bandwidth and the wavelength capacity is an integral multiple of that bandwidth. However, the traffic routing problem is not considered part of the problem. The architecture proposed is one in which some nodes are opaque, and some are translucent. The assumption is also made that translucent nodes form an independent set in the graph of the physical topology; that is, no two translucent nodes can be neighbors. Because of this, the only lightpaths that can be formed are either one or two physical hops long. This assumption allows wavelength assignment to be treated in an implicit and integrated manner. Also, the number of wavelengths that can carry local traffic and must be terminated optically at a translucent node is constrained to be the same on each link incident on that node. Under these assumptions, it is shown that such a network is nonblocking in the wide sense.

As posed, the problem is one of placing the translucent nodes in the network with the above-described design. This reduces the amount of electronic routing that has to be performed, and also the number of LTE. This latter is the objective of the study, thus the goal is related to the first cost function. Placing nodes optimally is easily seen to be NP-complete. The key contribution in this work is an algorithm, called DROP, to decide which nodes of a network should be translucent. In essence, it is a gradient-descent algorithm, with a perturbation parameter. Starting from the empty set, new nodes are added to the candidate independent set if they tend to increase the benefit to the network. Results are shown for several values of the perturbation parameter, so a simulated annealing type of algorithm is followed. It is seen that after only a few iterations of the algorithm a solution close to the optimal is typically found.

In [30], the problem of minimizing LTE is considered for a general topology. The authors concentrate on the topology subproblem integrated with the routing subproblem. The RWA problem is neglected, since it is assumed that the network is large enough so that the RWA problem is solvable for any virtual topology. The authors make the interesting intuitive observation that minimizing endpoints for a given traffic should be related closely to maximizing traffic for a given amount of endpoint equipment. This notion is formalized into the key concept of a precise duality, and thus it is shown that the LTE minimization problem is equivalent to a commodity flow problem. However, the duality is precise only under the assumption that the RWA problem can always be solved and the ability to carry traffic is limited only by available network capacity (in other words, the grooming granularity has no effect). Based on this, a heuristic algorithm is presented that involves starting with an initial topology and successively deleting lightpaths after rerouting the traffic they carry.

An arbitrary topology is also considered in [31]. The grooming problem is formulated as an ILP. An interesting feature of this formulation is that the objective is defined not in terms of network cost but rather the amount of the traffic demands that are actually satisfied. In other words, solutions are not constrained to serve all the traffic offered to the network. Thus, problem instances that are infeasible for formulations constrained to serve all traffic become feasible for this one. This is a goal somewhat similar to those traditionally associated with dynamic traffic scenarios, although the static traffic model is used in this work. Another feature of this formulation is that the authors disallow the common *bifurcation assumption*, and require each traffic component to be routed as a whole. This is a nontrivial difference from other formulations since the authors allow for different traffic requirements on the network to be of different bandwidth. The authors provide two heuristic approaches based on their own earlier work; one attempts to maximize the amount of traffic that can be carried in one logical hop, and the other to maximize utilization of the lightpaths. The RWA of the lightpaths formed is carried out using established algorithms; specifically, adaptive routing and first-fit wavelength assignment are used.

Future Challenges

A significant amount of research has been carried out in the area of traffic grooming in WDM networks. However, most of the work has concentrated on areas of the most immediate practical interest (e.g., on ring networks). Various generalizations of the traffic grooming problem, some of which we discuss in this section, remain uninvestigated. Even for the specific case of the ring network, the various approaches are based on different sets of assumptions regarding the problem setting and cost functions, and are not directly comparable. While there is reason to believe we have reached the point where there is proper appreciation of the problem in all its guises, many open questions exist and much serious work remains to be performed. In this section we discuss a number of research directions that appear to us to be both important and challenging. Our purpose is not to offer an all-inclusive list of topics, but rather to generate interest in an area of practical importance for emerging MPλS networks.

Min-max Cost Function — To the best of our knowledge, the objective of traffic grooming research so far has focused on total network cost (first or second cost functions). We believe that a min-max objective whereby it is desirable to minimize the maximum cost at any node (third cost model) is of great practical interest. Define the *degree of opacity* of a translucent node as the number of wavelengths terminated (dropped) at the node. One approach is to devise an efficient algorithm for (near-)optimally placing translucent nodes of degree of opacity k in the ring. On the other hand, if we insist on optimal solutions without any completely opaque nodes, not all traffic patterns are allowable (since, e.g., a single node originating or terminating a large amount of traffic may need to be completely opaque). In this case, it is desirable to be able to characterize all the traffic patterns that are allowable when no node has a degree of opacity greater than k. Since traffic patterns will change over timescales much smaller than those at which the physical infrastructure of the network is expected to change, it would be useful to be able to define "families" of traffic patterns such that the same degree of opacity at each node can handle any traffic pattern in a given family. Similar work exists in the context of the RWA problem [8], but not for the traffic grooming problem.

Rings — While traffic grooming has been studied extensively in rings, most studies do not address the routing aspect of the problem. Therefore, there is a need to extend these studies by incorporating routing decisions, and to investigate what is achievable with simple routing strategies such as shortest path routing. Since carrying a traffic component over a larger number of links consumes more bandwidth and may increase the need for electronic switching, shortest path routing may not be far from optimal. Another unexplored issue is whether there is any gain by allowing part of a traffic component to be routed in one direction and the remaining part in the opposite direction.

Topologies with a Special Structure — Many network topologies encountered in practice have a special structure in that they are built by interconnecting a number of the elemental topologies such as rings, stars, and/or trees. Typically, such interconnection is deliberate because of either the infrastructure already in place or the need to exploit some desirable property of the elemental topology (e.g., protection switching in SONET rings). Examples include interconnected rings, rings of trees, and rings of stars. A straightforward way to address the traffic grooming problem in such networks is to decompose them into their constituent elemental topologies by placing opaque nodes at the points where these topologies attach to each other. It is possible to improve on this solution by considering translucent rather than opaque nodes at interconnection points. These nodes can be transparent with respect to traffic that stays within each of the topologies they interconnect, but opaque with respect to traffic traveling between the topologies. With the proliferation of topologies such as interconnected rings, this appears to be a fruitful area of research.

Significant Traffic Patterns — In many situations, the traffic matrix has a special form one may be able to exploit in order to develop better algorithms than are possible in the general case. Special traffic patterns (other than the rather unrealistic uniform all-to-all pattern) arise often in practice. Consider the tree structure of a cable TV network. When all traffic originates at the root (head-end) and the leaf nodes are the only destination nodes (e.g., when delivering digital TV signals), we have the special case of a *distribution tree*. If the cable TV network is also used for Internet access, we have a similar situation in which all traffic originates and terminates only at the root and leaf nodes. The tree network may also be used for leaf-to-leaf traffic, for example, to connect businesses accessing the network at a leaf node. In all these cases, the internal tree nodes do not originate or terminate any traffic but have an important role for grooming purposes. Other examples include situations where the traffic exhibits one (or a small number of) dominant pattern(s) (e.g., locality patterns or hot spots) [27]. In our own studies, we have found the most challenging instances to be those in which the traffic is more or less uniformly distributed in the network. Thus, we have reason to believe that good solutions tailored to specific patterns are likely to exist.

Virtual Topology Constraints — Constraints on the solutions rather than on the problem instances are also of practical interest. For example, virtual topologies that contain only up to k-hop lightpaths (where k is a small integer, say, 2–4) have been studied in the context of the RWA problem; it is obviously beneficial to extend these studies and characterize the quality of the solution in the presence of traffic grooming. Such topologies are of interest because so constraining the search space may reduce the benefit only marginally: under moderate to heavy traffic loads, long lightpaths are unlikely to be formed in optimal topologies because longer lightpaths are likely to lock away more bandwidth in return for a given grooming benefit. At the same time, the computational requirements may be drastically reduced due to the smaller state space. Furthermore, short lightpaths may be necessary due to physical layer constraints (e.g., loss, dispersion, or other optical layer impairments), while the resulting virtual topologies are simpler and easier to manage.

Constraints involving metrics other than lightpath length are also possible. For instance, some wavelengths may be set aside for traffic components that travel "short" distances; such components do not offer a great deal of grooming reward and are not groomed. Other wavelengths can be set aside to groom traffic components that have to be carried over "long" paths. One potential implementation of such an approach is to place opaque nodes that are separated by no more than, say, k hops; then "short" traffic components are routed directly to the destination, while "long" traffic components are locally groomed and routed to the nearest opaque node, where they undergo grooming and are aggregated into lightpaths to their destination. This approach may significantly reduce the size of the problem, at the cost of suboptimality of the solution.

Arbitrary Physical Topologies — While there has not been much work on traffic grooming in general topologies, the related virtual topology design problem [1] has been extensively studied over the last decade. However, most of the research in this area has been almost entirely experimental in nature, with little or no formal analysis of the heuristics developed. There is a simple reason the study of virtual topologies has taken this direction: *formally reasoning about virtual topology algorithms is extremely hard*.

More than heuristics, a formal systematic approach to the traffic grooming problem is needed whose performance can be characterized (e.g., by tight upper and lower bounds). The need is for graph-theoretic techniques that *have* been analytically verified, coupled with algorithms with provable properties for more elementary networks, to attack the grooming problem in networks of arbitrary topology. A worthy ultimate goal is to develop algorithms with formally verified properties that can be flexibly and efficiently applied within a variety of optical network and cost models.

Grooming of Multicast Traffic — Recently, the problem of carrying multicast traffic over optical networks has received significant attention. The main approach that has been considered is to provide multicast capabilities at the optical layer by employing power splitting devices at (some of) the OXCs in the network. The objective has been to construct multicast trees (or light-trees) that optically carry the multicast traffic from the source to the destination nodes. While optical multicasting is certainly an option for a small number of traffic streams that can utilize the whole capacity of a lightpath, it is reasonable to expect that there will be a significantly larger number of lower-rate multicast streams. Since it would be inefficient to assign a whole lightpath to each such stream, multicasting has to be performed electronically. Therefore, the problem of traffic grooming in the presence of multicast traffic in elemental as well as general network topologies arises, which, to the best of our knowledge, has not been studied so far.

Concluding Remarks

Traffic grooming in WDM networks is an important and exciting research area, and the future is likely only to see it become more essential. Over the past few years, a number of traffic grooming problems have been defined and investigated in several network contexts. An exciting variety of directions in attacking the problem more ambitiously and in more general contexts remain. We are certain that many significant results of practical importance are forthcoming.

References

[1] R. Dutta and G. N. Rouskas, "A Survey of Virtual Topology Design Algorithms for Wavelength Routed Optical Networks," *Op. Net.*, vol. 1, no. 1, Jan. 2000, pp. 73–89.
[2] U. Gupta, D. Lee, and J. Leung, "Efficient Algorithms for Interval Graphs and Circular-arc Graphs," *Networks*, vol. 12, 1982, pp. 459–67.
[3] D. West, *Introduction to Graph Theory*, Prentice Hall, 1996.
[4] G. Wilfong and P. Winkler, "Ring Routing and Wavelength Translation," *Proc. 9th Annual ACM-SIAM Symp. Discrete Algorithms*, 1998, pp. 333–41.
[5] S. Corteel *et al.*, "The Permutation-path Coloring Problem on Trees," to appear, *Theoretical Comp. Sci.*
[6] K. Kaklamanis *et al.*, Constrained Bipartite Edge Coloring with Applications to Wavelength Routing," *Proc. 24th Int'l. Colloq. Automata, Lang., and Prog.*, Bologna, Italy, July 1997.
[7] E. Harder, S.-K. Lee, and H.-A. Choi, "On Wavelength Assignment in WDM Optical Networks," *Proc. 4th Int'l. Conf. Massively Parallel Processing Using Opt. Interconnections*, Geneva, Switzerland, Apr. 1997.
[8] R. Ramaswami and G. H. Sasaki, "Multiwavelength Optical Networks with Limited Wavelength Conversion," *IEEE/ACM Trans. Net.*, vol. 6, no. 6, Dec. 1998, pp. 744–54.
[9] L. Margara and J. Simon, "Wavelength Assignment Problem on All-optical Networks with k Fibres per Link," *Proc. 27th Int'l. Colloq. Automata, Lang., and Progr.*, 2000.
[10] G. Li. and R. Simha, "On the Wavelength Assignment Problem in Multifiber WDM Star and Ring Networks," *IEEE/ACM Trans. Net.*, vol. 9, no. 1, Feb. 2001, pp. 60–68.
[11] R. Dutta and G. N. Rouskas, "On Optimal Traffic Grooming in WDM Rings," *IEEE JSAC*, vol. 20, no. 1, Jan. 2002, pp. 110–21.
[12] R. A. Patterson *et al.*, Eds., "Special Issue on Telecommunications Grooming," *Op. Net.*, vol. 2, no. 3, May/June 2001.
[13] J. Simmons, E. Goldstein, and A. Saleh, "Quantifying the Benefit of Wavelength Add-drop in WDM Rings with Distance-independent and Dependent Traffic," *IEEE/OSA J. Lightwave Tech.*, Jan 1999, pp. 48–57.
[14] J. Simmons and A. Saleh, "The Value of Optical Bypass in Reducing Router Size in Gigabit Networks," *Proc. ICC*, Jun 1999, pp. 591–96.
[15] X. Zhang and C. Qiao, "On Scheduling All-to-all Personalized Connections and Cost-effective Designs in WDM Rings," *IEEE/ACM Trans. Net.*, vol. 7, no. 3, Jun 1999, pp. 435–45.
[16] A. L. Chiu and E. Modiano, "Traffic Grooming Algorithms for Reducing Electronic Multiplexing Costs in WDM Ring Networks," *IEEE/OSA J. Lightwave Tech.*, vol. 19, no. 6, June 2001, pp. 2146–68.
[17] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET Network Design," *Proc. IEEE INFOCOM*, 1999, pp. 734–43.
[18] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost-effective Traffic Grooming in WDM Rings," *IEEE/ACM Trans. Net.*, vol. 8, no. 5, Oct. 2000, pp. 618–30.
[19] C. Colburn and A. C. H. Ling, "Wavelength Add-Drop Multiplexing for Minimizing SONET ADMs, to appear, *Discrete Mathematics*.
[20] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength Assignment in a WDM Ring to Minimize Cost of Embedded Sonet Rings," *IEEE INFOCOM*, 1998, pp. 94–101.
[21] G. Li. and R. Simha, "Efficient Routing to Reduce the Cost of Add-Drop Multiplexers in WDM Optical Ring Networks," *SPIE Conf.*, vol. 3843, 1999, pp. 258–67.
[22] L. Liu *et al.*, "Wavelength Assignment in WDM Rings to Minimize SONET ADMs," *Proc. IEEE INFOCOM*, 2000, pp. 1020–25.
[23] P-J. Wan *et al.*, "Grooming of Arbitrary Traffic in SONET/WDM BLSRs. *IEEE JSAC*, vol. 18, no. 10, 2000, pp. 1995–2003.
[24] J. Wang *et al.*, "Improved Approaches for Cost-effective Traffic Grooming in WDM Ring Networks: ILP Formulations and Single-hop and Multihop Connections," *IEEE/OSA J. Lightwave Tech.*, Nov 2001, pp. 1645–53.
[25] X. Zhang and C. Qiao, "An Effective and Comprehensive Approach to Traffic Grooming and Wavelength Assignment in SONET/WDM Rings," *IEEE/ACM Trans. Net.*, vol. 8, no. 5, Oct. 2000, pp. 608–17.
[26] X-Y. Li *et al.*, Practical Traffic Grooming Scheme for Single-hub SONET/WDM Rings," *Proc. IEEE LCN*, 2000, pp. 556–63.
[27] R. Berry and E. Modiano, "Reducing Electronic Multiplexing Costs in SONET/WDM Rings with Dynamically Changing Traffic," *IEEE JSAC*, vol. 18, no. 10, 2000, pp. 1961–71.
[28] R. Dutta, "Virtual Topology Design for Traffic Grooming in WDM Networks," Ph.D. thesis, North Carolina State Univ., Raleigh, NC, Aug. 2001.
[29] G. Sasaki and T. Lin, "A Minimal Cost WDM Network for Incremental Traffic," *Proc. SPIE*, vol. 3843, 1999.
[30] V. R. Konda and T. Y. Chow, "Algorithm for Traffic Grooming in Optical Networks to Minimize the Number of Transceivers," *IEEE Wksp. High Perf. Switching and Routing*, 2001, pp. 218–21.
[31] K. Zhu and B. Mukherjee, "Traffic Grooming in an Optical WDM Mesh Network," *IEEE JSAC*, vol. 20, no. 1, Jan 2002, pp. 122–33.

Biography

RUDRA DUTTA (dutta@csc.ncsu.edu) received his B.E. degree in electrical engineering from Jadavpur University, Calcutta, India, in 1991, an M.E. degree in systems science and automation from Indian Institute of Science, Bangalore, in 1993, and a Ph.D. degree in computer science from North Carolina State University, Raleigh, in 2001. From 1993 to 1997 he worked for IBM as a software developer and programmer in various networking-related projects. He is employed as assistant professor in the Department of Computer Science at North Carolina State University, Raleigh.

GEORGE N. ROUSKAS [SM] (rouskas@csc.ncsu.edu) received his diploma in electrical engineering from the National Technical University of Athens, Greece, in 1989, and M.S. and Ph.D. degrees in computer science from the College of Computing, Georgia Institute of Technology, Atlanta, in 1991 and 1994, respectively. He joined the Department of Computer Science, North Carolina State University, in August 1994 as an assistant professor, and has been a professor since July 2002. During the 2000–2001 academic year he spent a sabbatical term at Vitesse Semiconductor, Morrisville, North Carolina, and in May and June 2000 he was an invited professor at the University of Evry, France. He is a recipient of a 1997 NSF Faculty Early Career Development (CAREER) Award, and aco-author of a paper that received the Best Paper Award at the 1998 SPIE Conference on All-Optical Networking. He also received the 1995 Outstanding New Teacher Award from the Department of Computer Science, North Carolina State University, and the 1994 Graduate Research Assistant Award from the College of Computing, Georgia Tech. He was a co-guest editor for *IEEE Journal on Selected Areas in Communications*, Special Issue on Protocols and Architectures for Next Generation Optical WDM Networks, published in October 2000, and is on the editorial boards of *IEEE/ACM Transactions on Networking*, *Computer Networks*, and *Optical Networks*. He is a member of the ACM and the Technical Chamber of Greece.